



Feasibility of 3D Reconstruction of Neural Morphology Using Expansion Microscopy and Barcode-Guided Agglomeration

Young-Gyu Yoon^{1,2}, Peilun Dai^{2,3}, Jeremy Wohlwend^{1,2}, Jae-Byum Chang^{2,4}, Adam H. Marblestone^{2*} and Edward S. Boyden^{2,3,5,6*}

¹ Department of Electrical Engineering and Computer Science, MIT, Cambridge, MA, United States, ² MIT Media Lab, MIT, Cambridge, MA, United States, ³ Department of Brain and Cognitive Sciences, MIT, Cambridge, MA, United States, ⁴ Department of Biomedical Engineering, Sungkyunkwan University, Seoul, South Korea, ⁵ Department of Biological Engineering, MIT, Cambridge, MA, United States, ⁶ McGovern Institute, MIT, Cambridge, MA, United States

We here introduce and study the properties, via computer simulation, of a candidate automated approach to algorithmic reconstruction of dense neural morphology, based on simulated data of the kind that would be obtained via two emerging molecular technologies—expansion microscopy (ExM) and *in-situ* molecular barcoding. We utilize a convolutional neural network to detect neuronal boundaries from protein-tagged plasma membrane images obtained via ExM, as well as a subsequent supervoxel-merging pipeline guided by optical readout of information-rich, cell-specific nucleic acid barcodes. We attempt to use conservative imaging and labeling parameters, with the goal of establishing a baseline case that points to the potential feasibility of optical circuit reconstruction, leaving open the possibility of higher-performance labeling technologies and algorithms. We find that, even with these conservative assumptions, an all-optical approach to dense neural morphology reconstruction may be possible via the proposed algorithmic framework. Future work should explore both the design-space of chemical labels and barcodes, as well as algorithms, to ultimately enable routine, high-performance optical circuit reconstruction.

OPEN ACCESS

Edited by:

Si Wu,
Beijing Normal University, China

Reviewed by:

Kai Du,
Karolinska Institute (KI), Sweden
Louis Tao,
Peking University, China

*Correspondence:

Adam H. Marblestone
amarbles@mit.edu
Edward S. Boyden
esb@media.mit.edu

Keywords: neural morphology, 3-D reconstruction, expansion microscopy, RNA barcode, convolutional neural network, agglomeration

Received: 02 August 2017

Accepted: 05 October 2017

Published: 24 October 2017

Citation:

Yoon Y-G, Dai P, Wohlwend J, Chang J-B, Marblestone AH and Boyden ES (2017) Feasibility of 3D Reconstruction of Neural Morphology Using Expansion Microscopy and Barcode-Guided Agglomeration. *Front. Comput. Neurosci.* 11:97. doi: 10.3389/fncom.2017.00097

INTRODUCTION

In order to analyze neural circuitry in a systematic manner, three-dimensional reconstruction of neural morphology is essential. Optical microscopy has been a key technology for mapping brain circuits throughout the history of neuroscience (Linás, 2003; Wilt et al., 2009), but it has been limited to imaging and tracing only a small subset of neurons located within a given brain region (Kasthuri and Lichtman, 2007; Miyasaka et al., 2014; Robles et al., 2014), in large part because of the limited spatial resolution of conventional optical microscopy. While it is possible to image thin axons and dendritic spine necks smaller than the diffraction limit of optical microscopes (Dumitriu et al., 2011), this necessitates sparseness of the neurons labeled (i.e., two objects that are each smaller than the resolution, but farther than the resolution from one another, can be resolved). Since only a subset of the neurons within a region can be imaged and resolved under optical microscopy, only

partial maps of brain circuitry have been obtained with light microscopy to date. Adding another dimension—color—to the image, by molecularly giving a random color to each neuron (Livet et al., 2007; Cai et al., 2013), increases the number of neurons that can be resolved (Sümbül et al., 2016), but the number is still much lower than the number of neurons in even a modest sized brain circuit because the color becomes locally ambiguous due to optical blurring (and hence color mixing) and the stochastic nature of the distribution of fluorescent molecules within neurons.

Expansion microscopy (ExM) provides a new approach to nanoscale-resolution optical microscopy optimized for application to large 3D brain circuits (Chen F. et al., 2015): in contrast to earlier methods of magnification, which rely on lenses to optically magnify images of preserved brain circuits, ExM physically magnifies samples directly. ExM fuses two sets of ideas: work going back to the 1980s on embedding preserved biological specimens in hydrogels such as polyacrylamide for imaging purposes (Hausen and Dreyer, 1981; Germroth et al., 1995), and work going back to the 1970s on responsive polymers capable of vastly changing size in response to changes chemical environment (Tanaka et al., 1980). By synthesizing a swellable polyelectrolyte gel throughout a specimen, which binds to key biomolecules or labels of interest, then mechanically homogenizing the specimen and dialyzing it in water, we can—by design, and established with extensive validation—expand brain tissues $\sim 4.5\times$ in linear dimension (in our original protocol), with low distortion ($\sim 1\text{--}4\%$). Thus, molecules initially located within a diffraction-limited volume are physically separated to distances great enough to be resolved with conventional microscopes, resulting in an effective resolution equal to the diffraction limit divided by the expansion factor, or $\sim 300\text{ nm}/\sim 4.5 \approx \sim 60\text{--}70\text{ nm}$. We have more recently shown that expansion can be performed multiple times iteratively on the same sample, for $4.5^2 \approx 20\times$ linear expansion, or even potentially $4.5^3 \approx 100\times$ linear expansion, for two and three iterations respectively (Chang et al., 2017). For the case of $20\times$ expansion, we have validated the ability to resolve nanoscale structures at a resolution of $\sim 25\text{ nm}$ (Chang et al., 2017); the fundamental resolution limit of ExM, the polymer mesh size, may be far smaller, in the range of 2 nm (O'Connell and Brady, 1976; Kurenkov et al., 2001) for the polymers used—in the ball park of the resolution of electron microscopy. Indeed, in our study on $20\times$ expansion, we estimated that the error added by gelation and expansion was on the order of $\sim 5\text{--}10\text{ nm}$, suggesting that the fundamental resolution of ExM could be quite good if other contributors to error (e.g., the sizes of antibodies or other labels), are minimized in magnitude.

These properties of ExM may make it optimal for *in-situ* readout of nucleic acid “barcodes” that have previously been proposed (Zador et al., 2012; Marblestone et al., 2014) to allow unique identification of single neurons (Lu et al., 2011; Kebschull et al., 2016). A “barcode” is defined here as a sequence of RNA used to uniquely “tag” a particular cell with a physical indicator of its identity. The number of possible nucleic acid barcodes grows exponentially with RNA sequence length: there are 4^N possible sequences of length N . For example, a barcode of length

30 has a potential diversity of $4^{30} \sim 10^{18}$, a number that vastly outstrips the number of neurons in the mouse brain ($\sim 10^8$) or even those of the largest mammals. Thus with sufficient diversity, ensured by a sufficiently long barcode sequence length, each neuron in a brain will almost certainly express a unique tag, even when barcodes are generated, or delivered to neurons, at random. Such barcodes could be read out optically in an expanded ExM specimen through the use of fluorescent *in-situ* sequencing (FISSEQ) (Lee et al., 2014; Chen et al., 2017) or fluorescence *in-situ* hybridization (FISH) (Langer-Safer et al., 1982; Chen et al., 2016). The barcodes, when read out *in-situ*, appear as spatially localized fluorescent dots (occupying $\sim 400\text{ nm}$ size post-expansion, or $\sim 20\text{ nm}$ effective size compared to other biological structures under $20\times$ ExM), which blink out a digital string of colors that uniquely encodes the identity of the cell that contains them, during the repeated steps of *in-situ* sequencing.

In this paper, we approach the question of dense 3-D neural morphology reconstruction theoretically, in light of the large-volume 3D optical super-resolution capability enabled by ExM, combined with the molecular identification capabilities supported by optical approaches. We ask: what are the limits of an optical approach in the context of 3-D reconstruction of neural morphology with ExM, and how can we approach those limits through novel chemical tags and analytical protocols that take advantage of ExM, as well as specific algorithms? Since such questions must be answered in the context of the detailed structural properties of neural tissue, we base our analysis on a simulation of ExM imaging which uses electron microscopic volume reconstructions of neural tissue (**Figure 1A**) as a “ground truth”—our simulation, then, produces images which emulate those which would be generated by ExM in an optical confocal microscope (**Figure 1B**), on a neural volume with a precisely known structure previously obtained by electron microscopy and analyzed. This approach is illustrated in **Figure 1**.

