Available online at www.sciencedirect.com

ScienceDirect

Current Opinion in Neurobiology

Progress and remaining challenges in high-throughput volume electron microscopy

Jörgen Kornfeld and Winfried Denk

Recent advances in the effectiveness of the automatic extraction of neural circuits from volume electron microscopy data have made us more optimistic that the goal of reconstructing the nervous system of an entire adult mammal (or bird) brain can be achieved in the next decade. The progress on the data analysis side — based mostly on variants of convolutional neural networks — has been particularly impressive, but improvements in the quality and spatial extent of published VEM datasets are substantial. Methodologically, the combination of hot-knife sample partitioning and ion milling stands out as a conceptual advance while the multi-beam scanning electron microscope promises to remove the data-acquisition bottleneck.

Address
Max Planck Institute of Neurobiology, Am Klopferspitz 18,
82152 Planegg, Germany

Corresponding author: Kornfeld, Jörgen (kornfeld@neuro.mpg.de)

Current Opinion in Neurobiology 2018, **50**:261–267

This review comes from a themed issue on **Neurotechnologies**

Edited by **Anikeeva** and **Luo**

https://doi.org/10.1016/j.conb.2018.04.030

Introduction

Only 30 years after their development began in earnest [1], self-driving cars, controlled by sophisticated artificial intelligence (AI) algorithms, now cause serious accidents at a rate that roughly matches that seen for human drivers, at least according to a manufacturer's press release (Tesla; URL: https://www.tesla.com/blog/tragic-loss). Natural intelligence (NI) depends on the brain, with algorithms (and probably memory as well [2]), that are encoded to some, possibly overwhelming, degree in the connection pattern between neurons. Those connections are made by synapses but depend on neural wires (axons and dendrites) for their link to the cell soma. When trying to reconstruct neural wiring from volume electron microscopy (VEM) data, a computer using the best current AI algorithms [3••] makes errors at about the rate of a moderately motivated human. However, most human tracing errors are due to inattention and thus uncorrelated [4], which means that they can be corrected by redundant tracing. Different computer algorithms, on the other hand, tend to fail at the same places in the data set, where disambiguation typically requires context and high-level knowledge that currently only human experts can provide (but see [5]). For VEM data of sufficient quality (with respect to resolution and signal-to-noise ratio (SNR)) highly motivated and knowledgeable humans, such as late-term graduate students and post-docs, can consistently (across experts) correct almost all of those errors. We can, therefore, assume that such data contain all the information needed to extract the neural circuit diagram and we can expect that algorithms will eventually be able to reconstruct at a human expert's error rate.

Progress in the analysis of VEM data

For the reconstruction of axons and dendrites, VEM-based connectomics (the only approach we will discuss here, for others see [6–8]) has in the past almost exclusively relied on a hybrid approach: first, the computer creates a candidate segmentation with parameters set to ensure that at most a few of the generated segments (supervoxels) span (merge) more than one real neurite, but each neurite is still broken up (split) into many supervoxels [9–12]. This situation is called an over segmentation and human proofreading can be performed by simply inspecting each segment for locations where it should be combined with one or more of its neighbors [13–15]. A different approach to 'proof reading' is to generate center-line tracings (skeletons), one for each neurite [4,12]. In a separate step, each skeleton is then used to combine overlapping supervoxels into a volume model of the corresponding neurite [10,12]. Both approaches yield similar results and consume about the same amount of human time (at least about 0.5 h per mm of neurite) with a possible advantage for skeletonization as long as the supervoxels are small [12,16] (for yet another approach see [17•]). This is not too surprising since both approaches require that all locations in the VEM data set are viewed by one or more humans. Note that proofreading is orders of magnitude faster than the rate at which segmentation proceeded in the case of *C. elegans* [18], where neither acquisition nor analysis used computers.

To move beyond the complete-viewing limit one has to stop proofreading decisions that the computer made on the basis of clear evidence and instead concentrate on decisions that could easily have gone the other way.

Confidence information, which may have to be generated separately [19,20], is needed to steer human inspection (Figure 1a) to ambiguous locations [13]. This has, in some cases, reduced the total proofreading effort needed to reach a given reconstruction reliability by orders of magnitude (J Kornfeld *et al.*, unpublished observations).

