

# Multi-Layered Maps of Neuropil with Segmentation-Guided Contrastive Learning

Sven Dorkenwald<sup>1,5,6,\*</sup>, Peter H. Li<sup>1,\*</sup>, Michał Januszewski<sup>2</sup>, Daniel R. Berger<sup>3</sup>, Jeremy Maitin-Shepard<sup>1</sup>, Agnes L. Bodor<sup>4</sup>, Forrest Collman<sup>4</sup>, Casey M. Schneider-Mizell<sup>4</sup>, Nuno Maçarico da Costa<sup>4</sup>, Jeff W. Lichtman<sup>3</sup>, Viren Jain<sup>1</sup>

<sup>1</sup> Google Research, Mountain View, CA

<sup>2</sup> Google Research, Zürich, Switzerland

<sup>3</sup> Dept. of Molecular and Cellular Biology, Center for Brain Science, Harvard, Cambridge, MA

<sup>4</sup> Allen Institute for Brain Science, Seattle, WA

<sup>5</sup> Princeton Neuroscience Institute, Princeton University, Princeton, NJ

<sup>6</sup> Computer Science Department, Princeton University, Princeton, NJ

\* co-first authors

Correspondence: [viren@google.com](mailto:viren@google.com)

## Abstract

Maps of the nervous system that identify individual cells along with their type, subcellular components, and connectivity have the potential to reveal fundamental organizational principles of neural circuits. Volumetric nanometer-resolution imaging of brain tissue provides the raw data needed to build such maps, but inferring all the relevant cellular and subcellular annotation layers is challenging. Here, we present Segmentation-Guided Contrastive Learning of Representations (“SegCLR”), a self-supervised machine learning technique that produces highly informative representations of cells directly from 3d electron microscope imagery and segmentations. When applied to volumes of human and mouse cerebral cortex, SegCLR enabled the classification of cellular subcompartments (axon, dendrite, soma, astrocytic process) with 4,000-fold less labeled data compared to fully supervised approaches. Surprisingly, SegCLR also enabled inference of cell types (neuron vs. glia and subtypes of each) from fragments with lengths as small as 10 micrometers, a task that can be difficult for humans to perform and whose feasibility greatly enhances the utility of imaging portions of brains in which many neuron fragments terminate at a volume boundary. These predictions were further augmented via Gaussian process uncertainty estimation to enable analyses restricted to high confidence subsets of the data.

## Introduction

Biological understanding has been enabled by annotating parts of organisms and elucidating their interrelationships. In the brain, numerous types of neuronal and glial cells have been discovered and cataloged according to their morphological, physiological, and molecular properties<sup>1–5</sup>, typically using methods that interrogate cells in a sparse or isolated setting. Further discoveries would be enabled by producing maps that contain dense assemblies of cells and multiple layers of annotation in the context of a neural circuit or region<sup>6–9</sup>.

Producing dense maps of neuropil is challenging due to the multiple scales of brain structures (e.g., nanometers for a synapse versus millimeters for an axon)<sup>10</sup>, and the vast number of objects in neuropil which must be individually segmented, typed, and annotated with the interrelationships. Volumetric electron microscopy (EM) has proven to be an effective way to image brain structures over both large and small scales<sup>11,12</sup> and automated segmentation of volume EM data has also shown significant progress<sup>13–17</sup>, including the demonstration of millimeter-scale error-free run lengths<sup>18</sup>.

Here, we address the problem of efficiently inferring types and annotations of segmented structures by introducing Segmentation-Guided Contrastive Learning of Representations (SegCLR), a machine learning approach that is scalable in three important respects: first, the representations produced by a single SegCLR model can be used for a diverse set of annotation tasks (e.g., local identification of cellular subcompartments, assigning type to an entire cell or fragment); second, the compact representations learned by SegCLR enable accurate downstream analyses with simple linear classifiers, removing the need to train large networks (such as 3d convolutional networks) for each additional task; and third, SegCLR reduces the amount of ground truth labeling required for specific tasks by up to 4 orders of magnitude. Perhaps most intriguingly, we show that SegCLR enables a type of annotation which is challenging for both automated methods as well as human experts: inferring the cell type from a short length (~10-100  $\mu\text{m}$ ) of cortical cell fragment. This capability has important implications for the utility of cortical EM datasets that so far encompass only subsets of whole brains.

Previous machine learning methods for neuropil annotation have primarily used features that were hand-designed or derived from supervised learning, including random forests trained on hand-designed features<sup>9,19</sup>, convolutional networks trained on 2d projections of neuropil (“Cellular Morphology Networks”)<sup>20,21</sup>, or 3d convolutional networks trained directly on voxels<sup>22</sup>. Schubert et al.<sup>20</sup> trained representations from 2d projections using a self-supervised triplet loss, but it was not reported whether these representations are suitable for downstream analyses. Previous results on cell type classification of small neurite fragments required larger spatial context, precomputed organelle masks, and manually engineered features<sup>19–21</sup>, or used a single local view to achieve modest classification accuracy on a limited set of classes<sup>20</sup>.

Self-supervised learning has emerged as a broadly successful technique for producing representations of text<sup>23</sup> and pixels<sup>24</sup> without the use of labeled data. Weis et al. introduced a method for self-supervised learning of neuronal morphology that operates on coarse skeleton representations of neurons<sup>25</sup>; however, similar to previous morphological clustering methods<sup>26–28</sup>, this approach did not demonstrate important capabilities of SegCLR (e.g. classification of subcellular structures and type inference from small cellular fragments). Self-supervised learning has also been previously explored for content-based image retrieval in unsegmented 3d connectomic datasets<sup>29</sup>, and for image retrieval and classification in 2d medical image datasets<sup>30</sup>.

SegCLR takes inspiration from recent advances in self-supervised contrastive learning<sup>24</sup> while introducing a segmentation-guided loss function in which positive example pairs are drawn from

nearby, but not necessarily overlapping, cutouts of the same segmented cell. In addition to raw volumetric data, this approach therefore also requires 3d segmentation of individual cells throughout the volume, which is a typical requirement for subcellular annotation methods. As we demonstrate, current automated segmentation methods are sufficiently accurate to be used to train SegCLR without further human proofreading.

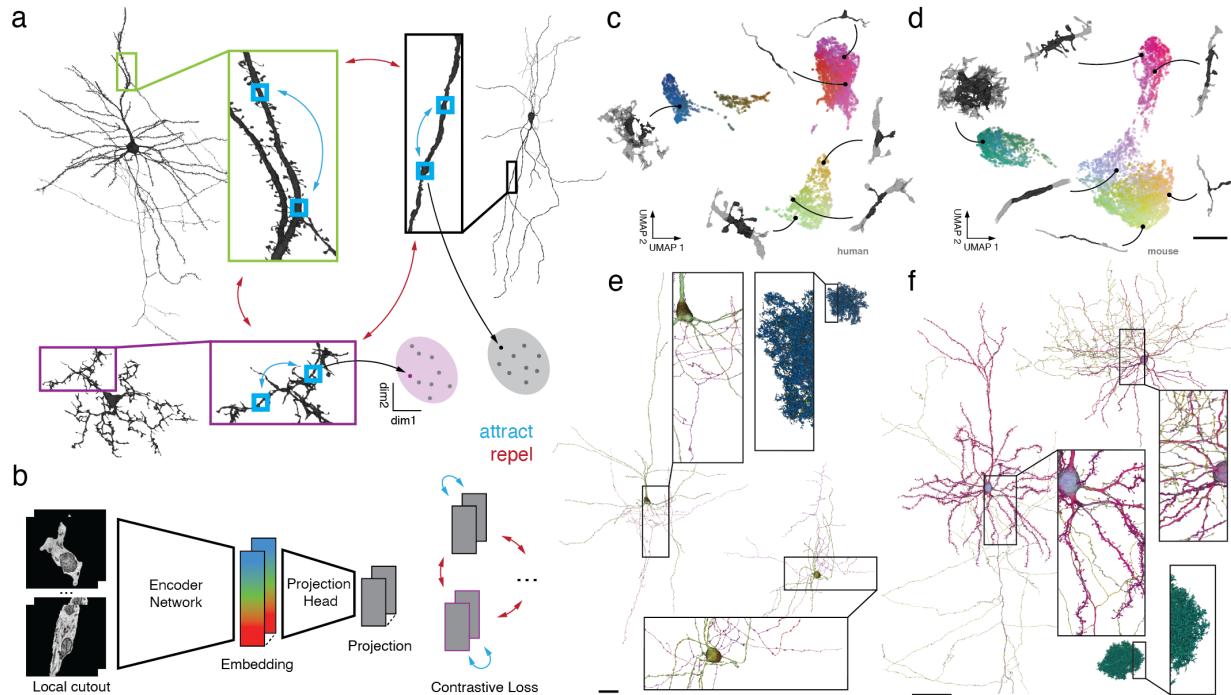
In addition to pursuing accurate classifications, it is critical to be able to estimate the uncertainty associated with specific predictions, such that biological inferences can be restricted to high confidence results. We show that SegCLR can be combined with Gaussian processes<sup>31</sup> to provide a highly practical means of uncertainty estimation.

