Published in final edited form as:

Curr Opin Neurobiol. 2019 April ; 55: 188–198. doi:10.1016/j.conb.2019.04.001.

Convolutional nets for reconstructing neural circuits from brain images acquired by serial section electron microscopy

Kisuk Lee^{1,*}, Nicholas Turner^{2,*}, Thomas Macrina², Jingpeng Wu³, Ran Lu³, and H. Sebastian Seung^{2,3,**}

¹Department of Brain and Cognitive Sciences, MIT, Cambridge, MA 02139, USA

²Department of Computer Science, Princeton University, Princeton, NJ, 08544, USA

³Neuroscience Institute, Princeton University, Princeton, NJ, 08544, USA

Abstract

Neural circuits can be reconstructed from brain images acquired by serial section electron microscopy. Image analysis has been performed by manual labor for half a century, and efforts at automation date back almost as far. Convolutional nets were first applied to neuronal boundary detection a dozen years ago, and have now achieved impressive accuracy on clean images. Robust handling of image defects is a major outstanding challenge. Convolutional nets are also being employed for other tasks in neural circuit reconstruction: finding synapses and identifying synaptic partners, extending or pruning neuronal reconstructions, and aligning serial section images to create a 3D image stack. Computational systems are being engineered to handle petavoxel images of cubic millimeter brain volumes.

Keywords

Connectomics; Neural Circuit Reconstruction; Artificial Intelligence; Serial Section Electron Microscopy

Introduction

The reconstruction of the *C. elegans* nervous system by serial section electron microscopy (ssEM) required years of laborious manual image analysis [1]. Even recent ssEM reconstructions of neural circuits have required tens of thousands of hours of manual labor [2]. The dream of automating EM image analysis dates back to the dawn of computer vision in the 1960s and 70s (Marvin Minsky, personal communication) [3]. In the 2000s, connectomics was one of the first applications of convolutional nets to dense image prediction [4]. More recently, convolutional nets were finally accepted by mainstream

^{**}Corresponding author.

^{*}Equal Contribution

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

computer vision, and enhanced by huge investments in hardware and software for deep learning. It now seems likely that the dream of connectomes with minimal human labor will eventually be realized with the aid of artificial intelligence (AI).

More specific impacts of AI on the technologies of connectomics are harder to predict. One example is the ongoing duel between 3D EM imaging approaches, serial section and block face [5]. Images acquired by ssEM may contain many defects, such as damaged sections and misalignments, and axial resolution is poor. Block face EM (bfEM) was originally introduced to deal with these problems [6]. For its fly connectome project, Janelia Research Campus has invested heavily in FIB-SEM [7], a bfEM technique that delivers images with isotropic 8 nm voxels and few defects [8]. FIB-SEM quality is expected to boost the accuracy of image analysis by AI, thereby reducing costly manual image analysis by humans. Janelia has decided that this is overall the cheapest route to the fly connectome, even if FIB-SEM imaging is expensive (Gerry Rubin, personal communication).

It is possible that the entire field of connectomics will eventually switch to bfEM, following the lead of Janelia. But it is also possible that powerful AI plus lower quality ssEM images might be sufficient for delivering the accuracy that most neuroscientists need. The future of ssEM depends on this possibility.

The question of whether to invest in better data or better algorithms often arises in AI research and development. For example, most self-driving cars currently employ an expensive LIDAR sensor with primitive navigation algorithms, but cheap video cameras with more advanced AI may turn out to be sufficient in the long run [9].

AI is now highly accurate at analyzing ssEM images under “good conditions,” and continues to improve. But AI can fail catastrophically at image defects. This is currently its major shortcoming relative to human intelligence, and the toughest barrier to creating practical AI systems for ssEM images. The challenge of robustly handling edge cases is universal to building real-world AI systems, and makes the difference between a self-driving car for research and one that will be commercially successful.

In this journal, our lab previously reviewed the application of convolutional nets to connectomics [10]. The ideas were first put into practice in a system that applied a single convolutional net to convert manually traced neuronal skeletons into 3D reconstructions [11], and a semiautomated system (Eyewire) in which humans reconstructed neurons by interacting with a similar convolutional net [12].¹ A modern connectomics platform (Fig. 1) now contains multiple convolutional nets performing a variety of image analysis tasks.

This review will focus on ssEM image analysis. At present ssEM remains the most widely used technique for neural circuit reconstruction, judging from number of publications [5]. The question for the future is whether AI will be able to robustly handle the deficiencies of ssEM images. This will likely determine whether ssEM will remain popular, or be eclipsed by bfEM.

¹Other early semiautomated reconstruction systems did not use convolutional nets [13,14].

Serial section EM was originally done using transmission electron microscopy (TEM). More recently, serial sections have also been imaged by scanning electron microscopy (SEM) [15]. Serial section TEM and SEM are more similar to each other, and more different from bfEM techniques. Both ssEM methods produce images that can be reconstructed by very similar algorithms, and we will typically not distinguish between them in the following.

Alignment

The connectomic image processing system of Fig. 1 accepts as input a 3D image stack. This must be assembled from many raw ssEM image tiles. It is relatively easy to stitch or montage multiple image tiles to create a larger 2D image of a single section. Adjacent tiles contain the same content in the overlap region, and mismatch due to lens distortion can be corrected fairly easily [16]. More challenging is the alignment of 2D images from a series of sections to create the 3D image stack. Images from serial sections actually contain different content, and physical distortions of sections can cause severe mismatch.

The classic literature on image alignment has included two major approaches. One is to find corresponding points between image pairs, and then compute transformations of the images that bring corresponding points close together. In the iterative rerendering approach, one defines an objective function such as mean squared error or mutual information after alignment, and then iteratively searches for the parametrized image transformations that optimize the objective function. The latter approach is popular in human brain imaging [17]. It has not been popular for ssEM images, possibly because iterative rerendering is computationally costly.

The corresponding points approach has been taken by a number of ssEM software packages, including TrakEM2 [18], AlignTK [19], NCR Tools [20], FijiBento [21], and EM_aligner [22]. As an example of the state-of-the-art, it is helpful to examine a recent whole fly brain dataset [23], which is publicly available. The alignment appears outstanding at most locations. Large misalignments do occur rarely, however, and small misalignments can also be seen. This level of quality is sufficient for manual reconstruction, as humans are smart enough to compensate for misalignments. However, it poses challenges for automated reconstruction.