In our current study, we take a conservative approach, both in our estimates of the performance of ExM and in our algorithmic design. Indeed, we aim to demonstrate the feasibility of a baseline optical reconstruction approach using already-demonstrated ExM performance in terms of resolution and noise, already-demonstrated neuronal labeling methods compatible with ExM, and straightforward algorithms from the electron microscopy connectomics field. Although the future will likely bring substantial improvements in the performance of all these technology components, our aim here is to demonstrate a baseline case that establishes the basic feasibility of a general class of optical approaches. Specifically, we simulate realistic, often pessimistic, conditions with important informational limitations, e.g., limited fluorophore density, nonzero background fluorescence, limited spatial precision of fluorophore targeting of membranes, and modest expansion factors (see **Table 1**). Such parameters have been roughly matched to the properties of an ExM protocol with $\sim 20\times$ linear expansion factor that we recently demonstrated experimentally (Chang et al., 2017). We also choose a conservative algorithmic approach, taking one that is similar to that used by the electron microscopy connectomics community, namely boundary




FIGURE 1 | Simulation of ExM volume reconstruction pipeline. **(A)** Ground truth input dataset obtained from fully reconstructed electron microscopic volume. **(B)** Simulated optical ExM image. **(C)** Boundary detection performed on simulated optical ExM image. **(D)** 3D neuron segmentation obtained from simulated optical ExM image.

TABLE 1 | Simulation parameters.

	Parameter	Value
Confocal microscopy	Magnification (M) of objective lens	40
	Numerical aperture (NA) of objective lens	1.15
	Physical pixel pitch	4.8 (μm)
	z step	120 (nm)
	Signal-to-noise ratio (read noise; at the brightest pixel)	50–100
	Signal-to-noise ratio (Poisson noise; at the brightest pixel)	7–12
Expansion	Expansion factor	20
Fluorophore	Density (membrane)	4,000–10,000 (puncta/ μm^2)
	Density (cytosol)	2,000–4,000 (puncta/ μm^3)
	Density (background)	1,000–2,000 (puncta/ μm^3)
	Density (barcodes)	5–300 (molecules/ μm^3)
	Localization accuracy [standard deviation (SD)]	20 (nm)
	Size of local clusters (SD)	1–48 (nm)

detection using a convolutional neural network (ConvNet) (Figure 1C). Subsequent steps of the segmentation pipeline, such as watershed oversegmentation and supervoxel merging, are also performed in a manner analogous to that used in electron microscopy connectomics. We characterized the error rate of this ConvNet-based image segmentation process on ExM data simulated from ground-truth electron microscopy data, and found that it was possible to obtain Rand scores and information theoretic scores above 0.9.

Taking advantage of the ability of ExM to support multiplexed molecular interrogation, we further studied whether this performance could be enhanced via the use of previously proposed nucleic acid sequencing approaches for cell identity readout based on cell-specific RNA barcodes (Zador et al., 2012; Marblestone et al., 2014; Peikon et al., 2014; Kechschull et al., 2016). We found that such a combined approach could allow higher-accuracy 3D reconstructions of neural volumes, with

both Rand scores and information theoretic scores above 0.95, by using cell identity information from barcodes to guide the merging of supervoxels derived from the boundary detector. We studied the performance of this integrated pipeline—which combines machine learning for boundary detection with *in-situ* nucleic acid barcode sequencing for supervoxel merging—as a function of barcode density, and found that barcodes could improve the accuracy of reconstructions with realistic numbers of RNA barcodes per cell, assuming that barcodes could be randomly and uniformly distributed throughout the cytosol (see section Discussion for details on barcode targeting).

In summary, in what follows, we first introduce the ExM simulation approach. We then study the performance of a reconstruction pipeline using simulated 20x ExM data with a protein-based plasma membrane label similar to the membrane-anchored BrainBow constructs used with ExM in Tillberg et al. (2016) and Chang et al. (2017). We then introduce an integrated reconstruction pipeline that incorporates both membrane morphology and barcode readout. The morphology information from the 20x ExM membrane label can be processed in a similar way to how electron microscopy images are processed for 3-D reconstruction of neural circuits (Liu et al., 2012; Nunez-Iglesias et al., 2013): machine learning is employed to obtain a boundary probability map (BPM) from a membrane image, and then the BPM is then processed to find the connected components. The BPM is first over-segmented with a conservative threshold to ensure that there is no false merging of different neurons, and then the resulting segments are merged to obtain the final segmentation. In addition, we propose to read out neuronal identity information via *in-situ* multiplexed readout of barcodes distributed randomly throughout the neuronal cytosol. The barcode identities are then used as an input to the merging step, and we describe how to use barcodes as a means to guide segment mergers. In the discussion, we argue that such a pipeline could extend 3-D reconstruction to much larger volumes by allowing error correction over short- and long-range distances and allow facile integration with *in-situ* readout of cellular and synaptic molecular markers. Thus, our theoretical and computational analysis may contribute to the formulation of novel, useful optical approaches to high-resolution brain mapping.

RESULTS

Simulation of Expansion Microscopy Images

Simulation of ExM images was performed by first simulating stochastic fluorescent labeling of a volume of neural tissue and then projecting the tissue onto a 3-D image stack using an optical model of a confocal microscope. We used electron microscopic volume reconstructions of mouse cortex from (Berger et al., 2013; Kasthuri et al., 2015) as an input for the simulation.

Figure 2 shows an input slice and the corresponding output slice from our simulation. All neuronal membranes were labeled with red fluorophores, with varying labeling density across neurons. As a result, some of the neurons had relatively clear boundaries whereas the boundaries of other neurons were less clear. In addition, due to the discrete nature of fluorophores, they appeared at this magnification (see Chang et al., 2017 for actual data) punctate rather than as continuous contours (**Figure 2B**). Interestingly, we found that the discernibility of neurons with very low labeling density was improved when their membranes touched the membranes of neighboring neurons.

We include a real 20x ExM image side by side with a simulated ExM image to illustrate the realism of our simulation. **Figure 3A** shows z-slices from a confocal image stack of immunostained Emx1-Cre mouse hippocampus with neurons expressing membrane-bound fluorescent proteins (Brainbow 3.0; Cai et al., 2013) obtained with the 20x ExM protocol and **Figure 3B** shows z-slices from an image stack simulated from electron microscopy data. The average color for each neuron was randomly selected, and then red, green, and blue fluorophores were randomly assigned to the voxels in a mutually exclusive fashion, to simulate BrainBow.

Segmentation

Boundary Detection

We implemented and used a 3-D ConvNet for boundary detection by adopting and modifying the neural network design from VD2D3D (Lee et al., 2015); our network architecture is shown in **Figure 4**. A BPM was obtained as the output of the

ConvNet where each pixel in the BPM represents the probability of the pixel being on the boundary of a neuron (Ciresan et al., 2012). As the boundary at certain locations was not clear from a single z-slice, due to the limited number of fluorophores and resolution, using a 3-D ConvNet architecture was important to provide volumetric context to the network. It should be noted that seven z-slices were fed into the 3-D ConvNet instead of just a single slice, as shown in the figure; also, every third slice was picked instead of picking consecutive slices. By sampling every third slice, the amount of redundant information in the slices was reduced and the volumetric context from a larger volume could be provided to the ConvNet. A few input slices to the ConvNet and their corresponding outputs are shown in **Figure 5**.