## Results

### Training and inference of SegCLR embeddings

Representations that track separate cells through dense neuropil, such as instance segmentations<sup>18</sup> or skeletonizations<sup>32</sup>, have proven fundamental to biological interpretation, as have annotations of select features of interest such as synapses<sup>19,33</sup>, organelles<sup>34</sup>, and cellular subcompartments<sup>22</sup>. In contrast, the aim of SegCLR is to produce tractable “embeddings”: vector representations that capture rich biological features in a dimensionally reduced space, and in which vector distance maps to a concept of biological distinctness (Fig. 1). These embeddings capture features relevant to a range of downstream tasks, and can be trained without manual feature engineering. Depending on the downstream application, embeddings can also be deployed without any requirement for manual proofreading or ground truth labeling, or with these requirements significantly reduced<sup>35</sup>. Each SegCLR embedding represents a local 3d view of EM data, and is focused on an individual cell or cell fragment within dense neuropil via an accompanying segmentation. Computed for billions of local views across large connectomic datasets, embeddings can directly support local annotation tasks (Fig. 2), or be flexibly combined at larger scales to support annotation at the level of whole cells or cell fragments (Figs. 3-4).

SegCLR builds on recent advances in contrastive learning of image representations<sup>24,35</sup>, with modifications that leverage freely-available dense automated instance segmentations of neurons and glia<sup>11,12</sup>. Contrastive learning approaches aim to learn representations by maximizing agreement between matched (“positive”) examples in a learned latent space. SegCLR selects example pairs with respect to the segmentation: positive pairs were drawn from nearby locations (within 150  $\mu$ m skeleton path length) on the same object and trained to have similar representations, while negative pairs were drawn from separate objects and trained to have dissimilar representations (Fig. 1a). Input preprocessing also leveraged the segmentation: local 3d views of EM data, 4-5  $\mu$ m on a side at 32-40 nm voxel resolution, were presented to the embedding network after being masked to feature only the segmented object at the center of the field of view (Fig. 1b, left). The network architecture was based on ResNet-18<sup>36</sup>, with convolutional filters extended to 3d. We added three bottleneck layers reducing the output layer to a 64-dimensional embedding layer, and a projection head further reducing the representations to 16 dimensions on which the contrastive training loss is applied<sup>35</sup> (Fig. 1b, right).

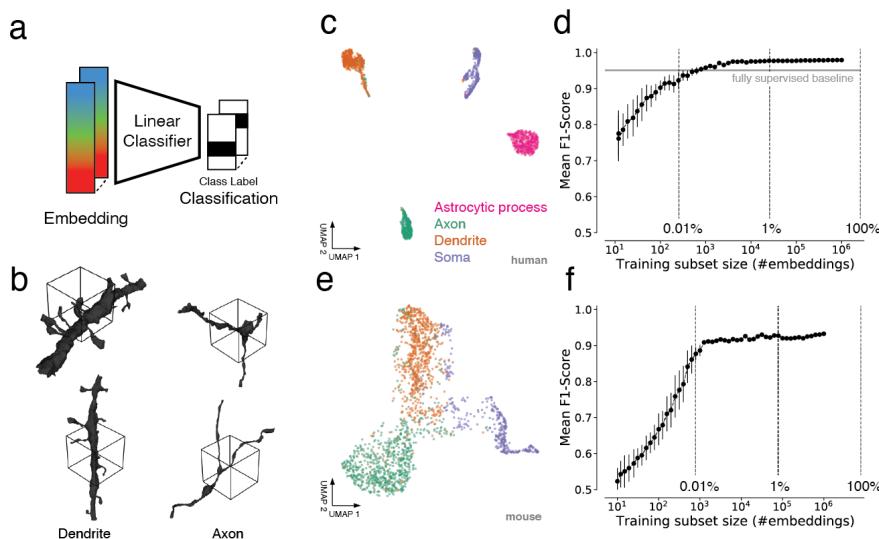


**Figure 1. SegCLR: Segmentation-Guided Contrastive Learning of Representations.**

**a.** In SegCLR, positive pairs (blue double-headed arrows) are chosen from proximal but not necessarily overlapping 3d views (small blue boxes) of the same segmented cell, while negative pairs (red double-headed arrows) are chosen from different cells. The SegCLR network is trained to produce an embedding vector for each local 3d view such that embeddings are more similar for positive pairs compared to negative pairs (cartoon of clustered points). **b.** The input to the embedding network is a local 3d view ( $4.1 \times 4.1 \times 4.3 \mu\text{m}$  at  $32 \times 32 \times 33 \text{ nm}$  resolution for human data;  $4.1 \times 4.1 \times 5.2 \mu\text{m}$  at  $32 \times 32 \times 40 \text{ nm}$  resolution for mouse) from the EM volume, masked by the segmentation for the object at the center of the field of view. A encoder network based on a ResNet-18 is trained to produce embeddings, via projection heads and a contrastive loss that are used only during training. **c-d.** Visualization via UMAP projection of the SegCLR embedding space for the human temporal cortex and mouse visual cortex datasets. Points for a representative sample of embeddings are shown, colored via 3d UMAP RGB, with the corresponding 3d morphology illustrated for 6 locations (network input in black,  $10 \times 10 \times 10 \mu\text{m}$  context with lower opacity). Biologically related objects cluster naturally in this space without any ground truth label supervision, e.g. myelinated axons and initial axons occupy a space adjacent to the rest of the axons in the mouse dataset. Scale bars (c,d):  $5 \mu\text{m}$  **e-f.** Embeddings visualized along the extent of representative human (e) and mouse (f) cells. Each mesh rendering is colored according to the 3d UMAP RGB of the nearest embedding for the surrounding local 3d view. The embedding colors readily distinguish different types of cells, as well as different subcellular structures. Some axons are cut off to fit. Scale bars (e,f):  $100 \mu\text{m}$ .

We trained SegCLR separately on two large-scale, publicly available EM connectomic datasets, one from human temporal cortex<sup>12</sup> and one from mouse visual cortex<sup>11</sup>, that were produced via different imaging and segmentation techniques. We then inferred SegCLR embeddings with

partially overlapping fields of view over all non-trivial objects (at least 1,000 voxels) within each volume. This produced a 64-dimensional embedding vector for each masked local 3d view, for a total of 3.9 billion and 4.2 billion embeddings for the human and mouse datasets respectively. SegCLR thus adds modest storage overhead relative to the full EM dataset size (human: 980 GB versus 1.4 PB at 4x4x33 nm; mouse: 1 TB versus 234 TB at 8x8x40 nm). Visualizing an illustrative subset of the resulting embeddings after dimensionality reduction via UMAP projection<sup>37</sup> revealed structure across each embedding space (Fig. 1c-d). Visualizing embeddings over individual cells also revealed structure within and between them (Fig. 1e-f), suggesting the potential for embeddings to solve diverse downstream tasks.