To improve alignment quality, several directions are being explored. One can add extensive human-in-the-loop capabilities to traditional computer vision algorithms, enabling a skilled operator to adjust parameters on a section-by-section basis, as well as detect and remove false correspondences [24]. Another approach is to reduce false correspondences from normalized cross correlation using weakly supervised learning of image encodings [25].

The iterative rerendering approach has been extended by Yoo et al. [26], who define an objective function based on the difference between image encodings rather than images. The encodings are generated by unsupervised learning of a convolutional autoencoder.

Another idea is to train a convolutional net to solve the alignment problem, with no optimization at all required at run-time. For example, one can train a convolutional net to take an image stack as input, and return an aligned image stack as output [27]. Alternatively,

one can train a convolutional net to take an image stack as input, and return deformations that align it [28]. Similar approaches are already standard in optical flow [29,30].

Boundary detection

Once images are aligned into a 3D image stack, the next step is *neuronal boundary detection* by a convolutional net (Fig. 1). In principle, if the resulting *boundary map* were perfectly accurate, then it would be trivial to obtain a segmentation of the image into neurons [4,10].

Detecting neuronal boundaries is challenging for several reasons. First, many other boundaries are visible inside neurons, due to intracellular organelles such as mitochondria and endoplasmic reticulum. Second, neuronal boundaries may fade out in locations, due to imperfect staining. Third, neuronal morphologies are highly complex. The densely packed, intertwined branches of neurons make for one of the most challenging biomedical image segmentation problems.

A dozen years ago, convolutional nets were already shown to outperform traditional image segmentation algorithms at neuronal boundary detection [4]. Since then advances in deep learning have dramatically improved boundary detection accuracy [31–34], as evidenced by two publicly available challenges on ssEM images, SNEMI3D² and CREMI.³

How do state-of-the-art boundary detectors differ from a dozen years ago? Jain et al. [4] used a net with six convolutional layers, eight feature maps per hidden layer, and 34,041 trainable parameters. Funke et al. [34] use a convolutional net containing pathways as long as 18 layers, as many as 1,500 feature maps per layer, and 83,998,872 trainable parameters.

State-of-the-art boundary detectors use the U-Net architecture [35,36] or variants [31,33,34,37,38]. The multiscale architecture of the U-Net is well-suited for handling both small and large neuronal objects, i.e., detecting boundaries of thin axons and spine necks, as well as thick dendrites. An example of automated reconstruction with a U-Net style architecture is shown in Fig. 2.

State-of-the-art boundary detectors make use of 3D convolutions, either exclusively or in combination with 2D convolutions. Already a dozen years ago, the first convolutional net applied to neuronal boundary detection used exclusively 3D convolutions [4]. However, this net was applied to bfEM images, which have roughly isotropic voxels and are usually well-aligned. It was popular to think that 3D convolutional nets were not appropriate for ssEM images, which had much poorer axial resolution and suffered from frequent misalignments. Furthermore, many deep learning frameworks did not support 3D convolutions very well or at all. Accordingly, much work on ssEM images has used 2D convolutional nets [39,40], relying on agglomeration techniques to link up superpixels in the axial direction. Today it has become commonly accepted that 3D convolutions are also useful for processing ssEM images.

²SNEMI3D challenge; URL: <http://brainiac2.mit.edu/SNEMI3D/>

³CREMI challenge; URL: <https://cremi.org/>

The current SNEMI3D [33] and CREMI leaders [32,34] both generate nearest neighbor *affinity graphs* as representations of neuronal boundaries. The affinity representation was introduced by Turaga et al. [41] as an alternative to classification of voxels as boundary or non-boundary. It is especially helpful for representing boundaries in spite of the poor axial resolution of ssEM images [10]. Parag et al. [42] have published empirical evidence that the affinity representation is helpful.

Turaga et al. [43] pointed out that using the Rand error as a loss function for boundary detection would properly penalize topological errors, and proposed the MALIS method for doing this. Funke et al. [34] have improved upon the original method through constrained MALIS training. Lee et al. [33] used long-range affinity prediction as an auxiliary objective during training, where prediction error for the affinities of a subset of long-range edges can be viewed as a crude approximation to the Rand error.

Handling of image defects

The top four entries on the SNEMI3D leaderboard have surpassed human accuracy as previously estimated by disagreement between two humans. But the claim of “superhuman” performance comes with many caveats [33]. The chief one is that the SNEMI3D dataset is relatively free of image defects. Robust handling of image defects is crucial for real-world accuracy, and is where human experts outshine AI.

One can imagine several ways of handling image defects: (1) Heal the defect by some image restoration computation. (2) Make the boundary detector more robust to the defect. (3) Correct the output of the boundary detector by subsequent processing steps. Below, we detail different types of image defects, their effects on reconstruction, and efforts to account for them.

Missing sections are not infrequent in ssEM images. Entire sections can be lost during cutting, collection, or imaging. The rate of loss varies across datasets, e.g., 0.17% in Zheng et al. [23] and 7.56% in Tobin et al. [44]. It is also common to lose part of a section. For example, one might exclude from imaging the regions of sections that are damaged or contaminated. One might also throw out image tiles after imaging, if they are inferior in quality. Or the imaging system might accidentally fail on certain tiles.

In an ssEM image stack, a partially missing section typically appears as an image with a black region where data is missing. An entirely missing section might be represented by an image that is all black, or might be omitted from the stack.

Traceability of neurites by human or machine is typically high if only a single section is missing. Traceability drops precipitously with long stretches of consecutive missing sections.⁴

⁴The longest stretch of missing sections was six in Tobin et al. [44] and eight in Lee et al. [2], which amounts to the loss of about 300 nm-thick tissue.

Partially and entirely missing sections are easy to simulate during training; one simply blacks out part or all of an image. Funke et al. [34] simulated entirely missing sections during training at a 5% rate. Lee et al. [33] simulated both partially and entirely missing sections at a higher rate, and found that convolutional nets can learn to “imagine” an accurate boundary map even with several consecutive missing sections.

Misalignments are not image defects, strictly speaking. They arise during image processing, after image acquisition. From the perspective of the boundary detector, however, misalignments appear as input defects. Progress in alignment software and algorithms is rapidly reducing the rate of misalignment errors. Nevertheless, misalignments can be the dominant cause of tracing errors, because boundary detection in the absence of image defects has become so accurate.