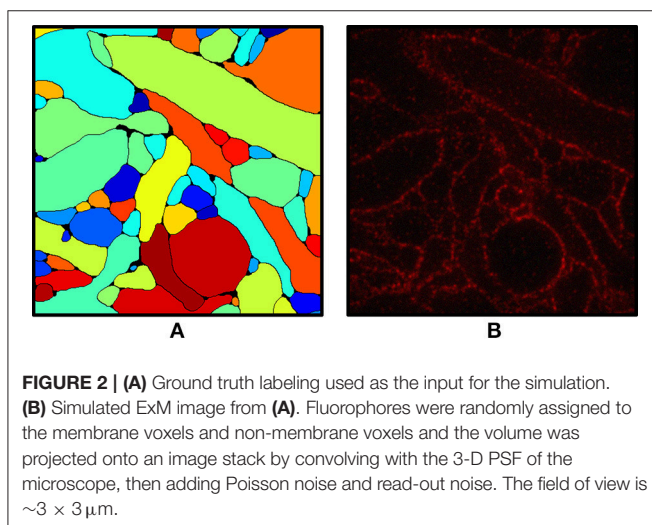
Over-Segmentation and Agglomeration

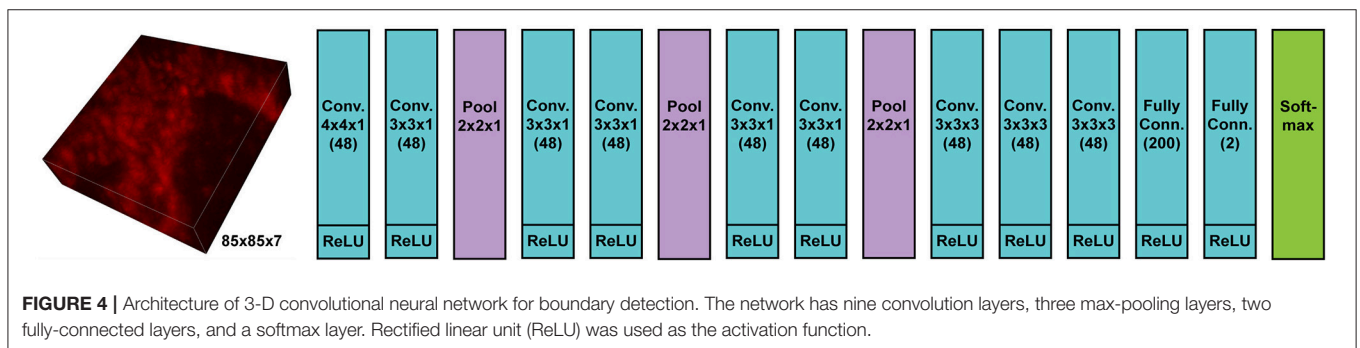
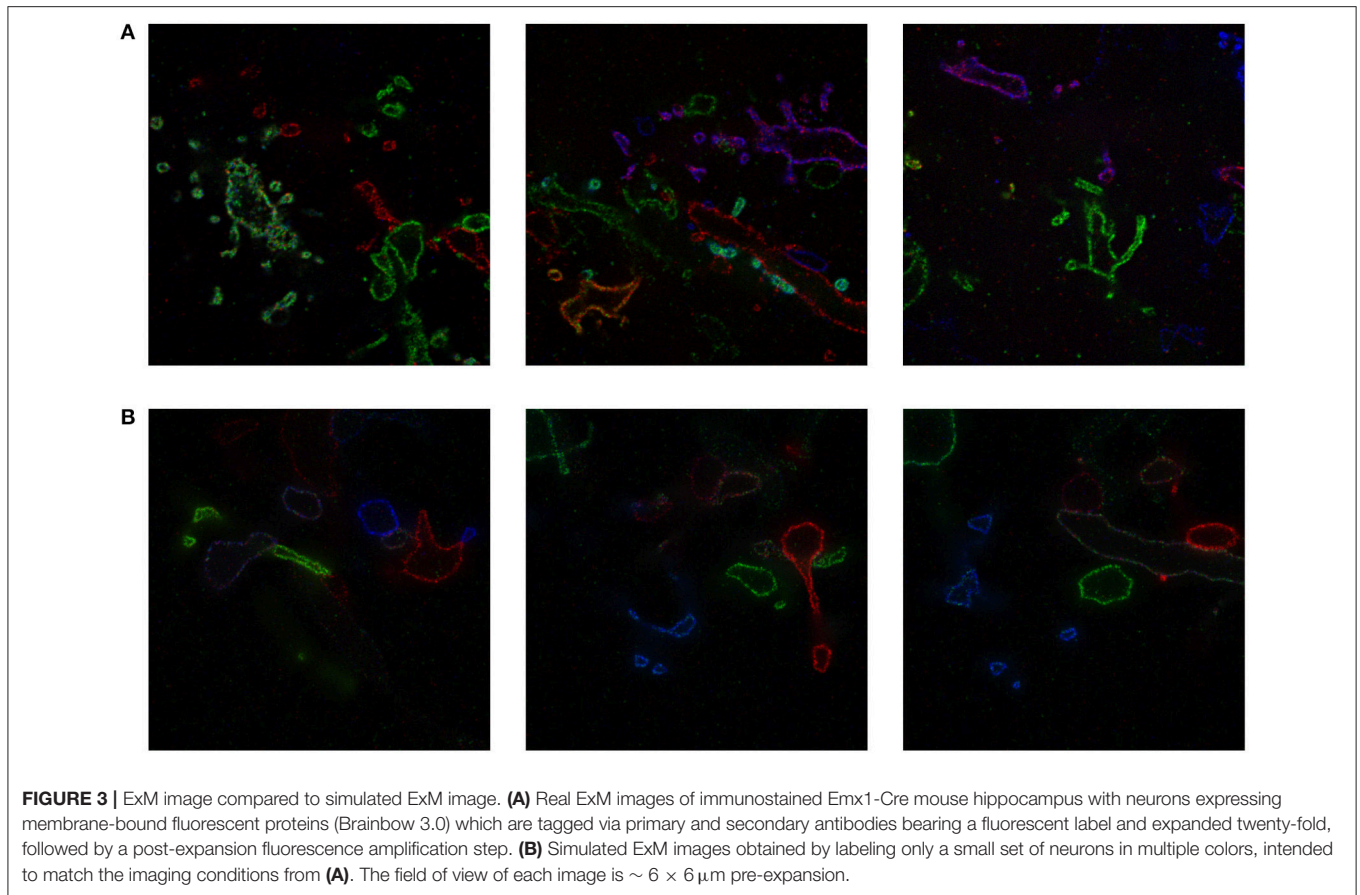
We first performed over-segmentation of the volume by applying a watershed transform to the BPM from the ConvNet and then the resulting segments were merged using the graph-based active learning of agglomeration (GALA) algorithm (Nunez-Iglesias et al., 2013, 2014), an agglomeration algorithm based on supervised learning of a policy that determines whether each pair of adjacent supervoxels should be merged or not. A hierarchical segmentation tree was obtained by applying GALA to the test image with multiple threshold values (**Figure 6A**), where a higher threshold value provided a result with more merges. This tree can be interpreted as a pool of segmentations in which different neurons are better reconstructed in different levels, i.e., there is no single threshold value that is ideal for all neurons. In other words, each node in the tree serves as a segmentation hypothesis and this hypothesis is either approved or rejected by the barcode consistency within the node (i.e., a node is inconsistent if more than one distinct barcode is found in a single node). Finally, nodes with the same barcode were merged. This procedure is illustrated in **Figure 6**. It should be noted that this algorithm allowed the agglomeration of segments that were disconnected by mistake, or not even connected in the field of view. This is important for large-scale reconstruction, as it shows how barcodes can prevent “propagation” of local errors and can help support the correction of errors without human intervention.

Intermediate and final segmentation results from our segmentation pipeline are shown in **Figure 7** along with the input image and the ground truth; 3-D surface renderings of the ground truth and the final segmentation are shown in **Figure 8**. An H-minima transform, followed by a watershed transform, yielded a good initial over-segmentation for subsequent agglomeration steps in the sense that both under-segmentation and severe over-segmentation were avoided. Most over-segmentation occurred in thin neurites running parallel to the z-planes, where the ConvNet failed to provide clear boundaries due to the anisotropic resolution ($\sim 200 \times 200 \times 600$ nm) of optical confocal microscopy that we assumed in our simulations, although this resolution is better for more modern microscope architectures (see section Discussion).

Quantification of Segmentation Accuracy

For quantitative assessment of segmentation quality, we used F-scores of the Rand index and variation of information (VI) (Arganda-Carreras et al., 2015), which are measures of the





similarity between two segmentations. The Rand split score and Rand merge score are defined as follows, where n_{ij} denotes the number of joint voxels of the i^{th} segment of the predicted segmentation X and the j^{th} segment of the ground truth segmentation Y , and the Rand F-score is the harmonic mean of the Rand split score and the Rand merge score:


$$V_{split}^{Rand} = \frac{\sum_{ij} n_{ij}^2}{\sum_i (\sum_j n_{ij})^2}, \quad V_{merge}^{Rand} = \frac{\sum_{ij} n_{ij}^2}{\sum_j (\sum_i n_{ij})^2}. \quad (1)$$

Similarly, the information theoretic split score and the information theoretic merge score are defined as follows

where $I(X,Y)$, $H(X)$, and $H(Y)$ denote the mutual information between X and Y , the entropy of X , and the entropy of Y , respectively, and their harmonic mean is the information theoretic score:

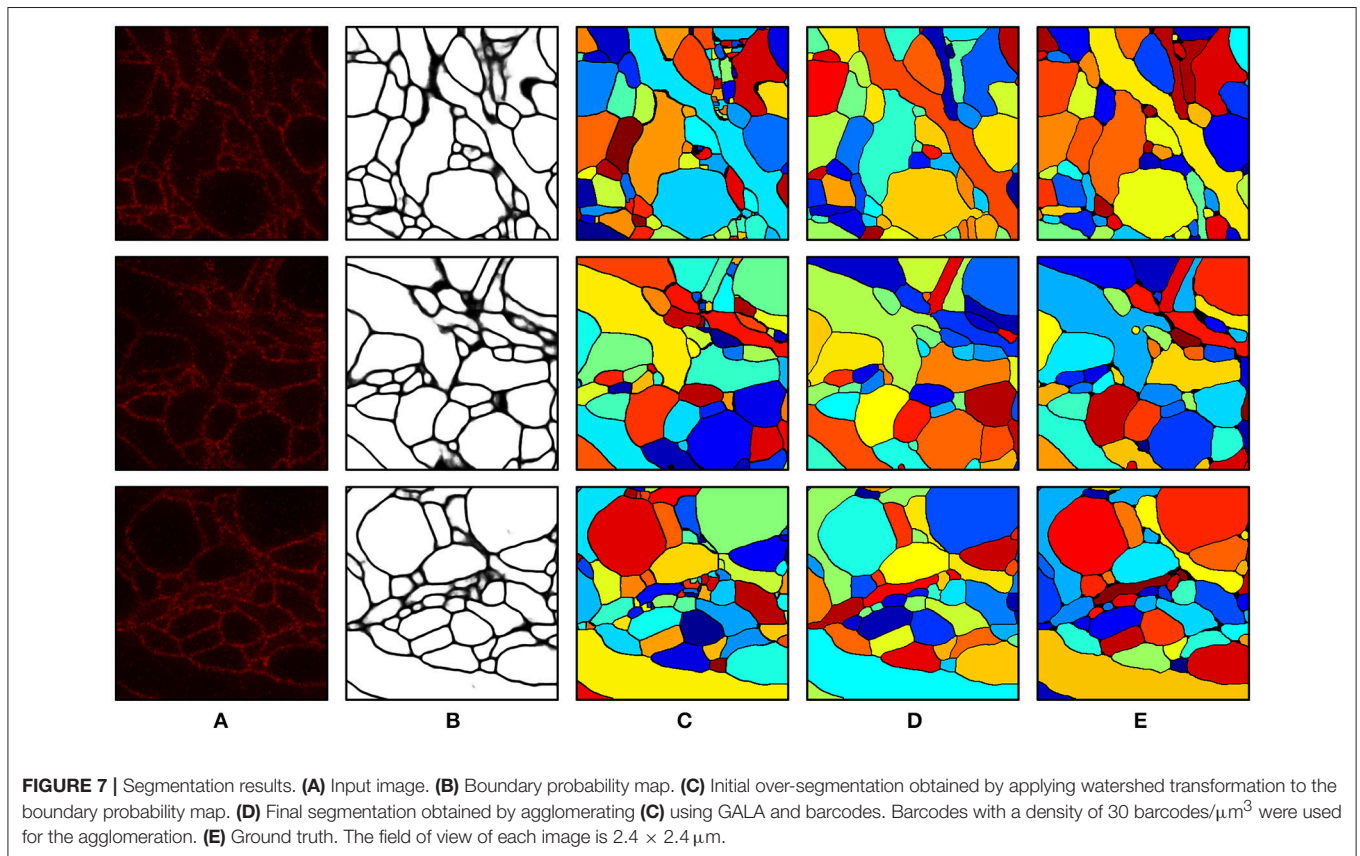
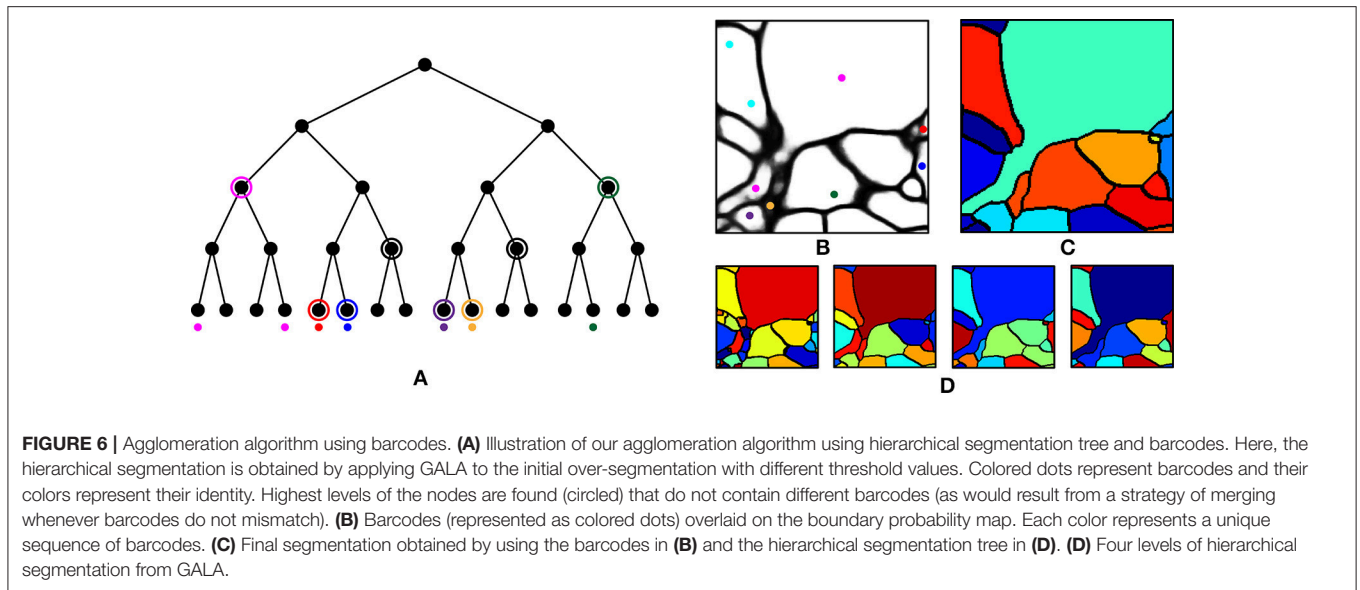
$$V_{split}^{VI} = \frac{I(X,Y)}{H(X)}, \quad V_{merge}^{VI} = \frac{I(X,Y)}{H(Y)}. \quad (2)$$

All scores are bounded between zero and one, by construction, with higher scores indicating more accurate segmentation. Foreground-restricted F-scores (i.e., F-scores computed only for voxels that are non-boundary in the ground truth data) were evaluated, as they were less sensitive to small border variations



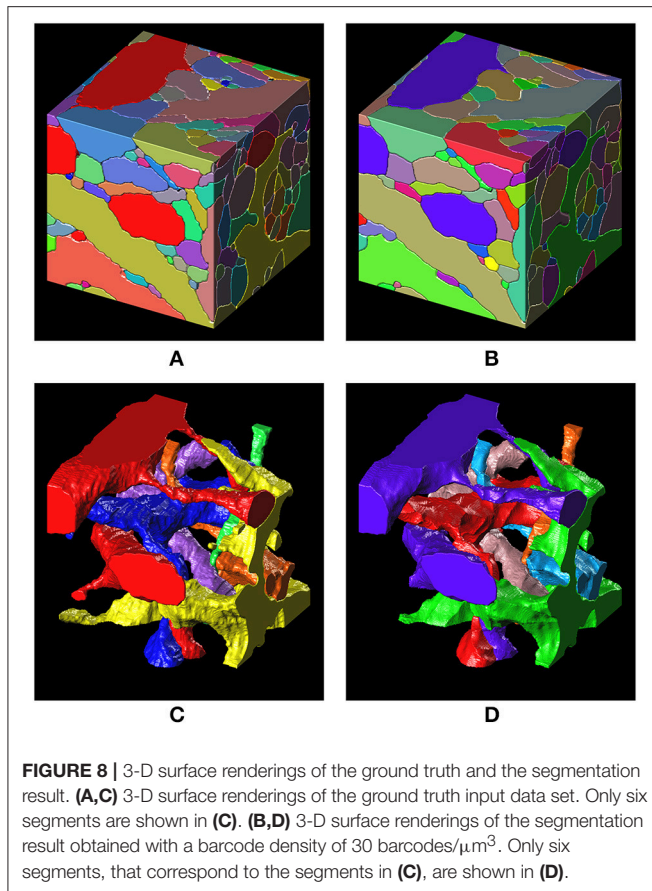
(Arganda-Carreras et al., 2015); the results are summarized in **Figure 9**. To quantify the segmentation quality achievable by the watershed transformation, the precision-recall curve, or the split score vs. merge score curve, was obtained by applying the H-minima transform with various threshold values as in Arganda-Carreras et al. (2015); Lee et al. (2015), and Nunez-Iglesias et al. (2014). To create a precision-recall curve for segmentation quality when barcodes were not used, each level of the hierarchical segmentation tree, which was obtained using GALA with a specific threshold value, was compared to the ground truth to measure its score. From **Figures 9A,B**, it can be seen that the split scores were significantly improved at each step of the pipeline (e.g., as barcodes are added), without compromising the merge score. The segmentation quality as

a function of the barcode density was also measured and the results are shown in **Figures 9C,D**. For each trial, we generated a new set of randomly distributed barcodes with a given density for the otherwise identical data set and used it for the agglomeration step. The result was different for each trial even with the same density of barcodes, as in some cases barcodes were located at critical locations and corrected major errors, whereas in some cases they were not. With the higher barcode density, it was more likely to have barcodes in such critical locations, and thus, scores tended to increase as the barcode density was increased. The F-scores did not asymptotically reach unity as the barcode density was increased, as the quality of the initial over-segmentation imposed an upper bound on the accuracy of the agglomeration result. To solve this problem,



future algorithms could explore methods to incorporate barcode information earlier in the segmentation pipeline, e.g., during the initial generation of the over-segmentation, or to use improved tags, e.g., with higher localization accuracy, to improve the quality of the initial BPM and over-segmentation (see section Discussion for details). Concretely, barcodes with pre-expansion

densities of 5–300 barcodes/ μm^3 (the density is $\sim 8,000$ -fold lower following 20x expansion) led to a significant improvement in segmentation quality scores (**Figures 9A,B**). These are high densities, but it should be noted that RNA readout with a density of 20 molecules/ μm^3 has been previously demonstrated using multiplexed RNA fluorescent *in-situ* hybridization (FISH)