**Figure 2. Subcompartment classification of SegCLR embeddings.**

**a.** Embedding vectors computed across the extent of the EM datasets can be used as input to downstream tasks, such as subcompartment classification. Each embedding represents a single local view (~4-5  $\mu$ m on a side). **b.** Ground truth examples of axon and dendrite subcompartment classes from the human temporal cortex dataset. The local 3d views for single embeddings is indicated by the wireframe cubes. **c.** Embedding clusters from the human cortical dataset visualized via 2d UMAP. Each point is an embedding, colored by its ground truth subcompartment class as judged without reference to the embeddings. **d.** Evaluation of linear classifiers trained for the subcompartment task on the human dataset. The mean F1 score across classes was computed for networks trained using varying-sized subsets of the full available training data. For smaller subsets, the mean and standard deviation of multiple subset resamplings is shown, while for larger subsets only a single sampling was evaluated. The horizontal line indicates the performance of a fully supervised ResNet-18 classifier trained on the full available training data. **e.** As in (c), for the mouse visual cortex dataset and three ground truth classes (axon, dendrite, soma). **f.** As in (d), for the mouse dataset.

### Cellular subcompartment classification

Embedding vectors representing local segment views throughout the EM datasets can be applied to a variety of downstream tasks, such as clustering, similarity search, or classification (Fig. 2a). We first examined using SegCLR to distinguish cellular subcompartments such as axons, dendrites, and somas (Fig. 2b). In the human cortical dataset we also included astrocytic

processes in this task, as a distinct kind of subcompartment for which we had ground truth labeling. On a set of segmented object locations with expert labeled subcompartment identities, the respective SegCLR embeddings formed largely separable clusters in embedding space (Fig. 2c, e). A linear classifier trained to distinguish embeddings from the human cortical dataset could identify subcompartments in a held out test set with 0.979 mean F1-Score, while on the mouse dataset, axon, dendrite, and soma classification reached 0.933 mean F1-Score.

We also tested reducing the ground truth labeling requirements, and compared the performance of subcompartment classification using SegCLR embeddings versus directly training a fully supervised subcompartment classification network<sup>22</sup>. The supervised network input data and network architecture (ResNet-18) were identical to the SegCLR setup, except that we replaced the SegCLR bottleneck and contrastive projection head with a classification softmax. On the 4-class subcompartment task, the embedding-based classification matches the performance of direct supervised training while requiring roughly 4,000 times less labeled training data, and exceeds supervised performance when using all labeled data (Fig. 2d). The number of training examples needed to match fully supervised performance is roughly 700, and having embeddings available could facilitate active learning strategies<sup>38</sup> for further gains in efficiency and performance.

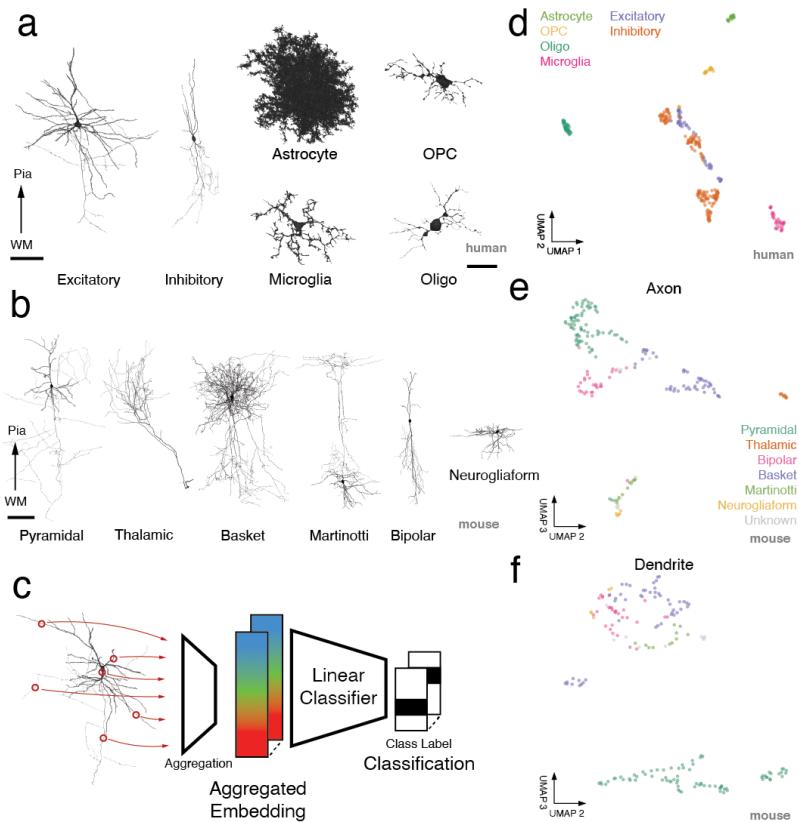
#### Classification of subpopulations of neuron and glia

We next examined using SegCLR for cell type classification, starting with larger segmented cells for which we had ground truth classifications (Fig. 3a,b). These cells typically have somas within the volume, but are not complete due to processes that exit the volume as well as split errors in the reconstructions. Some cells also suffer from merge errors, which complicate evaluations against ground truth, especially when evaluating “local” classifications (see below); therefore we manually proofread some large objects to remove merge errors (see Methods).

While each SegCLR embedding vector represents a masked local 3d view 4-5  $\mu\text{m}$  on a side, large reconstructed objects comprise hundreds or thousands of such overlapping views, each with a corresponding embedding. To support object-level analyses, we therefore aggregated the set of embedding vectors for each cell prior to object-level classification or clustering (Fig. 3c). We found that a simple aggregation approach of computing the mean embedding value over each feature dimension was effective for the linear classification networks used here and the shallow two-layer ResNet classifiers used below.

Visualization of the whole-cell embedding space demonstrates good separation of ground truth neuronal and glial subtypes in the human dataset (Fig. 3d), even for subtypes that experts find difficult to distinguish, such as microglia versus oligodendrocyte precursor cells<sup>12</sup>. In leave-one-out cross-validation, a linear classifier achieves a 0.97 mean F1-Score for the six human cell types (balanced evaluation). In the mouse dataset, the cell types for which ground truth was available were primarily neurons. In this case, we could leverage the prior subcompartment classification (Fig. 2) to compute separate aggregated embeddings for the axonal and dendritic subcompartments of each cell (Fig. 3e-f). In leave-one-out cross-validation, a linear classifier achieves 0.933 mean F1-Score for the five mouse dendritic types, and 0.920

for the six mouse axon types, including a putative thalamocortical axon type for which no corresponding dendrites are present in the volume. Embeddings reliably distinguish putative thalamocortical axons from the axons of pyramidal cells and other neuron types (Fig. 3e), a capability that could enhance understanding of cortical circuitry and has not been previously demonstrated for automated methods.