Lee et al. [33] introduced a novel type of data augmentation for simulating misalignments during training. Injecting simulated misalignments at a rate much higher than the real rate was shown to improve the robustness of boundary detection to misalignment errors.

Januszewski et al. [45] locally realigned image subvolumes prior to agglomeration, in an attempt to remove the image defect before further image processing.

Cracks and folds are common in ssEM sections. They may involve true loss of information, or cause misalignments in neighboring areas. We expect that improvements in software will be able to correct misalignments neighboring cracks or folds.

Knife marks are linear defects in ssEM images caused by imperfect cutting at one location on the knife blade. They can be seen in publicly available ssEM datasets [15].⁵ They are particularly harmful because they occur repeatedly at the same location in consecutive serial sections, and are difficult to simulate. Even human experts have difficulty tracing through knife marks.

Agglomeration

There are various approaches to generate a segmentation from the output of the boundary detector. The naïve approach of thresholding the boundary map and computing connected components can lead to many merge errors caused by “noisy” prediction of boundaries. Instead, it is common to first generate an oversegmentation into many small supervoxels. Watershed-type algorithms can be used for this step [46]. The number of watershed domains can be reduced by size-dependent clustering [46], seeded watershed combined with distance transformation [32,34], or machine learning [47].

Supervoxels are then agglomerated by various approaches. A classic computer vision approach is to use statistical properties of the boundary map [48], such as mean affinity [33] or percentiles of binned affinity [34]. A score function can be defined for every pair of contacting segments. At every step, the pair with the highest score is merged. This simple

⁵<https://neurodata.io/data/kasthuri15/>

procedure can yield surprisingly accurate segmentations when starting from high quality boundary maps, and can be made computationally efficient for large scale segmentation.

Agglomeration can also utilize other information not contained in the boundary map, such as features extracted from the input images or the segments [49–52]. Machine learning methods can also be used directly without defining underlying features to serve as the scoring function to be used in the agglomeration iterations [53].

Supervoxels can also be grouped into segments by optimization of global objective functions [32]. Success of this approach depends on designing a good objective function and algorithms for approximate solution of the typically NP-hard optimization problem.

Error detection and correction

Convolutional nets are also being used to automatically detect errors in neuronal segmentations [54–57]. Dmitriev et al. [56] leverage skeletonization of candidate segments, applying convolutional nets selectively to skeleton joints and endpoints to detect merge and split errors, respectively. Rolnick et al. [55] train a convolutional net to detect artificially induced merge errors. Zung et al. [57] demonstrate detection and localization of both split and merge errors with supervised learning of multiscale convolutional nets.

Convolutional nets have also been used to correct morphological reconstruction errors. Zung et al. [57] propose an error-correcting module which prunes an “advice” object mask constructed by aggregating erroneous objects found by an error-detecting module. The “object mask pruning” task (Fig. 3b) is an interesting counterpoint to the “object mask extension” task implemented by other methods (Fig. 3a) [45].

Synaptic relationships

To map a connectome, one must not only reconstruct neurons, but also determine the synaptic relationships between them. The annotation of synapses has traditionally been done manually, yet this is infeasible for larger volumes [58]. Most research on automation has focused on chemical synapses. This is because large volumes are typically imaged with a lateral resolution of 4 nm or worse, which is insufficient for visualizing electrical synapses. Higher resolution would increase both imaging time and dataset size.

A central operation in many approaches is the classification of each voxel as “synaptic” or “non-synaptic” [59–63]. Increasingly, convolutional nets are being used for this voxel classification task [58,64–69]. Synaptic clefts are then predicted by some hand-designed grouping of synaptic voxels.

A number of approaches are used to find the partner neurons at a synaptic cleft [58,64,66,67,70]. For example, Dorkenwald et al. [58] extract features relating the predicted cleft segments to candidate partners by overlap with these partners, as well as their contact site, and feed these features to a random forest classifier for a final “synaptic” or “non-synaptic” classification for inferring connectivity. Parag et al. [70] pass a candidate cleft, local image context, and a candidate pair of partner segments to a convolutional net to make

a similar “synaptic” or “non-synaptic” judgment. Turner *et al.* use a similar model, yet they instead use the cleft as an attentional input to predict the voxels of relevant presynaptic and postsynaptic partners (N Turner *et al.*, submitted).

Another group of recent approaches do not explicitly represent synaptic clefts [69,71]. Buhmann et al. [69] detect sites where presynaptic and postsynaptic terminals are separated by specific spatial intervals, using a single convolutional net to predict both location and directionality of synapses. Staffler et al. [71] extract contact sites for all adjacent pairs of segments within the region of interest, and use a decision tree trained on hand-designed features of the “subsegments” close to the contact site in order to infer synaptic contact sites and their directionality of connection.

Distributed chunk-wise processing of large images

In connectomics, it is often necessary to transform an input image into an output image (Fig. 1). The large input image is divided into overlapping chunks, each of which is small enough to fit into RAM of a single worker (CPU node or GPU). Each chunk is read from persistent storage by a worker, which processes the chunk and writes the result back to persistent storage. Additional operations may be required to insure consistency in the regions of overlap between the chunks. The processing of the chunks is distributed over many simultaneous workers for speedup.

To implement the above scheme, we need a “chunk workflow engine” for distributing chunk processing tasks over workers, and handling dependencies between tasks. We also need a “cutout service” that enables workers to “cut” chunks out of the large image. This includes both read and write operations, which may be requested by many workers in parallel. The engine and cutout service may function with a local cluster and networkattached storage, or with cloud computing and storage.

DVID is a cutout service used with both network-attached and Google Cloud Storage [72]. Requests are handled by a server running on a compute node in the local cluster or cloud. For ease of use, DICED is a Python wrapper that makes a DVID store behave as if it were a NumPy array [73]. bossDB uses microservices (AWS Lambda) to fulfill cutout requests from an image in Amazon S3 cloud storage [74]. ndstore shares common roots with bossDB [75], but uses Amazon EC2 servers rather than microservices. It is employed by neurodata.io, formerly known as the Open Connectome Project.

For cloud storage, the cutout service can also be implemented on the client side (Fig. 4). CloudVolume [76] and BigArrays.jl [77] convert cutout requests into native cloud storage requests. Both packages provide convenient slicing interfaces for Python NumPy arrays [76] and Julia arrays [77] respectively.