(Chen K. H. et al., 2015), even without the benefit of ExM. We expect that these kinds of barcode densities could be achievable by combining highly expressed RNA barcodes with an efficient readout step, e.g., leveraging the high efficiencies of single-molecule RNA FISH, although the details of the subcellular targeting and distribution of barcodes will be a key question for most barcoding approaches, e.g., whether barcodes are excluded from certain regions under a given targeting scheme. Moreover, other schemes for incorporating barcode information may place less stringent requirements on barcode density (see Discussion for details).

DISCUSSION

We have shown via simulations that, using demonstrated methods of 20x physical expansion, methods of labeling the neuronal cell membrane with fluorescent dyes, and standard confocal microscopes, it would be possible to use convolutional neural networks for boundary detection, and graph-based super-voxel agglomeration, to obtain dense 3-D reconstruction of neural morphologies in brain circuits. Our work thus opens up the possibility of an optical approach to dense reconstruction of neural circuitry, using a combination of demonstrated methodologies. As discussed in the introduction, an optical approach could be valuable for brain mapping by allowing dense

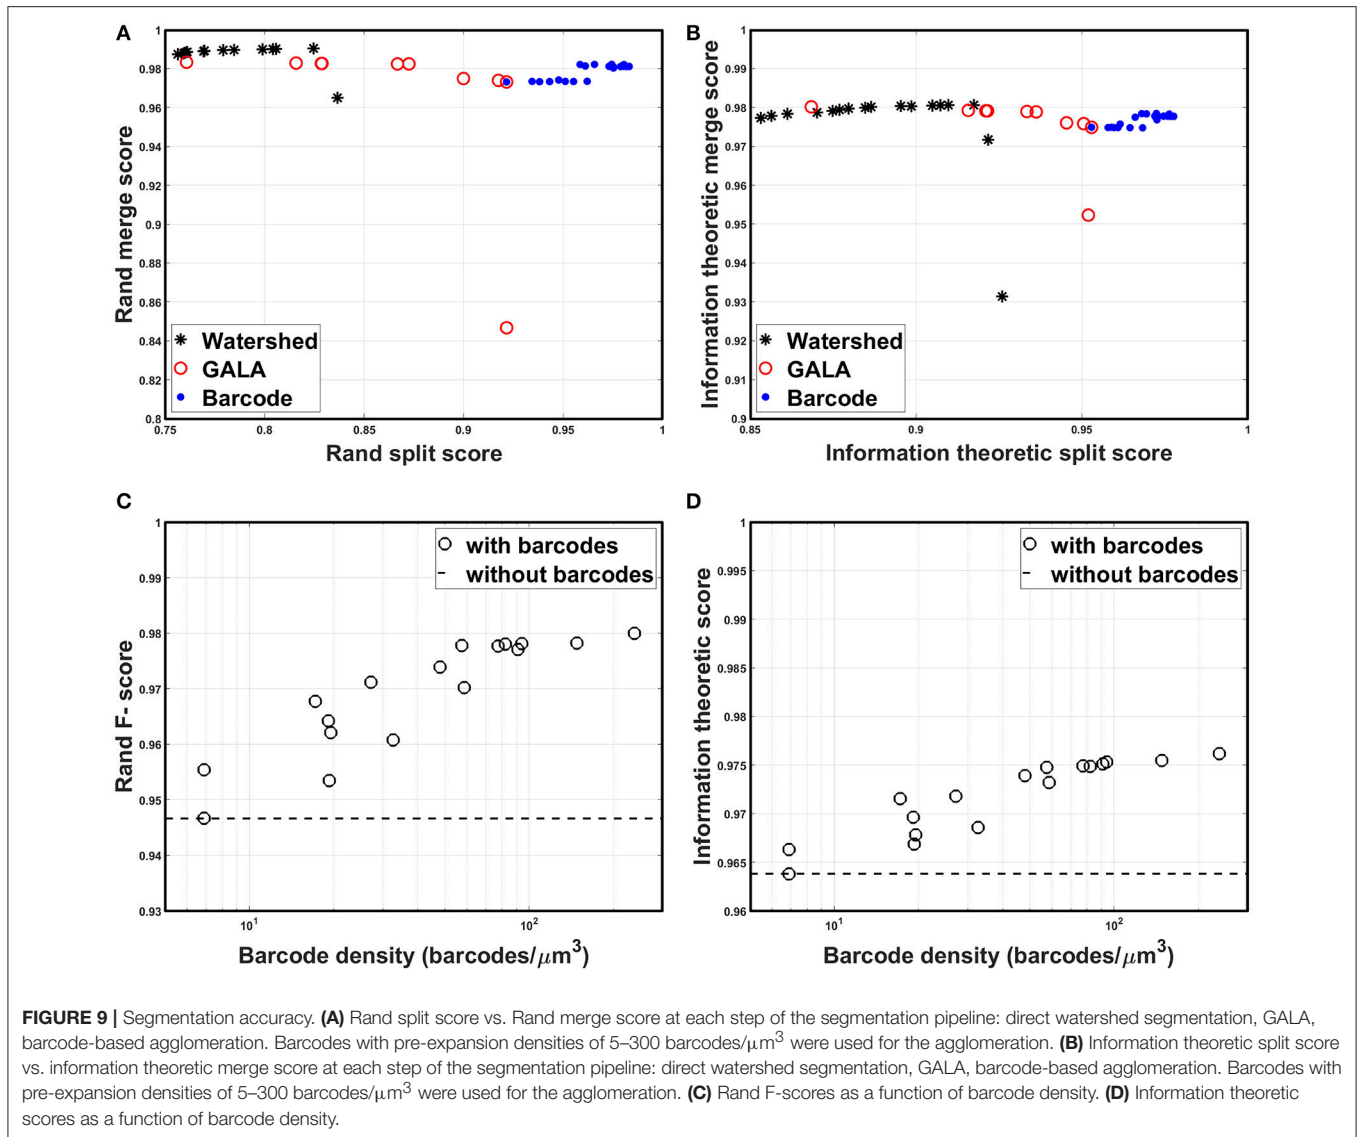
3-D reconstruction of neural circuitry, by enabling multi-color labels such as antibody labels for distinct synaptic sub-types, by enabling integration with *in-situ* transcriptomic and proteomic measurements, and by allowing the use of simple and widely available optical hardware such as confocal microscopes.

We also showed how the use of information-rich cell-identifying RNA barcodes could be incorporated into standard pipelines for image segmentation, by using them to guide super-voxel mergers to improve the accuracy of local reconstruction. We have been deliberately conservative here, focusing on a baseline implementation of an optical reconstruction pipeline, and making no attempt to construct a truly optimal pipeline, nor to push the performance of ExM beyond what has been demonstrated, aside from assuming, in the later sections of the paper, the possibility of high-accuracy barcode readout *in-situ*. There are therefore many directions that could be taken which could lead to major improvements, or to more powerful optical pipelines.

Our simulation study focused on a case where the membranes of all neurons were labeled in a single color for dense reconstruction of neural circuits, but there can be many other scenarios thanks to a diversity of fluorescent labeling strategies. For example, antibody labeling synaptic proteins would allow us to detect synapses directly, and hence, would be a relatively simple, yet powerful, addition to membrane labeling. Pan-neuronal BrainBow could also be useful as it could provide boundary information and cell identity information at the same time (although of course BrainBow is not able to make as many distinct colors as can be created through RNA barcoding, so some neuron-neuron contacts could be ambiguous to segment should they be the same color). Both boundary detection and agglomeration could benefit from such color information, as a multi-color ConvNet is likely to perform better than a single-color counterpart, and agglomeration of colored segments is likely to be an easier task (i.e., BrainBow can serve as an analog form of barcoding that has less information per fluorophore but is spatially more abundant than RNA barcodes).