In addition to an increased use of targeted proofreading, there has been a recent surge in the efficacy of machine segmentation itself. It is likely that this has to do with a renewed focus on the key step in segmentation: the detection of cell boundaries (CREMI challenge; URL: https://cremi.org/). Boundary locations can sometimes be identified simply by testing whether the voxel intensity is above or below a certain threshold (some early connectomics data sets even made use of special staining protocols that enhanced boundary contrast [10,21]). When that fails, one can, in addition, consider the staining pattern surrounding the voxel in question, a task well suited for machine learning based on convolutional neural networks (CNNs). Finally, one can ask whether making the currently considered voxel a boundary voxel makes the local boundary shape more plausible or less so. That is what any human annotator's visual system does after having seen some amount of data form neural tissue. Expert annotators, in addition, draw on neurobiological knowledge (typical neurite diameters, neurite continuity and straightness among them) to navigate difficult regions of neuropil. Similarly, algorithms that agglomerate supervoxels make decisions on whether the resulting shape is more or less plausible on a medium length-scale. The required shape priors can either be designed using explicit knowledge about neurite shapes [22] or trained using the shapes of actual neurites [19,23]. Rather than considering one decision after another, one can consider the combined effect of multiple merge decisions to select a globally optimal set of mergers [9]. Now, back to single-voxel classification.

Flood-filling networks (FFNs), which are currently leading in segmentation performance [3••,24], use an iterative voxel-classification process: first, using only the image intensities, an initial estimate for each voxel's probability of being part of the current neurite is generated. A feedback path makes the previous iteration's 'in'-probability part of the classifier input, automatically incorporating implicit shape priors into the primary voxel classification process.

Progress in VEM acquisition methods

Before we can start analyzing data we have to generate them. High-throughput VEM remains the state-of-the-art for resolving nanometer-sized structures, the development of sub-diffraction (super-resolution) light microscopy [25,26] and of methods that expand the tissue before imaging [27] notwithstanding. During the last decade, four VEM techniques were used for most connectomic data sets (reviewed in [28]). Two of those techniques are based on serial block-face electron microscopy (SBEM). SBEM using a diamond knife (DiK-SBEM) was originally introduced more than three decades ago [29], but became an effective tool about a decade and a half ago when it was shown that it can be used on nervous tissue [30]. Since then SBEM data have been used to reconstruct neural circuits on a large scale (Table 1). A variant of SBEM was introduced that employs a focused ion beam (FIB) for the periodic removal of material from the block face [31,32]. With FIB-SBEM the resolution is improved in all directions. Block-face increments as low as 2 nm have been demonstrated [33•] (for DiK-SBEM, the thinnest layer that can be reliably removed is about 20 nm under realistic imaging conditions [17•]). The usable lateral resolution is almost on par with what can be achieved with TEM [33•], because FIB-based material removal tolerates a much larger electron exposure than DiK-SBEM [30]. The one big drawback of FIB-SBEM is that the field-of-view (FOV) is limited to a few tens of microns along the axis of the ion beam [32].

One of the most exciting recent advances in VEM was, therefore, the demonstration that epoxy-embedded heavy metal-stained brain-tissue samples can be cut — with only minimal loss of material at the interface — into slabs between 10 and 30 μm thick using a lubricated and heated diamond knife [34••] (Figure 1b,c). By combining the hot-knife slab-cutting (HoKS?) technique with FIB-SBEM, a method (HoKS FIBSBEM?) is created that should have no limit on sample size (note that HoKS has not yet been demonstrated for mouse or bird whole-brain samples) and at the same time provides data that are of sufficient resolution and SNR to allow virtually error-free automatic neural circuit extraction [34••].

It would take thousands of years to acquire FIB-SBEM data for a whole mouse (bird) brain with a single single-beam microscope, but if one takes advantage of the fact that the brain would be cut into hundreds or thousands of slabs, which can be imaged in parallel, the wall-clock time required can be reduced by simply using multiple machines, limited mainly by the available budget.

Setting budget issues aside for the moment, we could ask: Has any of the four approaches to VEM discussed by Briggman and Bock in their 2012 paper [28] emerged as the winner and is on the verge of making other approaches obsolete? Posing the question in a different way: if one were to embark on a large-scale effort to reconstruct one whole mouse (bird) brain now, is it time to commit to one of these approaches?

Based only on its track records, serial-section TEM (SSTEM) should be the method of choice (Table 1). SSTEM was used for all currently available whole-brain datasets: *C. elegans* with 302 neurons [18], the adult fruit