**Figure 3. SegCLR clustering of subpopulations of neurons and glia.**

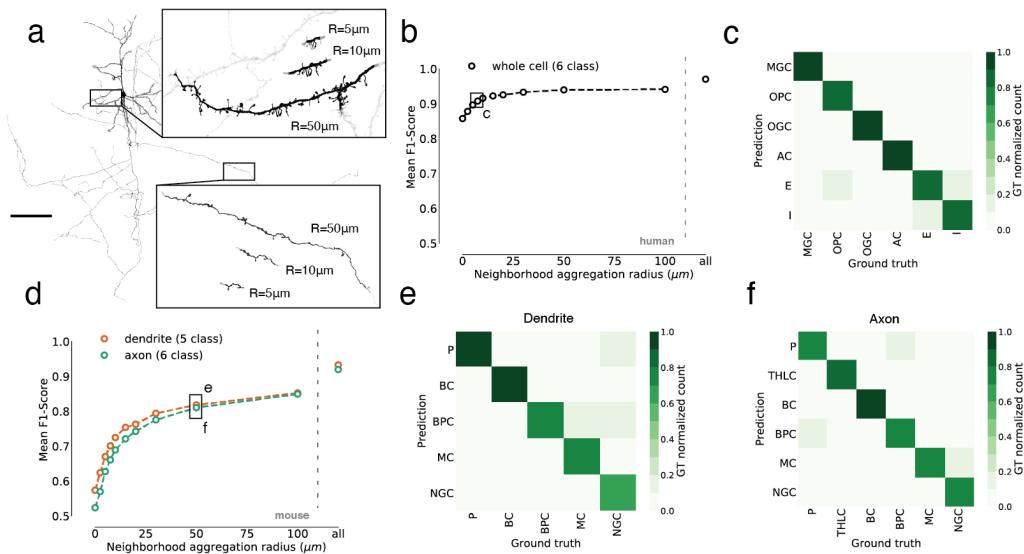
**a.** 3d renderings of representative neuron and glia cell types from the human dataset. Pyramidal cell axon is cut off to fit. Scale bars neuronal: 100 µm; glia: 25 µm. **b.** Renderings of representative neuron types from the mouse dataset. Pyramidal cell axon is cut off to fit. Scale bar 100 µm. **c.** Embedding representations for entire cells can be computed by aggregating all the local view embeddings across the cell. We used simple mean aggregation for classification experiments. **d.** UMAP visualized embedding clusters from human cortical cells. Each point represents the aggregated embedding for a whole cell. The points are colored according to the ground truth cell type labels, judged by expert visual inspection of the cells without reference to their embeddings. **e-f.** UMAP visualized embedding clusters from mouse neurons. When focusing on neuronal types, we considered the aggregated embeddings across the axonal (e) and dendritic (f) subcompartments separately. UMAP dimensions 2 and 3 show separation more clearly in this case than dimension 1 (not shown).

#### Locally aggregated embeddings for classification of cell fragments

Aggregated embeddings can reliably distinguish many cell types for large reconstructed cells (Fig. 3), and the human and mouse datasets each contain thousands of such cells. However, tools for reasoning about the millions of smaller cell fragments contained in these datasets

would also have many applications. For example, when tracing neural circuits a common goal is to identify the cell types of all synaptic partners upstream or downstream of the circuit, which requires reasoning about fragments in typical datasets where many cells pass out of the volume or are incompletely reconstructed.

To test whether embeddings aggregated from small cell fragments could reliably distinguish their cell types, we created artificial fragments within the ground truth labeled cells used above (Fig. 3) by restricting the aggregation window to a path length radius of 0-100  $\mu\text{m}$  along the cell skeleton (Fig. 4a). We then evaluated the cell type accuracy of a linear classifier trained on these features at different aggregation radii (Fig. 4b, d). On the human dataset, mean 6-class F1-Score including glial and neuronal cell types reached 0.916 at a radius of 10  $\mu\text{m}$ , with most residual error in distinguishing neuronal types (Fig. 4c).



**Figure 4. Locally aggregated embeddings for fragment classification.**

**a.** Starting from large reconstructed cells with ground truth labels, the embedding aggregation window is restricted to a local radius (insets: top, dendrite; bottom, axon) to evaluate the representational power of embeddings for smaller fragments. Scale bar 100  $\mu\text{m}$ . **b.** Classification performance for the six glia and neuronal types in the human ground truth set over different aggregation radii. Zero radius corresponds to a single unaggregated embedding node, while “all” corresponds to aggregation over the entire cell. Panel (c) uses the 10  $\mu\text{m}$  radius indicated by the black box. **c.** Confusion matrix for human cortex cell type classification with a 10  $\mu\text{m}$  aggregation radius. **d.** Classification accuracy for the mouse axon and dendrite aggregated embeddings. Panels (e-f) use the 50  $\mu\text{m}$  radius indicated by the black box. **e.** Confusion matrix for five mouse neuronal types in the dendrite ground truth set, with a 50  $\mu\text{m}$  aggregation radius. **f.** Confusion matrix for six mouse axon types.

**Human cell types.** MGC: microglia cell; OPC: oligodendrocyte precursor cell; OGC: oligodendrocyte cell; AC: amacrine cell; E: excitatory neuron; I: inhibitory interneuron.

**Mouse neuron types.** P: pyramidal; BC: basket cell; BPC: bipolar cell; MC: Martinotti cell; NGC: neurogliaform cell; THLC: thalamocortical axon.

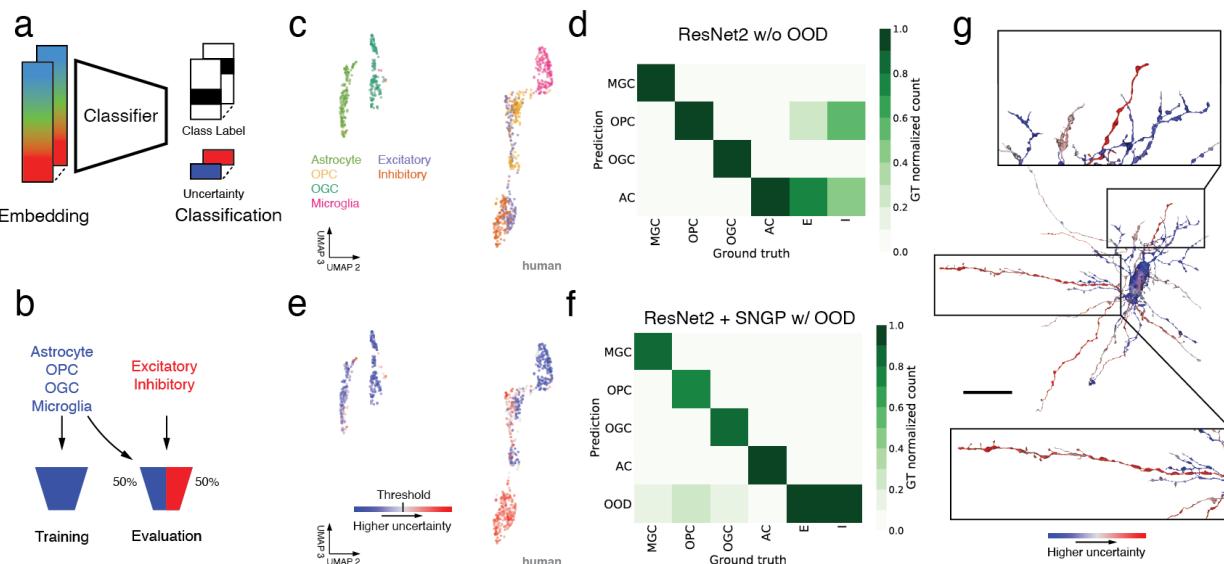
On the mouse dataset, we focused on neuronal cell types and considered the classification task separately for dendritic and axonal subcompartments as above. At an aggregation radius of 50  $\mu\text{m}$ , the 5-class dendritic F1-Score reached 0.818, while 6-class axonal F1-Score was 0.810 (Fig. 4d). Interestingly, residual error modes were largely distinct between dendritic versus axonal subcompartments (Fig. 4e-f). Overall, the results demonstrate that embeddings aggregated over small cellular fragments, which are common in most EM volumes, can support robust cell type classification. SegCLR embeddings could in the future be used to attempt to assign additional information to such cell fragments, such as molecular identities revealed by spatially resolved transcriptomics<sup>39</sup>.

#### Out-of-distribution input detection via Gaussian processes

A remaining issue with applications such as cell typing on large-scale datasets is how to gracefully handle image content that falls far outside the distribution of labeled examples. These “out-of-distribution” (OOD) input examples could include locations containing imaging artifacts or segmentation merge errors, or they could represent genuine biological structures that were simply absent in the training set. One example of the latter case is the common analysis situation in which only a few cell types have confident ground truth labels, but one wishes to classify these types while avoiding spurious classifications of a potentially large majority of surrounding segments belonging to diverse unknown types.

We addressed OOD inputs via Spectral-normalized Neural Gaussian Processes<sup>31</sup> (SNGP), which add a prediction uncertainty to the model output (Fig. 5a), and calibrate that uncertainty to reflect the distance between the test data and the training distribution. This allows OOD inputs to be detected and rejected, rather than spuriously classified, while requiring no extra labeling effort. To evaluate SNGP, we constructed a human cortical cell type dataset in which only the glial labeled types were used to train classifiers, while both glial and neuronal types were presented for testing (Fig. 5b). The neuronal types, making up 50% of the constructed test set, thus served as an OOD pool.