We have demonstrated one way to incorporate cell-identifying digital barcodes into the segmentation pipeline, by using barcodes as indicators for agglomeration of supervoxels, but the potential of barcodes in the context of brain mapping is much wider. For example, they could be used to guide reinforcement learning of the agglomeration algorithm in place of human annotation, since “correctness” of the merges could be inferred from the barcode consistency at each step of the agglomeration. Barcodes could also be used in a context similar to the “skeleton trace” (Berning et al., 2015), as it directly provides the ground truth for a set of points down the long axis of each neurite. One recent segmentation algorithm (Januszewski et al., 2016) performs image segmentation by “flood-filling” from an initial seed point using a learned flood-filling policy; barcodes could be directly used as seed points as well as a set of points that could be used to check for flood-fills that were overaggressive or that fell short.

Devising new algorithms from scratch specifically to exploit barcode information optimally may prove a particularly fruitful direction for future research. In particular, a framework that can



deal with erroneous barcode readout will be of great utility. Our algorithm assumed that barcodes can be read out without errors for simplicity, but in practice barcodes will have a nonzero error rate determined by sequence errors, the chemical specificity of fluorescent probing, and the optical noise level. Such sequence errors could be handled by many different approaches. For example, maximum a posteriori estimation could be used to perform error correction by using a known pool of barcodes. With such an approach, barcode error rate, analogous to the symbol error rate in digital communication, could be made arbitrarily low if the RNA sequence lengths greatly exceeded $\log_4(\text{number of neurons in the brain})$. Instead of digitizing and error-correcting each barcode, unsupervised learning could also be employed to classify all barcodes in the volume. For instance, the unsupervised clustering algorithm DBSCAN (Ester et al., 1996) could be performed, with the data points defined by the spatial locations of barcodes and their color space intensity values

during sequencing, to identify the entire pool of barcodes as well as the identity of each barcode.

One assumption utilized here is that barcodes are distributed throughout the entire cell, including fine processes such as spines and presynaptic terminals. While fusion of barcodes to synapse-targeting motifs has been shown to lead to their dispersal throughout long axons (at least for some cell types), it remains to be demonstrated whether barcodes can be made completely cell filling (Kebschull et al., 2016), e.g., including spine necks or other fine processes, and whether these molecular trafficking pathways are conserved across all neuron types. We have chosen in the current study not to make use of barcodes located near to the plasma membrane, because the finite localization accuracy of microscopy can contribute to a kind of error wherein barcodes near the cell boundary are mistakenly found on the other side of the boundary. Indeed, here we chose to ignore barcodes close to boundaries (but this approach is sub-optimal as it makes

barcodes in fine processes less likely to be analyzed). A future version of the segmentation algorithm would benefit from a framework that incorporates barcode error-correction at the sequence level, as well as a rating of the confidence with which each barcode should be trusted, as determined by its spatial proximity to the nearest boundary.

Future molecular and algorithmic innovations could directly confront these limitations concerning achievable barcode density, targeting accuracy, and localization uniformity. For example, barcodes could be deliberately targeted to membranes. Then, even if the membrane localization accuracy is such that it is unclear which side a barcode is on, the barcodes on a membrane could still be used to establish a pair of neurons that contacts this boundary. This use of barcodes for constraining neuron pairs at boundaries would lead to a distinct algorithmic framework from the preliminary one demonstrated here, where barcodes are used to guide supervoxel mergers within a BPM that is established without using barcode information. Alternatively, an approach that “flood-fills” (Januszewski et al., 2016) outward from identified barcodes would not require barcodes to be present in every subcellular compartment, instead merely limiting the size of the volume over which accurate flood-filling would be necessary. This might allow far fewer barcodes to be used, yet still provide a benefit by making any segmentation errors into local perturbations that do not propagate over long distances, greatly limiting the impact of local errors on attributions of global cell shape. Thus, different algorithmic frameworks place different requirements on the spatial distributions of barcodes, and on the ways in which barcodes benefit the pipeline, and often, these requirements are more favorable than in the pessimistic baseline case studied here.

Previous work has shown neuronal barcoding to provide information relevant to identifying long-range projections: one can identify a process purely from a barcode located distally, without having to trace it all the way from the parent cell body via image analysis (Kebschull et al., 2016). The use of digital barcodes converts segmentation from a problem in which only local information is available, and hence in which errors in neuronal tracing accumulate with the total distance traced along an axon or dendrite, to a problem in which global information is available, and hence it is only necessary to trace an axon over a short distance from one barcode to another. What is less clear is to what extent barcodes can also assist with the local dense reconstruction problem, i.e., aid the accuracy with which a small volume can be densely reconstructed. We have here through simulations shown that barcodes can also aid in the local reconstruction problem, if the barcodes can be expressed in the cell at a sufficiently high density.

While our analysis was based on current forms of ExM, we expect ExM to continue to improve and make optical brain mapping even more feasible. A higher expansion factor—which may be achieved through additional iterations of expansion on the sample (Chang et al., 2017)—if it results in better resolution, may improve the reconstruction accuracy of thin processes. The resolution of double expansion is ~ 25 nm, but the error due to the gelation and expansion chemistry is ~ 5 – 10 nm (Chang et al., 2017), suggesting that higher expansion factors could, in

principle, reveal yet more biological information. Improving the localization accuracy of the fluorescent probes themselves will become increasingly important with higher expansion factors, as it imposes an upper limit on the effective resolution. For example, small molecule fluorophores could be devised which tile the cell membrane more densely than expressed fluorescent proteins as used here. In this way, we would achieve a dense, spatially precise labeling of cell membrane boundaries, with higher localization accuracy than we have pessimistically assumed here for the case of secondary-antibody-tagged membrane-targeted fluorescent proteins (e.g., <5 vs. ~ 20 nm). Moreover, while we have assumed the use of standard confocal microscopes here, note that more isotropic, compact point-spread functions can be obtained in optical microscopy, which would improve resolution (Gustafsson et al., 1999; Shao et al., 2008; Chen et al., 2014). All of these approaches are in principle compatible with ExM.

In summary, we have demonstrated, via simulations of a conservative baseline case, that dense optical 3-D neural morphology reconstruction using ExM may be feasible, and have illustrated preliminary examples of algorithmic approaches that can be used to implement such an approach, both with and without digital barcodes. This should be viewed as the beginning of a much broader exploration of a vast design space for optical and molecularly annotated connectomics that is enabled by ExM, *in-situ* molecular multiplexing, machine learning, and other emerging techniques.

MATERIALS AND METHODS

Simulation of Expansion Microscopy Images

We implemented the simulator in MATLAB R2016A and it was run on a workstation with an NVIDIA K40c GPU and 128GB of RAM. First, the input volume was re-sampled with 6 nm of isotropic resolution and the voxels that correspond to the membrane of each neuron were identified by detecting the boundary of each segment. In order to reflect the finite localization accuracy in ExM (which originates from non-zero mesh size of the hydrogel and the non-zero size of antibodies used for tagging specific proteins Chang et al., 2017), fluorophores were assigned with random labeling densities near the boundary voxels with a standard deviation of 20 nm. Fluorophores were also added to non-membrane parts of the volume, with random densities, to simulate fluorophores with non-specific binding and auto-fluorescence. Then, local clusters of fluorophores, or puncta, were created by convolving each fluorophore with a random size Gaussian kernel to reflect local clusters of fluorophores which result from multiple fluorophores being present on a single antibody-delivered tag, and from multiple antibodies binding to a single fluorescent protein target. The synthetic volume was projected onto a 3-D image stack by convolving the volume with the 3-D point-spread-function (PSF) of a confocal microscope with an objective lens with 40x magnification and 1.15 numerical aperture. The PSF was obtained by combining the excitation PSF and the emission PSF, which were in turn computed by evaluating the following

Debye integral (Gu, 2000) where k , λ , v , u , α , $P(\theta)$, $J_0(\cdot)$ denote the wave number, the wavelength, normalized radial coordinate, normalized axial coordinate, the half-angle of the numerical aperture, the apodization function of the microscope and the zeroth order Bessel function of the first kind, respectively:

$$U(\mathbf{v}, \mathbf{u}) = \frac{2\pi i}{\lambda} e^{-ikz} \int_0^\alpha P(\theta) J_0\left(\frac{v \sin \theta}{\sin \alpha}\right) \exp\left(\frac{i u \sin^2 \theta}{2 \sin^2 \alpha}\right) \sin \theta d\theta. \quad (3)$$

For modeling 20x ExM, the PSF was scaled down by a factor of 20 instead of expanding the volume by a factor of 20, for efficient computation. Finally, Poisson noise and read-out noise were added to the image (Supplementary Figure 1). Almost every parameter (e.g., labeling density, noise level) in each step of the simulation was randomized to reflect the stochastic nature of fluorescent labeling and imaging and to avoid over-fitting our analysis to a specific imaging condition. For simplicity, it was assumed that barcodes were randomly distributed over the volume, including throughout fine processes, and also that they can be read out without error. **Table 1** summarizes on the simulation parameters. All parameters are in units corresponding to the original scale of the tissues (i.e., before expansion). The confocal microscopy parameters were matched to values of a typical confocal microscope, and the fluorophore density parameters were chosen to match the qualitative appearance of the simulated images to that of actual 20x ExM images (Chang et al., 2017). Localization accuracy of 20 nm was assumed, as our current 20x ExM protocol uses both primary and secondary antibodies to deliver the fluorescent label, each of which has a size around 14 nm (i.e., if we assume that two localization error vectors are uncorrelated, the added localization accuracy can be calculated as $\sqrt{(14 \text{ nm})^2 + (14 \text{ nm})^2} \cong 20 \text{ nm}$).