The chunk workflow engine handles task scheduling, the assigning of tasks to workers. In the simplest case, processing the large image can be decomposed into chunk processing tasks that do not depend on each other at all. Then scheduling can be handled by a task queue. Tasks are added to the queue in arbitrary order. When a worker is idle, it takes the first available task from the queue. The workers continue until the queue is empty. The

preceding can be implemented in the cloud using AWS Simple Queue Service [78,79]. Task scheduling is more complex when there are dependencies between tasks. If workflow dependencies are represented by a directed acyclic graph (DAG), the workflow can be executed by, for example, Apache Airflow [80].

In large scale distributed processing, errors are almost inevitable. The most common error is that a worker will fail to complete a task. Without proper handling of this failure mode, chunks will end up missing from the output image. A workflow engine typically judges a task to have failed if it is not completed within some fixed waiting time. The task is then reassigned to another worker. This “retry after timeout” strategy can be implemented in a variety of ways. In the case of a task queue, when one worker starts a task, the task is made invisible to other workers for a fixed waiting time. If the worker completes the task in time, it deletes the task from the queue. If the worker fails to complete the task in time, the task becomes visible in the queue again, and another worker can take it.

One can reduce cloud computing costs by using unstable instances (called “preemptible” on Google and “spot” on Amazon). These are much cheaper than on-demand instances, but can be killed at any time in favor of customers using on-demand instances or bidding with a higher price. When an instance is killed, its tasks will fail to complete. Therefore, the error handling described above is not only important for insuring correct results, but also for lowering cost.

Beyond boundary detection?

The idea of segmenting a nearest neighbor affinity graph generated by a convolutional net that detects neuronal boundaries is now standard (Fig. 1). One intriguingly different approach is the flood-filling network (FFN) of Januszewski et al. [45]. The convolutional net receives the EM image as one input, and a binary mask representing part of one object as an “attentional” input. The net is trained to extend the binary mask to cover the entire object within some spatial window (Fig. 3a). Iteration of the “flood filling” operation can be used to trace an object that is much larger than the spatial window. “Flood filling,” also called “object mask extension” in the computer vision literature, directly reconstructs an object. This is a beautiful simplification, because it eliminates postprocessing by a segmentation algorithm (e.g. connected components or watershed) as required by boundary detection.

That being said, whether FFNs deliver accuracy superior to state-of-the-art boundary detectors is unclear. On the FIB-25 dataset [81], FFNs [45] outperformed the U-Net boundary detector of Funke et al. [34], but this U-Net may no longer be state-of-the-art according to the CREMI leaderboard. In the SNEMI3D challenge, the accuracies of the FFN and boundary detector approaches are merely comparable [45].

Currently FFNs have high computational cost relative to boundary detectors, because every location in the image must be covered so many times by the FFN [45]. Nevertheless, one can imagine that FFNs could become much less costly in the future, riding the tailwinds of huge industrial investment in efficient convolutional net inference. Boundary detectors will at the

same time become less costly, but watershed-type postprocessing of boundary maps may not enjoy similar cost reduction.

Another intriguing direction is the application of deep metric learning to neuron reconstruction. As noted above, Lee et al. [33] have shown that prediction of nearest neighbor affinities can be made more accurate by also training the network to predict a small subset of long-range affinities. For predicting many or all long-range affinities, it is more efficient for a convolutional net to assign a feature vector to every voxel, and compute the affinity between two voxels as a function of their feature vectors. Voxels within the same object receive similar feature vectors and voxels from different objects receive dissimilar feature vectors. This idea was first explored in connectomics by Zung et al. [57] in an error detection and correction system.

Given the representation from deep metric learning, it is tempting to reconstruct an object by finding all feature vectors that are close in embedding space to a seed vector. This approach has been followed by the computer vision literature [82,83], effectively portraying deep metric learning as a “kill watershed” approach like FFNs. However, seed-based segmentation turns out to be outperformed by segmenting a nearest-neighbor affinity graph generated from the feature vectors [84]. Seed-based segmentation often fails on large objects, because the convolutional net cannot learn to generate sufficiently uniform feature vectors inside them.

Whether novel “object-centered” approaches like FFNs and deep metric learning can render boundary detectors obsolete remains to be seen. The showdown ensures suspense in the future of AI for connectomics.

Conclusion

The automation of ssEM image analysis is important because ssEM is still the dominant approach to reconstructing neural circuits [5]. It is also highly challenging for AI because ssEM image quality is poor relative to FIB-SEM. The field has made dramatic progress in increasing accuracy, but human effort is still required to correct the remaining errors of AI. Due to lack of space, we are unable to describe the considerable engineering effort currently going into semi-automated platforms that enable human proofreading of automated reconstructions. These platforms must adapt to leverage the increasing power of AI. As the error rate of automated reconstruction decreases, it may become more timeconsuming for the human expert to both detect and correct errors. To make efficient use of human effort, the user interface must be designed and implemented with care. Automated error detection may be useful for focusing human effort on particular locations [54], and automated error correction may be useful for suggesting interventions.

In spite of the gains from AI, the overall amount of human effort employed by connectomics may increase because the volumes being reconstructed are also increasing. Even a cubic millimeter has never been fully reconstructed yet, and a human brain is larger by six orders of magnitude. The demand for neural circuit reconstructions should also grow dramatically as they become less costly. Therefore AI may actually increase the total amount of human

effort devoted to connectomics, at least in the short term, bucking the conventional wisdom that AI is a jobs destroyer.

Automation of image analysis is spurring interest in increasing the throughput of image acquisition. When image analysis was the bottleneck, increasing image quality was perhaps more important. Now that the bottleneck is being relieved, there will be increased demand for high throughput SEM and TEM image acquisition systems [23,85]. With continued advances in both image acquisition and analysis, it seems likely that ssEM reconstruction of cubic millimeter volumes will eventually become routine. It is unclear, however, whether ssEM image acquisition can scale up to cubic centimeter volumes (whole mouse brain) or larger. New bfEM approaches are being proposed for acquiring such exascale datasets [86]. The battle between ssEM and bfEM continues.