We first trained a small conventional network (i.e., lacking SNGP capabilities) on the glial classification task. Specifically, a shallow two-module ResNet classifier (“ResNet-2”) was trained on locally aggregated embeddings (radius 10  $\mu\text{m}$ ) from only the glial labeled cells. This network performed with high accuracy on the in-distribution glial half of the test set, but also (by construction) spuriously classified all OOD neuronal examples (Fig. 5d). The goal of OOD detection is to retain the strong in-distribution performance while selectively filtering out OOD inputs. To achieve this we converted the ResNet-2 to a SNGP by making drop-in substitutions to the trained network’s hidden and output layers<sup>31</sup>. These changes equip the classifier with an uncertainty output that estimates in part the degree to which each example is OOD with respect to the training distribution (Fig. 5e). The SNGP uncertainty can then be thresholded at a task-appropriate level to determine how aggressively to reject OOD inputs.



**Figure 5. Out-of-distribution input detection via Gaussian processes.**

**a.** We handled OOD inputs by computing prediction uncertainties alongside class labels, and calibrated the uncertainties to reflect the distance between the test example and the training distribution. **b.** To evaluate OOD detection, we trained classifiers on glial cell type labels, and then evaluated the classifiers on a 50/50 split between glial and OOD neuronal cell types. **c.** UMAP projection of locally aggregated embeddings (radius 10  $\mu\text{m}$ ) from the human cortical dataset, colored by ground truth labeled cell type. **d.** Confusion matrix for a ResNet-2 classifier trained on only the four glia types, with OOD neuronal examples mixed in at test time. While in-distribution glia classification is almost flawless, all neuronal examples (half the constructed test set) receive spurious classifications. **e.** As in (c) with the UMAP projected embeddings now colored by their SNGP uncertainty. The colormap transitions from shades of blue to shades of red at the threshold level used to reject OOD in our experiment. **f.** Confusion matrix for the SNGP-ResNet-2, assembled from five-fold cross-validations. Examples that exceed the uncertainty threshold are now treated as their own OOD predicted class. The strong in-distribution glia performance of (d) is largely retained, while neuronal fragments are effectively filtered out. **g.** Spatial distribution of local uncertainty over an unproofread segment that suffers from reconstruction merge errors between the central OPC glia and several neuronal fragments. The uncertainty signal distinguishes the merged neurites (red: high uncertainty / OOD) from the glia cell (blue: low uncertainty) with a spatial resolution of roughly the embedding aggregation distance. Scale bar 25  $\mu\text{m}$ .

The SNGP-ResNet-2 retains strong in-distribution glia classification performance while effectively filtering out OOD neuronal examples (Fig. 5f). With uncertainty thresholded at a level that optimizes F1 on a validation set and the resulting OOD examples treated as a separate class, overall mean F1-Score reaches 0.888. Note that the network layer substitutions for SNGP apply only to the small classifier network, with no modifications required to the underlying SegCLR embeddings. Furthermore, the neuronal ground truth labels used here were only needed to validate the results, while to train and deploy a classifier with SNGP OOD detection requires no extra ground truth labeling beyond the in-distribution set. These results show that SegCLR can be effectively deployed in common settings where only a small selection of expert labels may be available for application to a dataset containing a large proportion of OOD

objects. We also note that when used with SegCLR, OOD classification may prove simpler to implement than in typical non-embedding settings (see Methods).

Finally, we also evaluated the spatial distribution of local uncertainty over larger segments. This is particularly relevant for unproofread segments that contain reconstruction merge errors between a labeled and an OOD type. For example, the uncertainty of our SNGP classifier can distinguish neuronal fragments erroneously merged onto a central glia (Fig. 5g). Because the embeddings are aggregated with a local 10  $\mu\text{m}$  radius, the distinction between neuronal and glial branches of the segment can be resolved to within roughly 20  $\mu\text{m}$ . Automated merge error correction based on these branch distinctions<sup>12,22</sup>, potentially combined with direct detection of merge-specific features from embeddings at the merge point<sup>40</sup>, would be an attractive application for future investigation.

## Discussion

We have introduced SegCLR, a self-supervised method for training rich representations of local cellular morphology and ultrastructure, and demonstrated applications for biological annotation in human and mouse cortical volumes. Beyond the requirement for an accompanying instance segmentation, the current SegCLR formulation has some limitations. First, the 32-40 nm voxel resolution of input views impedes capture of finer EM ultrastructural features, such as vesicle subtypes or ciliary microtubule structure. The relatively large distances allowed between positive training pairs (up to 150  $\mu\text{m}$ ) also may bias embeddings toward larger-scale features; the current embeddings appear less effective at representing very local structures, such as dendritic spines, compared to preliminary results trained with shorter pair distances. Training SegCLR on higher resolution inputs, with shorter pair distances, or with multiscale capabilities, is worth detailed exploration.

Another limitation is that explicit input masking excludes EM context outside the current segment, while in some cases retaining surrounding context could be useful, e.g. for myelin sheaths, synaptic clefts, and synaptic partners. We therefore tested a version of SegCLR that receives the unmasked EM block and the segmentation mask as two separate input channels, rather than a single explicitly masked EM input. This variant performed similarly on the subcompartment classification task, but appeared more sensitive to subtle nonlinear photometric differences across the extent of the dataset. Given SegCLR's ability to reliably classify cell types that human experts find challenging to distinguish, including for small cellular fragments, exploration of attribution methods<sup>41</sup> to understand which SegCLR features are decisive could inform both future automated methods and human efforts.

SegCLR currently also focuses only on a single segment, whereas it could be useful to additionally learn representations of dual- or multi-cellular complexes. Training SegCLR on multi-segment inputs is conceptually straightforward, but running the network on all possible segment pairs or complexes within a large-scale dataset would be prohibitive. A strategy for limiting the set of targeted complexes, e.g. only using predefined synaptic partners, would be needed.

Finally, the local 4-5  $\mu\text{m}$  input field of view for each embedding could be considered a limitation; many tasks, including cell typing, are expected to benefit from larger contexts<sup>25,27,28</sup>. In the current work, we demonstrated that simple embedding aggregation strategies are sufficient to enable reasoning over larger contexts. More sophisticated aggregation methods using smoothed histograms<sup>42</sup>, or neural networks that process embedding point clouds<sup>43</sup>, skeleton graphs of embedding nodes<sup>25</sup>, or sparse voxel renderings of embeddings<sup>44</sup> could generate richer representations of larger contexts. There are also opportunities to extend representations beyond single cells. For example, neuronal embeddings could be extended with additional dimensions aggregated from pre- and post-synaptic partners, to create connectivity-enhanced cell type signatures or to form motif representations.

We expect SegCLR to be widely applicable across the growing breadth of volumetric EM datasets, and speculate that it could also apply in other settings where available segmentations can similarly serve as a guide to representation learning, e.g. brainbow fluorescence microscopy<sup>45,46</sup>, aerial and satellite imagery<sup>47</sup>, or dense cityscapes<sup>48</sup>. While the computational demands for training and deploying embeddings are considerable, this cost is incurred only once, rather than repeatedly for each downstream analysis. We will make available Python code for input preprocessing, network training, and evaluation, as well as pretrained TensorFlow network weights.

By providing rich and tractable representations of EM data, SegCLR greatly simplifies and democratizes downstream research and analysis. For example, the previous state of the art in subcompartment classification required millions of training examples assembled from thousands of manually validated segments, thousands of GPU hours to train a deep network classifier, and hundreds of thousands of CPU hours to evaluate over a large-scale dataset<sup>12,22</sup>. With SegCLR embeddings in hand, this benchmark is outperformed by a simple linear classifier, trained in minutes on a single CPU, with a few hundred manually labeled examples. We release full embedding datasets for the [human](#) and mouse cortical volumes to the community, to enhance exploration and understanding of these rich and complex resources.