Four different ground truth volumes (Supplementary Figure 2) were used for the simulation and the size of each volume was $\sim 3 \mu\text{m}$ in each dimension. Twelve images from three ground truth volumes were used for training the segmentation algorithm and one image from a different ground truth volume was used for testing the algorithm. The simulated images from a single ground truth volume were significantly different as the simulation involved randomized fluorophore allocation as well as many randomized parameters, and hence, we took advantage of this as a part of the data augmentation for training the algorithms.

Boundary Detection

We implemented a ConvNet for boundary detection (Ciresan et al., 2012) using MATLAB R2016A with MatConvNet (Sun, 2015; Vedaldi and Lenc, 2015) and the computation was performed on a workstation with a NVidia K40c GPU and a medium size cluster with multiple GPUs. To train this network, we first implemented a 2-D ConvNet that had the same architecture except that it had only 2-D convolution layers (i.e., the filter size in 11th, 12th, and 13th convolution layers was $3 \times 3 \times 1$ instead of $3 \times 3 \times 3$), and trained this network. Then, the weights in the first to eighth convolution layers were taken to initialize the same part of the 3-D ConvNet. Then, the

3-D ConvNet was trained on a patch-by-patch basis with back-propagation and stochastic gradient descent with a learning rate of 0.001, a momentum of 0.9 and a batch size of 100. The patches were randomly sampled from 12 training volumes. In order to deal with the training imbalance (i.e., there are more non-boundary pixels than boundary pixels), the input patches were sampled to match their numbers. Data augmentation was done by flipping and rotating the input patches. In order to exploit as much 3-D information as possible to further improve the accuracy of the BPM, two ConvNets with the same architecture were trained with two different input orientations—one in xy ($85 \times 85 \times 7$) and the other in xz/yz ($85 \times 7 \times 85$, $7 \times 85 \times 85$).

For inference (i.e., obtaining the BPM from the image), we implemented a forward pass function that can be applied to a volume of arbitrary size instead of using a sliding window for patch-by-patch evaluation, similar to ZNN (Zlateski et al., 2016), to reduce the number of redundant convolution operations for speed up. Using two ConvNets, three BPMs were obtained for a single volume, and then the median value of three was taken for each voxel to obtain the final BPM (Supplementary Figure 3).

Over-Segmentation and Agglomeration

We applied a 3-D median filter to the BPM to make the boundary smooth and an H-minima transform to remove local minima that were too shallow with a threshold value T of 0.01 as pre-processing steps followed by watershed transformation to obtain the initial over-segmentation. Once over-segmentation was performed with a watershed transform, the segments were merged using GALA. To train GALA, the watershed segmentation result, the BPM and the input image were fed in and moment and histogram were chosen as the training features. The data sets used for training the ConvNets were used for training GALA as well. After training, we applied the algorithm to the test image with nine threshold values ranging from 0.1 to 0.9 to obtain a hierarchical segmentation tree. This tree was combined with the barcodes by assigning the barcodes to the segments, or nodes, at the lowest level of the segmentation tree. Barcodes that were within 30 nm from the boundary of segments were ignored, as the barcodes that are near a boundary might be assigned to neurons on the other side due to the finite localization accuracy of ExM assumed here and the potential for boundary shifts in the ConvNet output. After assigning the barcodes to the nodes at the lowest level, the numbers of distinct barcodes in the nodes at a higher level were examined. If the number of different barcodes of a node did not exceed one, all child nodes of the node were merged. This process was repeated until the top level of the tree was reached and the nodes with only 0 or 1 barcodes at the highest possible levels were identified. Finally, the nodes that had the same barcode were merged.

AUTHOR CONTRIBUTIONS

YY and PD designed the segmentation algorithm. YY and JW built the ExM simulator. JC experimentally characterized ExM. AM and EB conceived of and led the project. YY, AM, and EB wrote the manuscript.

FUNDING

YY acknowledges funding by the Samsung Scholarship. PD acknowledges the Singapore National Science Scholarship. JC acknowledges the Simons Fellowship. AM acknowledges funding by IARPA D16PC00008. EB acknowledges funding by the HHMI-Simons Faculty Scholars Program, NIH 1R41MH112318, the Open Philanthropy Project, NIH 1R01MH110932, U.S. Army Research Laboratory and the U. S. Army Research Office under contract/grant number W911NF1510548, NIH 1RM1HG008525, Jeremy and Joyce Wertheimer, NIH Director's Pioneer Award 1DP1NS087724, John Doerr, and the MIT Media Lab.

REFERENCES

- Arganda-Carreras, I., Turaga, S. C., Berger, D. R., Cireşan, D., Giusti, A., Gambardella, L. M., et al. (2015). Crowdsourcing the creation of image segmentation algorithms for connectomics. *Front. Neuroanat.* 9:142. doi: 10.3389/fnana.2015.00142
- Berger, D., Schalek, R., Kasthuri, N., Tapia, J., Hayworth, K., and Lichtman, J. (2013). *SNEMI3D: 3D Segmentation of Neurites in EM Images*. Available online at: <http://brainiac2.mit.edu/SNEMI3D>
- Berning, M., Boergens, K. M., and Helmstaedter, M. (2015). SegEM: efficient image analysis for high-resolution connectomics. *Neuron* 87, 1193–1206. doi: 10.1016/j.neuron.2015.09.003
- Cai, D., Cohen, K. B., Luo, T., Lichtman, J. W., and Sanes, J. R. (2013). Improved tools for the brainbow toolbox. *Nat. Methods* 10, 540–547. doi: 10.1038/nmeth.2450
- Chang, J. B., Chen, F., Yoon, Y. G., Jung, E., Babcock, H., Kang, J. S., et al. (2017). Iterative expansion microscopy. *Nat. Methods* 14, 593–599. doi: 10.1038/nmeth.4261
- Chen, B. C., Legant, W. R., Wang, K., Shao, L., Milkie, D. E., Davidson, M. W., et al. (2014). Lattice light-sheet microscopy: imaging molecules to embryos at high spatiotemporal resolution. *Science* 346:1257998. doi: 10.1126/science.1257998
- Chen, F., Tillberg, P. W., and Boyden, E. S. (2015). Expansion microscopy. *Science* 347, 543–548. doi: 10.1126/science.1260088
- Chen, F., Wassie, A. T., Cote, A. J., Sinha, A., Alon, S., Asano, S., et al. (2016). Nanoscale imaging of RNA with expansion microscopy. *Nat. Methods* 13, 679–684. doi: 10.1038/nmeth.3899
- Chen, K. H., Boettiger, A. N., Moffitt, J. R., Wang, S., and Zhuang, X. (2015). Spatially resolved, highly multiplexed RNA profiling in single cells. *Science* 348:aaa6090. doi: 10.1126/science.aaa6090
- Chen, X., Sun, Y.-C., Church, G. M., Lee, J. H., and Zador, A. M. (2017). Efficient *in situ* barcode sequencing using padlock probe-based BaristaSeq. *bioRxiv*. doi: 10.1101/180323
- Ciresan, D., Giusti, A., Gambardella, L. M., and Schmidhuber, J. (2012). “Deep neural networks segment neuronal membranes in electron microscopy images,” in *Advances in Neural Information Processing Systems*, 2843–2851. Available online at: <https://papers.nips.cc/paper/4741-deep-neural-networks-segment-neuronal-membranes-in-electron-microscopy-images>
- Dumitriu, D., Rodriguez, A., and Morrison, J. H. (2011). High-throughput, detailed, cell-specific neuroanatomy of dendritic spines using microinjection and confocal microscopy. *Nat. Protoc.* 6, 1391–1411. doi: 10.1038/nprot.2011.389
- Ester, M., Kriegel, H.-P., Sander, J., and Xu, X. (1996). “A density-based algorithm for discovering clusters in large spatial databases with noise,” in *AAI Knowledge Discovery in Databases* (Portland, OR: AAAI Press), 226–231.
- Germroth, P. G., Gourdie, R. G., and Thompson, R. P. (1995). Confocal microscopy of thick sections from acrylamide gel embedded embryos. *Microsc. Res. Tech.* 30, 513–520. doi: 10.1002/jemt.1070300608
- Gu, M. (2000). *Advanced Optical Imaging Theory* (Berlin; Heidelberg: Springer-Verlag).