References

1. White JG, Southgate E, Thomson JN, Brenner S: The structure of the nervous system of the nematode *Caenorhabditis elegans*. *Philos Trans R Soc Lond B Biol Sci* 1986, 314:1–340. [PubMed: 22462104]
2. Lee W-CA, Bonin V, Reed M, Graham BJ, Hood G, Glatfelter K, Reid RC: Anatomy and function of an excitatory network in the visual cortex. *Nature* 2016, 532:370–374. [PubMed: 27018655]
3. Sobel I, Levinthal C, Macagno E: Special techniques for the automatic computer reconstruction of neuronal structures. *Annu Rev Biophys Bioeng* 1980, 9:347–362. [PubMed: 7396437]
4. Jain V, Murray JF, Roth F, Turaga S, Zhigulin V, Briggman KL, Helmstaedter MN, Denk W, Seung HS: Supervised Learning of Image Restoration with Convolutional Networks. In 2007 IEEE 11th International Conference on Computer Vision. 2007:1–8.
5. Kornfeld J, Denk W: Progress and remaining challenges in high-throughput volume electron microscopy. *Curr Opin Neurobiol* 2018, 50:261–267. [PubMed: 29758457]
6. Denk W, Horstmann H: Serial block-face scanning electron microscopy to reconstruct three-dimensional tissue nanostructure. *PLoS Biol* 2004, 2:e329. [PubMed: 15514700]
7. Xu CS, Hayworth KJ, Lu Z, Grob P, Hassan AM, Garcia-Cerdan JG, Niyogi KK, Nogales E, Weinberg RJ, Hess HF: Enhanced FIB-SEM systems for large-volume 3D imaging. *Elife* 2017, 6:e25916. [PubMed: 28500755]
8. Knott G, Marchman H, Wall D, Lich B: Serial section scanning electron microscopy of adult brain tissue using focused ion beam milling. *J Neurosci* 2008, 28:2959–2964. [PubMed: 18353998]
9. Bojarski M, Del Testa D, Dworakowski D, Firner B, Flepp B, Goyal P, Jackel LD, Monfort M, Muller U, Zhang J, et al.: End to end learning for self-driving cars. *ArXiv Prepr ArXiv160407316* 2016,
10. Jain V, Seung HS, Turaga SC: Machines that learn to segment images: a crucial technology for connectomics. *Curr Opin Neurobiol* 2010, 20:653–666. [PubMed: 20801638]
11. Helmstaedter M, Briggman KL, Turaga SC, Jain V, Seung HS, Denk W: Connectomic reconstruction of the inner plexiform layer in the mouse retina. *Nature* 2013, 500:168–174. [PubMed: 23925239]
12. Kim JS, Greene MJ, Zlateski A, Lee K, Richardson M, Turaga SC, Purcaro M, Balkam M, Robinson A, Behabadi BF, et al.: Space-time wiring specificity supports direction selectivity in the retina. *Nature* 2014, 509:331–336. [PubMed: 24805243]
13. Takemura S, Bharioke A, Lu Z, Nern A, Vitaladevuni S, Rivlin PK, Katz WT, Olbris DJ, Plaza SM, Winston P, et al.: A visual motion detection circuit suggested by *Drosophila* connectomics. *Nature* 2013, 500:175. [PubMed: 23925240]
14. Mishchenko Y, Hu T, Spacek J, Mendenhall J, Harris KM, Chklovskii DB: Ultrastructural Analysis of Hippocampal Neuropil from the Connectomics Perspective. *Neuron* 2010, 67:1009–1020. [PubMed: 20869597]

15. Kasthuri N, Hayworth KJ, Berger DR, Schalek RL, Conchello JA, Knowles-Barley S, Lee D, Vázquez-Reina A, Kaynig V, Jones TR, et al.: Saturated Reconstruction of a Volume of Neocortex. *Cell* 2015, 162:648–661. [PubMed: 26232230]
16. Kaynig V, Vazquez-Reina A, Knowles-Barley S, Roberts M, Jones TR, Kasthuri N, Miller E, Lichtman J, Pfister H: Large-scale automatic reconstruction of neuronal processes from electron microscopy images. *Med Image Anal* 2015, 22:77–88. [PubMed: 25791436]
17. Avants BB, Tustison NJ, Song G, Cook PA, Klein A, Gee JC: A reproducible evaluation of ANTs similarity metric performance in brain image registration. *NeuroImage* 2011, 54:2033–2044. [PubMed: 20851191]
18. Saalfeld S, Fetter R, Cardona A, Tomancak P: Elastic volume reconstruction from series of ultra-thin microscopy sections. *Nat Methods* 2012, 9:717. [PubMed: 22688414]
19. Wetzel AW, Bakal J, Dittrich M, Hildebrand DG, Morgan JL, Lichtman JW: Registering large volume serial-section electron microscopy image sets for neural circuit reconstruction using FFT signal whitening. In *Applied Imagery Pattern Recognition Workshop (AIPR), 2016 IEEE*. . IEEE; 2016:1–10.
20. Anderson JR, Jones BW, Watt CB, Shaw MV, Yang J-H, DeMill D, Lauritzen JS, Lin Y, Rapp KD, Mastronarde D, et al.: Exploring the retinal connectome. *Mol Vis* 2011, 17:355–379. [PubMed: 21311605]
21. Joesch M, Mankus D, Yamagata M, Shahbazi A, Schalek R, Suissa-Peleg A, Meister M, Lichtman JW, Scheirer WJ, Sanes JR: Reconstruction of genetically identified neurons imaged by serial-section electron microscopy. *eLife* 2016, doi:10.7554/eLife.15015.
22. Khairy K, Denisov G, Saalfeld S: Joint Deformable Registration of Large EM Image Volumes: A Matrix Solver Approach. *ArXiv Prepr ArXiv180410019* 2018,
23. Zheng Z, Lauritzen JS, Perlman E, Robinson CG, Nichols M, Milkie D, Torrens O, Price J, Fisher CB, Sharifi N, et al.: A Complete Electron Microscopy Volume of the Brain of Adult *Drosophila melanogaster*. *Cell* 2018, 174:730–743.e22. [PubMed: 30033368]
24. Macrina T, Ih D: Alembic (ALignment of Electron Microscopy By Image Correlograms): A set of tools for elastic image registration in Julia 2018.
25. Buniatyan D, Macrina T, Ih D, Zung J, Seung HS: Deep Learning Improves Template Matching by Normalized Cross Correlation. *ArXiv Prepr ArXiv170508593* 2017,
26. Yoo I, Hildebrand DGC, Tobin WF, Lee W-CA, Jeong W-K: ssEMnet: Serial-Section Electron Microscopy Image Registration Using a Spatial Transformer Network with Learned Features. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support* Edited by Cardoso MJ, Arbel T, Carneiro G, Syeda-Mahmood T, Tavares JMRS, Moradi M, Bradley A, Greenspan H, Papa JP, Madabhushi A, et al. Springer International Publishing; 2017:249–257.* This paper describes a novel application of deep learning to ssEM image alignment. Images are encoded by a convolutional autoencoder. The encodings are used to search for good alignment transforms, with promising performance even near image defects.
27. Jain V: Adversarial Image Alignment and Interpolation. *ArXiv Prepr ArXiv170700067* 2017,
28. Mitchell E, Keselj S, Popovych S, Buniatyan D, Seung HS: Siamese Encoding and Alignment by Multiscale Learning with Self-Supervision. *ArXiv Prepr ArXiv190402643* 2019,
29. Dosovitskiy A, Fischer P, Ilg E, Häusser P, Hazirbas C, Golkov V, Smagt P v d, Cremers D, Brox T: FlowNet: Learning Optical Flow with Convolutional Networks. In *2015 IEEE International Conference on Computer Vision (ICCV)*. . 2015:2758–2766.
30. Ranjan A, Black MJ: Optical Flow Estimation Using a Spatial Pyramid Network. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. . 2017:2720–2729.
31. Zeng T, Wu B, Ji S: DeepEM3D: approaching human-level performance on 3 D anisotropic EM image segmentation. *BMC Bioinformatics* 2017, 33:2555–2562.
32. Beier T, Pape C, Rahaman N, Prange T, Berg S, Bock DD, Cardona A, Knott GW, Plaza SM, Scheffer LK, et al.: Multicut brings automated neurite segmentation closer to human performance. *Nat Methods* 2017, 14:101–102. [PubMed: 28139671] * This paper describes the multicut method which starts from an oversegmentation, assigns attractive and repulsive scores to all relevant superpixel pairs, and determines the final segmentation via global optimization of the scores.