## Methods

### Datasets

We used two large-scale serial-section EM connectomic datasets for SegCLR experiments: one from the human temporal cortex, imaged via SEM<sup>12</sup>; and one from the mouse visual cortex, imaged via TEM<sup>11</sup>. Both datasets are freely available and provide an aligned and registered EM volume with an accompanying automated dense instance segmentation. For SegCLR experiments we downsampled the human and mouse data to 32x32x33 nm and 32x32x40 nm nominal resolution respectively. The human EM volume was also CLAHE normalized<sup>49</sup>. We skeletonized both segmentation volumes as previously described<sup>12,32</sup>.

For subcompartment classification (Fig. 2), ground truth human data was collected on an earlier pre-release version of the dataset, as well as two smaller cutouts<sup>12</sup>. All three regions are contained within the publicly released dataset, although they differ slightly in their alignment and photometric normalization. For evaluation of local cell type classification (Fig. 4), it was important to have cells with minimal reconstruction merge errors in the ground truth labeled set. We therefore proofread some cells in the human dataset to exclude regions close to observed merge errors from the embeddings cell type analysis (Figs. 3-4). On the mouse dataset, we restricted analysis to the set of ground truth labeled cells that were already expert proofread prior to release<sup>11</sup>; at the time of publication this set consisted of neurons but not glia.

### Training SegCLR embedding networks

SegCLR was inspired by SimCLR<sup>24,35</sup>. The embedding network was a ResNet-18 architecture<sup>36</sup> implemented in TensorFlow, with convolutions extended to 3d and 3 bottleneck layers prior to a 64-dimensional embedding output. During training, we added 3 projection layers prior to a normalized temperature-scaled cross entropy (“NT-Xent”) loss<sup>35</sup> with temperature 0.1. We also found that downstream task performance was enhanced by adding a decorrelation term to the loss, defined as:

$$L_{\text{decorrelation}} = \frac{1}{d^2 - d} \sum_{i=1}^d \sum_{j \neq i} C_{ij}^2 \quad (1)$$

where  $d$  is the embedding dimensionality and  $C$  is the correlation matrix between embeddings over the batch. We trained SegCLR networks on 8x8 v2 Cloud TPUs for up to 500,000 steps with a full batch size of 512 examples and learning rate decay schedule starting at 0.2. Separate networks were trained for the human and mouse datasets.

The input to the network was a local 3d cutout of EM data 129 voxels on a side, nominally corresponding to 4128x4128x4257 nm in the human EM dataset, and 4128x4128x5160 nm in the mouse dataset. We then masked the EM data by the segmentation for the object at the center of the field of view, so that possible confounds in the surrounding EM context would be excluded.

We also leveraged the segmentation and corresponding skeletonization to generate example pairs for contrastive training. For an arbitrary segment, we picked a 3d view to be centered on

an arbitrary skeleton node, and then picked a positive pair location centered on a second node within 150  $\mu\text{m}$  path length away on the same skeleton. Positive pairs were preprocessed before training for higher performance. As there are more possible pairs for larger distances, we sorted these positive pairs into four distance buckets from which we drew uniformly. The bucket boundaries were (0, 2500, 10000, 30000, 150000) nanometers.

As in SimCLR<sup>24</sup>, we used the 510 examples from the rest of the batch, which were drawn from 255 other segments throughout the volume, as negative pairs. We also applied random reflections, and photometric augmentations<sup>19</sup> to the inputs, to prevent the network from solving the contrastive task via trivial cues such as the orientation of processes, or the local voxel statistics.

### SegCLR inference

We inferred SegCLR embeddings over the full extent of the human and mouse datasets. After removing trivial segments, we extracted local 3d views centered on skeleton nodes for the remaining segments, with approximately 1500 nm path length spacing. The set of views for each segment therefore had substantial overlap of about 65-70% with typical nearest neighbor views. We then ran SegCLR on all selected views (3.9 billion and 4.2 billion for the human and mouse datasets respectively) via an Apache Beam Python pipeline running on a large CPU cluster, and stored the resulting embedding vectors keyed by segment ID and spatial XYZ coordinates. For visualization, we ran UMAP dimensionality reduction<sup>37</sup> to 2-4 dimensions on representative samplings or on subsets of interest from among the embeddings. When sampling from a large population of embeddings of local cutouts, we sampled such that all classes were represented with a significant number of examples.

### Subcompartment classification

We trained linear classifiers to identify subcompartment types from embedding inputs based on expert ground truth labels on each dataset (Table 1). For comparison, we also trained a fully-supervised subcompartment classifier directly on voxel inputs using an identical 3d ResNet-18 architecture and input configuration (photometric augmentation was omitted and random 3d rotations were added), and replaced the bottleneck, projection heads, and contrastive loss with a classification softmax and cross-entropy loss. The supervised network was trained on 8 GPUs with a total batch size of 64 via stochastic gradient descent with learning rate 0.001 for 1.5 M steps. During full supervised training, examples were rebalanced class-wise by upsampling all classes to match the most numerous class. However, in SegCLR experiments (Fig. 2) we showed that performance using embeddings was robust to substantial reductions in training data via random sampling of examples. During subsampling we ensured that every class was represented with at least 10% of the examples. We repeated each sampling round 20 times for sample sizes  $\leq 5000$ .

dataset	subcompartment	#training segments	total #training examples
human	axon	1,256	262,395
	dendrite	933	1,151,175
	soma	993	40,649
	astrocytic process	212	1,392,702
mouse	axon	121	4,653,326
	dendrite		2,343,169
	soma		13,830

**Table 1. Cellular subcompartment ground truth label sets.**

The number of distinct segments and total number of training examples of each type available for subcompartment classification (Fig. 2). For fully supervised training on the human dataset, the full set of training examples was used, while for linear classifiers trained on SegCLR embeddings, both full and reduced subsets of the training set were evaluated (Fig. 2d).

### Embedding aggregation

For object level analyses we aggregated local embeddings, either across entire cells (Fig. 3), or across entire axonal and dendritic subcompartments (Fig. 3), or within more limited 0-100  $\mu$ m radius windows along a cell's skeleton path length (Figs. 4-5). A simple aggregation method of taking the mean embedding value over each feature dimension performed well when using linear classifiers for downstream tasks (Figs. 2-4), as well as for the shallow ResNet-2 used for out-of-distribution input detection (Fig. 5).

### Cell type classification

We tested cell type classification via embedding fingerprints using a set of expert ground truth labeled neurons and glia from both human and mouse (Table 2). These ground truth cells are generally large and contain somas within the volume, but they are not completely reconstructed. We manually proofread human ground truth cells for merge errors by marking bad agglomeration edges in the agglomeration graph prior to evaluation. In the mouse dataset, we restricted analysis to proofread neurons included in the public v117 release, with further updates to the available set of labeled and proofread cells anticipated. We trained linear classifiers to predict cell subtypes (Fig. 3), and evaluated classification performance via leave-one-out cross-validation.

dataset	cell supertype	cell subtype and abbreviation	#segments	#embeddings total	
human	neuron	Excitatory (E)	159	863,644	
		Inhibitory (I)	52	91,283	
	glia	Microglia cell (MGC)	36	29,536	
		Oligodendrocyte precursor cell (OPC)	19	23,614	
		Oligodendrocyte glia cell (OGC)	17	4,449	
		Amacrine cell (AC)	45	1,394,056	
	mouse	Pyramidal cell (P)	78	1,117,794	
mouse		Basket cell (BC)	48	850,642	
		Bipolar cell (BPC)	29	114,022	
		Martinotti cell (MC)	11	160,122	
		Neurogliaform cell (NGC)	5	27,760	
		Putative thalamocortical axon (THLC)	9	73,385	

**Table 2. Cell type ground truth label sets.**

The number of distinct segments and total number of embeddings available for cell type classification (Figs. 3-5).

To test cell type classification of smaller cell fragments, we artificially broke the ground truth labeled cells into smaller pieces by restricting the embedding aggregation window to a specific  $\mu\text{m}$  radius along the skeleton path length from arbitrary starting points. We evaluated fragment classification performance by partitioning all cells into 5 buckets with similar type distributions between buckets. During evaluation we trained on embeddings drawn from the cells in four buckets and tested on embeddings drawn from the fifth bucket. We report the mean F1-Score from these five runs. Within each run we compute the mean F1-Score of all classes. Overall we sampled 50,000 embeddings across all buckets. During training and testing, imbalances between classes were balanced by repeating examples from minority classes. We filtered out cell fragments with fewer than  $1/R_{\mu\text{m}}$  (1 embedding node per micrometer radius) to exclude cases where artificially cut skeletons did not allow for a full aggregation of the embedding nodes.