ACKNOWLEDGMENTS

We thank Jeong Seuk (Louis) Kang, Sam Rodrigues, Daniel Goodwin, Grace Huynh, David Rolnick, Nick Barry, Paul Myers, Emmanouil (Manos) Karagiannis, and other members of the Synthetic Neurobiology Group for helpful discussions.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fncom.2017.00097/full#supplementary-material>

- Gustafsson, M. G., Agard, D. A., and Sedat, J. W. (1999). I5M: 3D widefield light microscopy with better than 100 nm axial resolution. *J. Microsc.* 195, 10–16. doi: 10.1046/j.1365-2818.1999.00576.x
- Hausen, P., and Dreyer, C. (1981). The use of polyacrylamide as an embedding medium for immunohistochemical studies of embryonic tissues. *Stain Technol.* 56, 287–293. doi: 10.3109/10520298109067329
- Januszewski, M., Maitin-Shepard, J., Li, P., Kornfeld, J., Denk, W., and Jain, V. (2016). Flood-filling networks. *arXiv:1611.00421*.
- Kasthuri, N., and Lichtman, J. (2007). The rise of the ‘projectome.’ *Nat. Methods* 4, 307–308. doi: 10.1038/nmeth0407-307
- Kasthuri, N., Hayworth, K. J., Berger, D. R., Schalek, R. L., Conchello, J. A., Knowles-Barley, S., et al. (2015). Saturated reconstruction of a volume of neocortex. *Cell* 162, 648–661. doi: 10.1016/j.cell.2015.06.054
- Kebschull, J. M., Garcia da Silva, P., Reid, A. P., Peikon, I. D., Albeanu, D. F., and Zador, A. M. (2016). High-throughput mapping of single-neuron projections by sequencing of barcoded RNA. *Neuron* 91, 975–987. doi: 10.1016/j.neuron.2016.07.036
- Kurenkov, V. F., Hartan, H.-G., and Lobanov, F. I. (2001). Alkaline hydrolysis of polyacrylamide. *Russ. J. Appl. Chem.* 74, 543–554. doi: 10.1023/A:1012786826774
- Langer-Safer, P. R., Levine, M., and Ward, D. C. (1982). Immunological method for mapping genes on *Drosophila* polytene chromosomes. *Proc. Natl. Acad. Sci. U.S.A.* 79, 4381–4385. doi: 10.1073/pnas.79.14.4381
- Lee, J. H., Daugharthy, E. R., Scheiman, J., Kalhor, R., Yang, J. L., Ferrante, T. C., et al. (2014). Highly multiplexed subcellular RNA sequencing *in situ*. *Science* 343, 1360–1363. doi: 10.1126/science.1250212
- Lee, K., Zlateski, A., Vishwanathan, A., and Seung, H. S. (2015). “Recursive training of 2D-3D convolutional networks for neuronal boundary detection,” in *Advances in Neural Information Processing Systems*, eds C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett (Montréal, QC: MIT Press), 3573–3581.
- Liu, T., Jurrus, E., and Seyedhosseini, M. (2012). “Watershed merge tree classification for electron microscopy image segmentation - IEEE Xplore document,” in *21st International Conference on Pattern Recognition (ICPR)*, Tsukuba: IEEE.
- Livet, J., Weissman, T. A., Kang, H., Draft, R. W., Lu, J., Bennis, R. A., et al. (2007). Transgenic strategies for combinatorial expression of fluorescent proteins in the nervous system. *Nature* 450, 56–62. doi: 10.1038/nature06293
- Llinás, R. R. (2003). The contribution of Santiago Ramon y Cajal to functional neuroscience. *Nat. Rev. Neurosci.* 4, 77–80. doi: 10.1038/nrn1011
- Lu, R., Neff, N. F., Quake, S. R., and Weissman, I. L. (2011). Tracking single hematopoietic stem cells *in vivo* using high-throughput sequencing in conjunction with viral genetic barcoding. *Nat. Biotechnol.* 29, 928–933. doi: 10.1038/nbt.1977
- Marblestone, A. H., Daugharthy, E. R., Kalhor, R., Peikon, I. D., Kebschull, J. M., Shipman, S. L., et al. (2014). Rosetta brains: a strategy for molecularly-annotated connectomics. *arXiv:1404.510*
- Miyasaka, N., Arganda-Carreras, I., Wakisaka, N., Masuda, M., Sümbül, U., Seung, H. S., et al. (2014). Olfactory projectome in the zebrafish forebrain revealed by genetic single-neuron labelling. *Nat. Commun.* 5:3639. doi: 10.1038/ncomms4639

- Nunez-Iglesias, J., Kennedy, R., Parag, T., Shi, J., and Chklovskii, D. B. (2013). Machine learning of hierarchical clustering to segment 2D and 3D images. *PLoS ONE* 8:e71715. doi: 10.1371/journal.pone.0071715
- 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. *Front. Neuroinformatics* 8:34. doi: 10.3389/fninf.2014.00034
- O'Connell, P. B., and Brady, C. J. (1976). Polyacrylamide gels with modified cross-linkages. *Anal. Biochem.* 76, 63–73. doi: 10.1016/0003-2697(76)90264-5
- Peikon, I. D., Gizatullina, D. I., and Zador, A. M. (2014). *In vivo* generation of DNA sequence diversity for cellular barcoding. *Nucleic Acids Res.* 42, e127. doi: 10.1093/nar/gku604
- Robles, E., Laurell, E., and Baier, H. (2014). The retinal projectome reveals brain-area-specific visual representations generated by ganglion cell diversity. *Curr. Biol.* 24, 2085–2096. doi: 10.1016/j.cub.2014.07.080
- Shao, L., Isaac, B., Uzawa, S., Agard, D. A., Sedat, J. W., and Gustafsson, M. G. (2008). 15S: wide-field light microscopy with 100-nm-scale resolution in three dimensions. *Biophys. J.* 94, 4971–4983. doi: 10.1529/biophysj.107.120352
- Sümbül, U., Roussien, D. Jr., Chen, F., Barry, N., Boyden, E. S., Cai, D., et al. (2016). Automated scalable segmentation of neurons from multispectral images. *arXiv:1611.00388*.
- Sun, P. (2015). *MexConv3D: Matlab Mex Implementation of the Basic Operations for 3D (Volume) Convolutional Neural Network*. Available online at: <https://github.com/pengsun/MexConv3D>
- Tanaka, T., Sun, S.-T., Nishio, I., Swislow, G., and Shah, A. (1980). Phase transitions in ionic gels. *Phys. Rev. Lett.* 45, 1636–1639. doi: 10.1103/PhysRevLett.45.1636
- Tillberg, P. W., Chen, F., Piatkevich, K. D., Zhao, Y., Yu, C. C., English, B. P., et al. (2016). Protein-retention expansion microscopy of cells and tissues labeled using standard fluorescent proteins and antibodies. *Nat. Biotechnol.* 34, 987–992. doi: 10.1038/nbt.3625
- Vedaldi, A., and Lenc, K. (2015). “MatConvNet: convolutional neural networks for MATLAB. MM,” in *15 Proceedings of the 23rd ACM International Conference on Multimedia* (Brisbane, QL), 689–692.
- Wilt, B. A., Burns, L. D., Wei Ho, E. T., Ghosh, K. K., Mukamel, A. E., and Schnitzer, M. J. (2009). Advances in light microscopy for neuroscience. *Annu. Rev. Neurosci.* 32, 435–506. doi: 10.1146/annurev.neuro.051508.135540
- Zador, A. M., Dubnau, J., Oyibo, H. K., Zhan, H., Cao, G., and Peikon, I. D. (2012). Sequencing the connectome. *PLoS Biol.* 10:e1001411. doi: 10.1371/journal.pbio.1001411
- Zlateski, A., Lee, K., and Seung, H. S. (2016). “ZNN – a fast and scalable algorithm for training 3D convolutional networks on multi-core and many-core shared memory machines,” in *IEEE International Parallel and Distributed Processing Symposium* (Chicago, IL: IPDPS).

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2017 Yoon, Dai, Wohlwend, Chang, Marblestone and Boyden. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.