33. Lee K, Zung J, Li P, Jain V, Seung HS: Superhuman Accuracy on the SNEMI3D Connectomics Challenge. ArXiv Prepr ArXiv170600120 2017,* This paper describes the first submission to the SNEMI3D challenge that surpassed the estimate of human accuracy. They train a variant of 3D U-Net to predict a nearest neighbor affinity graph as representation of neuronal boundaries, with two training tricks to improve accuracy: (1) simulation of image defects and (2) prediction of long-range affinities as an auxiliary objective.
34. Funke J, Tschopp FD, Grisaitis W, Sheridan A, Singh C, Saalfeld S, Turaga SC: Large Scale Image Segmentation with Structured Loss based Deep Learning for Connectome Reconstruction. IEEE Trans Pattern Anal Mach Intell 2018, doi:10.1109/TPAMI.2018.2835450.* This paper describes one of the top-ranked entries in the CREMI challenge. They train a 3D U-Net to predict a nearest neighbor affinity graph as representation of neuronal boundaries using a constrained version of MALIS, a structured loss to minimize topological errors by optimizing the Rand index. They also demonstrates the typical method of generating segmentation from the boundary map. An oversegmentation was generated first with seeded watershed. Agglomeration algorithms based on various statistical characters of affinities were then used to create the final segmentation.
35. Ronneberger O, Fischer P, Brox T: U-Net: Convolutional Networks for Biomedical Image Segmentation. In Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015 Edited by Navab N, Hornegger J, Wells WM, Frangi AF. Springer International Publishing; 2015:234–241.
36. Çiçek Ö, Abdulkadir A, Lienkamp SS, Brox T, Ronneberger O: 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation. In Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016 Edited by Ourselin S, Joskowicz L, Sabuncu MR, Unal G, Wells W. Springer International Publishing; 2016:424–432.
37. Quan TM, Hildebrand DGC, Jeong W-K: FusionNet: A deep fully residual convolutional neural network for image segmentation in connectomics. ArXiv Prepr ArXiv161205360 2016,
38. Fakhry A, Zeng T, Ji S: Residual Deconvolutional Networks for Brain Electron Microscopy Image Segmentation. IEEE Trans Med Imaging 2017, 36:447–456. [PubMed: 28113967]
39. Knowles-Barley S, Kaynig V, Jones TR, Wilson A, Morgan J, Lee D, Berger D, Kasthuri N, Lichtman JW, Pfister H: RhoanaNet Pipeline: Dense Automatic Neural Annotation. ArXiv Prepr ArXiv161106973 2016,
40. Meirovitch Y, Matveev A, Saribekyan H, Budden D, Rolnick D, Odor G, Knowles-Barley S, Jones TR, Pfister H, Lichtman JW, et al.: A Multi-Pass Approach to Large-Scale Connectomics. ArXiv Prepr ArXiv161202120 2016,
41. Turaga SC, Murray JF, Jain V, Roth F, Helmstaedter M, Briggman K, Denk W, Seung HS: Convolutional Networks Can Learn to Generate Affinity Graphs for Image Segmentation. Neural Comput 2010, 22:511–538. [PubMed: 19922289]
42. Parag T, Tschopp F, Grisaitis W, Turaga SC, Zhang X, Matejek B, Kamentsky L, Lichtman JW, Pfister H: Anisotropic EM Segmentation by 3D Affinity Learning and Agglomeration. ArXiv Prepr ArXiv170708935 2017,
43. Turaga SC, Briggman KL, Helmstaedter M, Denk W, Seung HS: Maximin Affinity Learning of Image Segmentation. In Proceedings of the 22Nd International Conference on Neural Information Processing Systems. . Curran Associates Inc.; 2009:1865–1873.
44. Tobin WF, Wilson RI, Lee W-CA: Wiring variations that enable and constrain neural computation in a sensory microcircuit. Elife 2017, 6:e24838. [PubMed: 28530904]
45. Januszewski M, Kornfeld J, Li PH, Pope A, Blakely T, Lindsey L, Maitin-Shepard J, Tyka M, Denk W, Jain V: High-precision automated reconstruction of neurons with flood-filling networks. Nat Methods 2018, 15:605–610. [PubMed: 30013046] ** This paper proposes an alternative framework to neuron segmentation that obviates the need for an explicit, intermediate representation of neuronal boundaries. They combine boundary detection and segmentation by attentive/iterative/recurrent application of a convolutional flood-filling module, and use the same module for agglomeration.
46. Zlateski A, Seung HS: Image Segmentation by Size-Dependent Single Linkage Clustering of a Watershed Basin Graph. ArXiv Prepr ArXiv150500249 2015,
47. Wolf S, Schott L, Köthe U, Hamprecht F: Learned Watershed: End-to-End Learning of Seeded Segmentation. In 2017 IEEE Int Conf Comput Vis (ICCV). . 2017:2030–2038.