For the mouse dataset, we first classified all embedding nodes as axon, soma, dendrite with a linear classifier trained on the entire subcompartment training set. We then completed the evaluation outlined above for dendrite and axon individually. For each category (axon, dendrite), we excluded examples where less than 90% of the embedding nodes in the aggregation window were classified as that category. Such cases usually represented scenarios were the

segmentation between the axon and a dendrite of the same neuron were merged and the skeleton falsely crossed between the compartments.

#### Cell-type ground truth

The different neuronal types in the mouse dataset were classified based on the morphological and synaptic criteria as described<sup>50</sup>. Pyramidal cells were identified by the presence of a spiny apical dendrite radiating toward the pia, spiny basal dendrites and an axon that formed asymmetric synapses. Putative thalamic axons also formed asymmetric synapses and though their soma was located outside the reconstructed volume, their gross morphology resembles previously described thalamic arbors<sup>51</sup> and their fine morphology at the ultrastructure levels<sup>52</sup>. Neuronal cells were classified as inhibitory interneurons if their axon formed symmetric synapses. Inhibitory interneurons were further assigned subtypes using their synaptic connectivity and the morphology of axons and dendrites<sup>53–55</sup>. Basket cells were identified by having a larger number of primary dendrites, and at least 12% of their postsynaptic targets were pyramidal cell somata. Martinotti cells were identified by its apical axon that projected to cortical layer 1 and targeted mostly distal dendritic shafts and spines of excitatory cells; consistent with<sup>53</sup>. Martinotti cells also were characterized by having a multipolar dendritic arbor that was usually spiny. Bipolar cells usually had two to three primary dendrites. The dendrites were usually spiny and showed a vertical bias. Neurogliaform cells were often (but not exclusively) in cortical layer 1. The axons of neurogliaform cells have usually lower density of synapses; consistent with<sup>53</sup>. Neurogliaform cells also had a large number of primary dendrites and have a pattern of synaptic inputs when compared with other inhibitory cell types.

#### Out-of-distribution input detection

We detected OOD (OOD) inputs via Spectral-normalized Neural Gaussian Processes<sup>31</sup> (SNGPs). As a baseline we trained a shallow ResNet-2 (2 ResNet modules) to classify glial cell types of 10  $\mu\text{m}$  radius fragments. We then modified the ResNet-2 by spectrally normalizing its hidden layers and replacing the output layer with a Gaussian process, using the SNGP package in TensorFlow. We used the BERT hyperparameter setting from Liu et al.<sup>31</sup> for our analysis.

We evaluated performance on a test set constructed of a 50/50 split between glia and OOD neuronal examples. First, we computed prediction uncertainty estimates for the test networks using the Dempster-Shafer metric<sup>31</sup>:

$$u(\mathbf{x}_i) = \frac{K}{K + \sum_{k=1}^K \exp(h_k(\mathbf{x}_i))} \quad (2)$$

where  $K$  is the number of classes (in our case: 4) and  $h_k$  are the classification logits (network outputs prior to probability normalization). As suggested in the SNGP tutorial, we replaced the Monte Carlo estimation of the output logits with the mean-field method<sup>56</sup>:

$$h_\lambda(\mathbf{x}_i) = \frac{h(\mathbf{x}_i)}{\sqrt{1 + \lambda\sigma^2(\mathbf{x})}} \quad (3)$$

using  $\lambda = 3/\pi^2$  and variances estimated with the GP module.

For the evaluation, we repeated the 5-fold cross-validation as outlined for the cell type fragment classification. We sampled 100,000 unique examples and upsampled them such that each glia class accounted for 12.5% and each OOD class (excitatory, inhibitory) accounted for 25%. During each fold we trained a new classifier on a subset of the glia examples and then predicted the examples in the hold-out test set as well as the set of neuronal fragments which we reused for all 5 folds. For each fold we found the uncertainty threshold that maximized the F1-Score of the in-distribution vs out-of-distribution task. For this, we set aside half the examples from the test-set which were then not used for calculating the scores. For each fold, we replaced the original class prediction (1 of 4 glia classes) with the OOD class when the uncertainty for an example exceeded this threshold. Finally, we calculated F1-Scores for each of the 5 classes after scaling them such that OOD examples accounted for 50% and averaged them for a final F1-Score per fold. We reported the mean F1-Score across the five folds.

### Embedding datasets

We make the embeddings for the mouse and human datasets available as sharded csv files.

mouse: [gs://iarpa\\_microns/minnie/minnie65/embeddings/segclr\\_csvzips/README](gs://iarpa_microns/minnie/minnie65/embeddings/segclr_csvzips/README)

human: [gs://h01-release/data/20220326/c3/embeddings/segclr\\_csvzips/README](gs://h01-release/data/20220326/c3/embeddings/segclr_csvzips/README)

## **Acknowledgements**

We are grateful to Simon Kornblith, Geoffrey Hinton, Chen Sun, Dilip Krishnan, and Balaji Lakshminarayanan for helpful discussions. ALB, FC, CSSM and NMdC were supported by the Intelligence Advanced Research Projects Activity (IARPA) of the Department of Interior/Interior Business Center (Dol/IBC) through contract number D16PC00004; and by Allen Institute for Brain Science. The views and conclusions contained herein are those of the authors and should not be interpreted as representing the official policies or endorsements, either expressed or implied, of the funding sources including IARPA, Dol/IBC, or the U.S. Government. ALB, FC, CSSM and NMdC wish to thank the Allen Institute founder, Paul G. Allen, for his vision, encouragement, and support.

## References

1. Ascoli, G. A., Donohue, D. E. & Halavi, M. NeuroMorpho.Org: A Central Resource for Neuronal Morphologies. *J. Neurosci.* **27**, 9247–9251 (2007).
2. DeFelipe, J. *et al.* New insights into the classification and nomenclature of cortical GABAergic interneurons. *Nat. Rev. Neurosci.* **14**, 202–216 (2013).
3. Hodge, R. D. *et al.* Conserved cell types with divergent features in human versus mouse cortex. *Nature* **573**, 61–68 (2019).
4. Nandi, A. *et al.* Single-neuron models linking electrophysiology, morphology and transcriptomics across cortical cell types. *bioRxiv* 2020.04.09.030239 (2020) doi:10.1101/2020.04.09.030239.
5. Kandel, E. R., Jessell, T. M. & Siegelbaum, S. A. *Principles of Neural Science, Sixth Edition*. (McGraw Hill Professional, 2021).
6. Li, F. *et al.* The connectome of the adult Drosophila mushroom body provides insights into function. *Elife* **9**, (2020).
7. Hulse, B. K. *et al.* A connectome of the Drosophila central complex reveals network motifs suitable for flexible navigation and context-dependent action selection. *eLife* vol. 10 (2021).
8. Eichler, K. *et al.* The complete connectome of a learning and memory centre in an insect brain. *Nature* vol. 548 175–182 (2017).
9. Motta, A. *et al.* Dense connectomic reconstruction in layer 4 of the somatosensory cortex. *Science* **366**, (2019).
10. Lichtman, J. W. & Denk, W. The Big and the Small: Challenges of Imaging the Brain's Circuits. *Science* vol. 334 618–623 (2011).
11. MICrONS Consortium *et al.* Functional connectomics spanning multiple areas of mouse visual cortex. *bioRxiv* 2021.07.28.454025 (2021) doi:10.1101/2021.07.28.454025.
12. Shapson-Coe, A. *et al.* A connectomic study of a petascale fragment of human cerebral

cortex. (2021) doi:10.1101/2021.05.29.446289.