48. Arbelaez P, Maire M, Fowlkes C, Malik J: Contour Detection and Hierarchical Image Segmentation. *IEEE Trans Pattern Anal Mach Intell* 2011, 33:898–916. [PubMed: 20733228]
49. Jain V, Turaga SC, Briggman K, Helmstaedter MN, Denk W, Seung HS: Learning to Agglomerate Superpixel Hierarchies. In *Adv Neural Inf Process Syst 24* Edited by Shawe-Taylor J, Zemel RS, Bartlett PL, Pereira F, Weinberger KQ. Curran Associates, Inc.; 2011:648–656.
50. Bogovic JA, Huang GB, Jain V: Learned versus Hand-Designed Feature Representations for 3d Agglomeration. *ArXiv Prepr ArXiv13126159* 2013,
51. Nunez-Iglesias J, Kennedy R, Parag T, Shi J, Chklovskii DB: Machine Learning of Hierarchical Clustering to Segment 2D and 3D Images. *PLoS One* 2013, 8:1–11.
52. Nunez-Iglesias J, Kennedy R, Plaza S, Chakraborty A, Katz W: Graph-based active learning of agglomeration (GALA): a Python library to segment 2D and 3 D neuroimages. *Front Neuroinform* 2014, 8:34. [PubMed: 24772079]
53. Maitin-Shepard JB, Jain V, Januszewski M, Li P, Abbeel P: Combinatorial Energy Learning for Image Segmentation. In *Adv Neural Inf Process Syst 29* Edited by Lee DD, Sugiyama M, Luxburg UV, Guyon I, Garnett R. Curran Associates, Inc.; 2016:1966–1974.
54. Haehn D, Kaynig V, Tompkin J, Lichtman JW, Pfister H: Guided Proofreading of Automatic Segmentations for Connectomics. *IEEE Conf Comput Vis Pattern Recognit CVPR 2018*,* This paper showcases the first interactive AI-guided proofreading system where AI detects morphological errors and guides human proofreaders to the errors, thereby reducing proofreading time and effort.
55. Rolnick D, Meirovitch Y, Parag T, Pfister H, Jain V, Lichtman JW, Boyden ES, Shavit N: Morphological error detection in 3d segmentations. *ArXiv Prepr ArXiv170510882* 2017,
56. Dmitriev K, Parag T, Matejek B, Kaufman A, Pfister H: Efficient Correction for EM Connectomics with Skeletal Representation. *Br Mach Vis Conf BMVC 2018*,* This work identifies errors based on the skeletonization of reconstructed segments. They apply one convolutional net at skeleton branch points to detect false merge errors, and a separate convolutional net at leaf nodes to detect false split errors. The output of the merge error detection network represents the cut plane which remedies the merger.
57. Zung J, Tartavull I, Lee K, Seung HS: An Error Detection and Correction Framework for Connectomics. In *Advances in Neural Information Processing Systems 30* Edited by Guyon I, Luxburg UV, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R. Curran Associates, Inc.; 2017:6818–6829.** This paper proposes a framework for both detecting and correcting morphological reconstruction errors based on the close interplay between the error-detecting and correcting modules. The error-correcting module “prunes” an “advice” object mask constructed by aggregating erroneous objects found by the error-detecting module. Both modules are based on supervised learning of 3D multiscale convolutional nets.
58. Dorkenwald S, Schubert PJ, Killinger MF, Urban G, Mikula S, Svara F, Kornfeld J: Automated synaptic connectivity inference for volume electron microscopy. *Nat Methods* 2017, 14:435–442. [PubMed: 28250467] * This paper proposes a new model for detection of synaptic clefts, and analyzes the performance of this model within the context of several other recent works. Their model also predicts several other intracellular objects such as mitochondria, which are used for cell type classification applied over a relatively large volume.
59. Kreshuk A, Straehle CN, Sommer C, Koethe U, Knott G, Hamprecht FA: Automated segmentation of synapses in 3D EM data. In *2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro. .* 2011:220–223.
60. Becker C, Ali K, Knott G, Fua P: Learning Context Cues for Synapse Segmentation. *IEEE Trans Med Imag* 2013, 32:1864–1877.
61. Kreshuk A, Koethe U, Pax E, Bock DD, Hamprecht FA: Automated Detection of Synapses in Serial Section Transmission Electron Microscopy Image Stacks. *PLOS ONE* 2014, 9:e87351. [PubMed: 24516550]
62. Jagadeesh V, Anderson J, Jones B, Marc R, Fisher S, Manjunath BS: Synapse classification and localization in Electron Micrographs. *Pattern Recognit Lett* 2014, 43:17–24.
63. Márquez Neila P, Baumela L, González-Soriano J, Rodríguez J-R, DeFelipe J, MerchánPérez Á: A Fast Method for the Segmentation of Synaptic Junctions and Mitochondria in Serial Electron Microscopic Images of the Brain. *Neuroinformatics* 2016, 14:235–250. [PubMed: 26780198]