13. Lee, K. *et al.* Convolutional nets for reconstructing neural circuits from brain images acquired by serial section electron microscopy. *Curr. Opin. Neurobiol.* **55**, 188–198 (2019).
14. Funke, J. *et al.* Large Scale Image Segmentation with Structured Loss Based Deep Learning for Connectome Reconstruction. *IEEE Trans. Pattern Anal. Mach. Intell.* **41**, 1669–1680 (2019).
15. Berning, M., Boergens, K. M. & Helmstaedter, M. SegEM: efficient image analysis for high-resolution connectomics. *Neuron* **87**, 1193–1206 (2015).
16. Meirovitch, Y. *et al.* Cross-Classification Clustering: An Efficient Multi-Object Tracking Technique for 3-D Instance Segmentation in Connectomics. in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* 8425–8435 (2019).
17. Lee, K., Zung, J., Li, P., Jain, V. & Sebastian Seung, H. Superhuman Accuracy on the SNEMI3D Connectomics Challenge. *arXiv [cs.CV]* (2017).
18. Januszewski, M. *et al.* High-precision automated reconstruction of neurons with flood-filling networks. *Nat. Methods* **15**, 605–610 (2018).
19. Dorkenwald, S. *et al.* Automated synaptic connectivity inference for volume electron microscopy. *Nat. Methods* **14**, 435–442 (2017).
20. Schubert, P. J., Dorkenwald, S., Januszewski, M., Jain, V. & Kornfeld, J. Learning cellular morphology with neural networks. *Nat. Commun.* **10**, 2736 (2019).
21. Kornfeld, J. M. *et al.* An anatomical substrate of credit assignment in reinforcement learning. *bioRxiv* 2020.02.18.954354 (2020) doi:10.1101/2020.02.18.954354.
22. Li, H., Januszewski, M., Jain, V. & Li, P. H. Neuronal Subcompartment Classification and Merge Error Correction. *bioRxiv* 2020.04.16.043398 (2020) doi:10.1101/2020.04.16.043398.
23. Devlin, J., Chang, M.-W., Lee, K. & Toutanova, K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. (2018).

24. Chen, T., Kornblith, S., Norouzi, M. & Hinton, G. A Simple Framework for Contrastive Learning of Visual Representations. in *International Conference on Machine Learning* 1597–1607 (PMLR, 2020).
25. Weis, M. A., Pede, L., Lüddecke, T. & Ecker, A. S. Self-supervised Representation Learning of Neuronal Morphologies. (2021).
26. Zhao, T. & Plaza, S. M. Automatic Neuron Type Identification by Neurite Localization in the *Drosophila* Medulla. *arXiv [q-bio.NC]* (2014).
27. Li, Y., Wang, D., Ascoli, G. A., Mitra, P. & Wang, Y. Metrics for comparing neuronal tree shapes based on persistent homology. *PLoS One* **12**, e0182184 (2017).
28. Kanari, L. *et al.* Objective Morphological Classification of Neocortical Pyramidal Cells. *Cereb. Cortex* **29**, 1719–1735 (2019).
29. Huang, G. B., Yang, H.-F., Takemura, S.-Y., Rivlin, P. & Plaza, S. M. Latent Feature Representation via Unsupervised Learning for Pattern Discovery in Massive Electron Microscopy Image Volumes. *arXiv [cs.CV]* (2020).
30. Azizi, S. *et al.* Big Self-Supervised Models Advance Medical Image Classification. *arXiv [eess.IV]* (2021).
31. Liu, J. Z. *et al.* Simple and Principled Uncertainty Estimation with Deterministic Deep Learning via Distance Awareness. (2020).
32. Silversmith, W., Alexander Bae, J., Li, P. H. & Wilson, A. M. *seung-lab/kimimaro: Zenodo Release v1*. (2021). doi:10.5281/zenodo.5539913.
33. Buhmann, J. *et al.* Automatic Detection of Synaptic Partners in a Whole-Brain *Drosophila* EM Dataset. *bioRxiv* 2019.12.12.874172 (2019).
34. Heinrich, L. *et al.* Automatic whole cell organelle segmentation in volumetric electron microscopy. *Cold Spring Harbor Laboratory* 2020.11.14.382143 (2020) doi:10.1101/2020.11.14.382143.
35. Chen, T., Kornblith, S., Swersky, K., Norouzi, M. & Hinton, G. Big Self-Supervised Models

are Strong Semi-Supervised Learners. *arXiv [cs.LG]* (2020).

36. He, K., Zhang, X., Ren, S. & Sun, J. Deep Residual Learning for Image Recognition. *arXiv [cs.CV]* (2015).
37. McInnes, L., Healy, J. & Melville, J. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *arXiv [stat.ML]* (2018).
38. Lin, Z. *et al.* Two Stream Active Query Suggestion for Active Learning in Connectomics. (2020).
39. Close, J. L., Long, B. R. & Zeng, H. Spatially resolved transcriptomics in neuroscience. *Nat. Methods* **18**, 23–25 (2021).
40. Rolnick, D. *et al.* Morphological Error Detection in 3D Segmentations. *arXiv [cs.CV]* (2017).
41. Sundararajan, M., Taly, A. & Yan, Q. Axiomatic Attribution for Deep Networks. *arXiv [cs.LG]* (2017).
42. Kearnes, S., McCloskey, K., Berndl, M., Pande, V. & Riley, P. Molecular graph convolutions: moving beyond fingerprints. *J. Comput. Aided Mol. Des.* **30**, 595–608 (2016).
43. Qi, C. R., Su, H., Mo, K. & Guibas, L. J. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. *arXiv [cs.CV]* (2016).
44. Graham, B., Engelcke, M. & van der Maaten, L. 3D Semantic Segmentation With Submanifold Sparse Convolutional Networks. in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* 9224–9232 (2018).
45. Duan, B. *et al.* Unsupervised Neural Tracing in Densely Labeled Multispectral Brainbow Images. *bioRxiv* 2020.06.07.138941 (2020) doi:10.1101/2020.06.07.138941.
46. Shen, F. Y. *et al.* Light microscopy based approach for mapping connectivity with molecular specificity. *Nat. Commun.* **11**, 4632 (2020).
47. Guérin, E., Oechslin, K., Wolf, C. & Martinez, B. Satellite Image Semantic Segmentation. *arXiv [cs.CV]* (2021).
48. Cordts, M. *et al.* The Cityscapes Dataset for Semantic Urban Scene Understanding. *arXiv*

[cs.CV] (2016).

49. Zuiderveld, K. Contrast Limited Adaptive Histogram Equalization. in *Graphics Gems IV* (ed. Heckbert, P. S.) 474–485 (Academic Press Professional, Inc., 1994).
50. Turner, N. L. *et al.* Reconstruction of neocortex: Organelles, compartments, cells, circuits, and activity. *Cell* (2022) doi:10.1016/j.cell.2022.01.023.
51. Antonini, A., Fagiolini, M. & Stryker, M. P. Anatomical correlates of functional plasticity in mouse visual cortex. *J. Neurosci.* **19**, 4388–4406 (1999).
52. Bopp, R., Holler-Rickauer, S., Martin, K. A. C. & Schuhknecht, G. F. P. An Ultrastructural Study of the Thalamic Input to Layer 4 of Primary Motor and Primary Somatosensory Cortex in the Mouse. *J. Neurosci.* **37**, 2435–2448 (2017).
53. Kawaguchi, S.-Y. & Hirano, T. Integrin  $\alpha 3\beta 1$  suppresses long-term potentiation at inhibitory synapses on the cerebellar Purkinje neuron. *Molecular and Cellular Neuroscience* vol. 31 416–426 (2006).
54. Kubota, Y., Karube, F., Nomura, M. & Kawaguchi, Y. The Diversity of Cortical Inhibitory Synapses. *Frontiers in Neural Circuits* vol. 10 (2016).
55. Kubota, Y. Untangling GABAergic wiring in the cortical microcircuit. *Current Opinion in Neurobiology* vol. 26 7–14 (2014).
56. Lu, Z., Ie, E. & Sha, F. Mean-Field Approximation to Gaussian-Softmax Integral with Application to Uncertainty Estimation. *arXiv [cs.LG]* (2020).