64. Roncal WG, Kaynig-Fittkau V, Kasthuri N, Berger DR, Vogelstein JT, Fernandez LR, Lichtman JW, Vogelstein RJ, Pfister H, Hager GD: Volumetric Exploitation of Synaptic Information using Context Localization and Evaluation. *ArXiv Prepr ArXiv14033724* 2014,
65. Huang GB, Plaza S: Identifying Synapses Using Deep and Wide Multiscale Recursive Networks. *ArXiv Prepr ArXiv14091789* 2014,
66. Huang GB, Scheffer LK, Plaza SM: Fully-Automatic Synapse Prediction and Validation on a Large Data Set. *ArXiv160403075 Cs* 2016,
67. Santurkar S, Budden DM, Matveev A, Berlin H, Saribekyan H, Meirovitch Y, Shavit N: Toward Streaming Synapse Detection with Compositional ConvNets. *ArXiv Prepr ArXiv170207386* 2017,
68. Heinrich L, Funke J, Pape C, Nunez-Iglesias J, Saalfeld S: Synaptic Cleft Segmentation in Non-isotropic Volume Electron Microscopy of the Complete *Drosophila* Brain. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018* Edited by Frangi AF, Schnabel JA, Davatzikos C, Alberola-López C, Fichtinger G. Springer International Publishing; 2018:317–325.
69. Buhmann J, Krause R, Lentini RC, Eckstein N, Cook M, Turaga S, Funke J: Synaptic Partner Prediction from Point Annotations in Insect Brains. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018* Edited by Frangi AF, Schnabel JA, Davatzikos C, Alberola-López C, Fichtinger G. Springer International Publishing; 2018:309–316.
70. Parag T, Berger D, Kamensky L, Staffler B, Wei D, Helmstaedter M, Lichtman JW, Pfister H: Detecting Synapse Location and Connectivity by Signed Proximity Estimation and Pruning with Deep Nets. *ArXiv Prepr ArXiv180702739* 2018,
71. Staffler B, Berning M, Boergens KM, Gour A, Smagt P van der, Helmstaedter M: SynEM, automated synapse detection for connectomics. *Elife* 2017, 6:e26414. [PubMed: 28708060]
72. Katz WT, Plaza SM: DVID: Distributed Versioned Image-Oriented Dataservice. *Front Neural Circuits* 2019, 13:5. [PubMed: 30804760]
73. Plaza S: DICED: Interface that allows versioned access to cloud-backed nD datasets janelia-flyem; 2018.
74. Kleissas D, Hider R, Pryor D, Gion T, Manavalan P, Matelsky J, Baden A, Lillaney K, Burns R, D'Angelo D, et al.: The Block Object Storage Service (bossDB): A CloudNative Approach for Petascale Neuroscience Discovery. *bioRxiv* 2017,
75. Burns R, Roncal WG, Kleissas D, Lillaney K, Manavalan P, Perlman E, Berger DR, Bock DD, Chung K, Grosenick L, et al.: The Open Connectome Project Data Cluster: Scalable Analysis and Vision for High-Throughput Neuroscience. *Sci Stat Database Manag Int Conf SSDBM Proc Int Conf Sci Stat Database Manag* 2013, doi:10.1145/2484838.2484870.
76. Silversmith W: CloudVolume: Client for reading and writing to Neuroglancer Precomputed volumes on cloud services seung-lab; 2018.
77. Wu J: BigArrays.jl: storing and accessing large Julia array locally or in cloud storage seung-lab; 2018.
78. Wu J, Seung S: Chunkflow: Distributed hybrid cloud processing of large 3 D images by convolutional nets. *Comput Vis Microsc Image Anal CVMI Submitt* 2019,
79. Silversmith W: Igneous: Python pipeline for Neuroglancer compatible Downsampling, Mesh, Remapping, and more seung-lab; 2018.
80. Wong W: Seuron: A curated set of tools for managing distributed task workflows seunglab; 2018.
81. Takemura S, Xu CS, Lu Z, Rivlin PK, Parag T, Olbris DJ, Plaza S, Zhao T, Katz WT, Umayam L, et al.: Synaptic circuits and their variations within different columns in the visual system of *Drosophila*. *Proc Natl Acad Sci U S A* 2015, 112:13711–13716. [PubMed: 26483464]
82. Fathi A, Wojna Z, Rathod V, Wang P, Song HO, Guadarrama S, Murphy KP: Semantic Instance Segmentation via Deep Metric Learning. *ArXiv Prepr ArXiv170310277* 2017,
83. Brabandere BD, Neven D, Gool LV: Semantic Instance Segmentation with a Discriminative Loss Function. *ArXiv Prepr ArXiv170802551* 2017,
84. Luther K, Seung HS: Learning Metric Graphs for Neuron Segmentation In Electron Microscopy Images. *ArXiv Prepr ArXiv190200100 ISBI 2019* 2019,

85. Eberle AL, Mikula S, Schalek R, Lichtman J, Tate MLK, Zeidler D: High-resolution, high-throughput imaging with a multibeam scanning electron microscope. *J Microsc* 2015, 259:114–120. [PubMed: 25627873]
86. Hayworth KJ, Peale D, Januszewski M, Knott G, Lu Z, Xu CS, Hess HF: GCIB-SEM: A path to 10 nm isotropic imaging of cubic millimeter volumes. *bioRxiv* 2019, doi:10.1101/563239.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Figure 1.

Example of connectomic image processing system. Neuronal boundaries and synaptic clefts are detected by convolutional nets. The boundary map is oversegmented into supervoxels by a watershed-type algorithm. Supervoxels are agglomerated to produce segments using hand-designed heuristics or machine learning algorithms. Synaptic partners in the segmentation are assigned to synapses using convolutional nets or other machine learning algorithms. Error detection and correction is performed by convolutional nets or human proofreaders. Gray dashed lines indicate alternative paths taken by some methods.

Figure 2.

Automated reconstruction of pyramidal neuron in mouse visual cortex by a system like that of Fig. 1. The reconstruction is largely correct, though one can readily find split errors (red, 4) and merge errors (yellow, 6) in proofreading. 6 of these errors are related to image artifacts, and merge errors often join an axon to a dendritic spine. Cyan dots represent automated predictions of presynaptic sites. Predicted postsynaptic terminals were omitted for clarity. The $200 \times 100 \times 100 \text{ m}^3$ ssEM dataset was acquired by the Allen Institute for Brain Science (A. Bodor, A. Bleckert, D. Bumbarger, N. M. da Costa, C. Reid) after calcium imaging in vivo by the Baylor College of Medicine (E. Froudarakis, J. Reimer, Andreas S. Tolias). Neurons were reconstructed by Princeton University (D. Ih, C. Jordan, N. Kemnitz, K. Lee, R. Lu, T. Macrina, S. Popovych, W. Silversmith, I. Tartavull, N. Turner, W. Wong, J. Wu, J. Zung, and H. S. Seung).

Figure 3.

Illustration of (a) object mask extension [40,45] and (b) object mask pruning [57]. Both employ attentional mechanisms, focusing on one object at a time. Object mask extension takes as input a subset of the true object and adds missing parts, whereas object mask pruning takes as input a superset of the true object and subtracts excessive parts. Both tasks typically use the raw image (or some alternative representation) as an extra input to figure out the right answer, though object mask pruning may be less dependent on the raw image [57].

Figure 4.

Two approaches to chunk cutout services. (a) Server-side cutout. A server responds to a cutout request by loading the required image blocks (light blue), performing the cutout (blue), and sending the result back to a client. (b) Client-side cutout. A client requests the chunk blocks from cloud storage, and performs the cutout. A read operation is shown here; write operations are also supported by cutout services.