

NEURD: automated proofreading and feature extraction for connectomics

Brendan Celii^{1,2,9,*}, Stelios Papadopoulos^{1,2}, Zhuokun Ding^{1,2}, Paul G. Fahey^{1,2}, Eric Wang^{1,2}, Christos Papadopoulos^{1,2}, Alexander B. Kunin^{1,2,8}, Saamil Patel^{1,2}, J. Alexander Bae^{4,5}, Agnes L. Bodor³, Derrick Brittain³, JoAnn Buchanan³, Daniel J. Bumbarger³, Manuel A. Castro⁴, Erick Cobos^{1,2}, Sven Dorkenwald^{4,6}, Leila Elabbady³, Akhilesh Halageri⁴, Zhen Jia^{4,6}, Chris Jordan⁴, Dan Kapner³, Nico Kemnitz⁴, Sam Kinn³, Kisuk Lee^{4,7}, Kai Li⁶, Ran Lu⁴, Thomas Macrina^{4,6}, Gayathri Mahalingam³, Eric Mitchell⁴, Shanka Subhra Mondal^{4,5}, Shang Mu⁴, Barak Nehoran^{4,6}, Sergiy Popovych^{4,6}, Casey M. Schneider-Mizell³, William Silversmith⁴, Marc Takeno³, Russel Torres³, Nicholas L. Turner^{4,6}, William Wong⁴, Jingpeng Wu⁴, Szi-chieh Yu⁴, Wenjing Yin³, Daniel Xenos¹⁵, Lindsey M. Kitchell¹⁵, Patricia K. Rivlin¹⁵, Victoria A. Rose¹⁵, Caitlyn A. Bishop¹⁵, Brock Wester¹⁵, Emmanouil Froudarakis^{1,2,12}, Edgar Y. Walker^{10,11}, Fabian Sinz^{1,2,13,14}, H. Sebastian Seung⁴, Forrest Collman³, Nuno Maçarico da Costa³, R. Clay Reid³, Xaq Pitkow^{1,2,9}, Andreas S. Tolias^{1,2,9}, and Jacob Reimer^{1,2,9}

¹Center for Neuroscience and Artificial Intelligence, Baylor College of Medicine, Houston, USA

²Department of Neuroscience, Baylor College of Medicine, Houston, USA

³Allen Institute for Brain Science, Seattle, USA

⁴Princeton Neuroscience Institute, Princeton University, Princeton, USA

⁵Electrical and Computer Engineering Department, Princeton University, Princeton, USA

⁶Computer Science Department, Princeton University, Princeton, USA

⁷Brain & Cognitive Sciences Department, Massachusetts Institute of Technology, Cambridge, USA

⁸Department of Mathematics, Creighton University, Omaha, USA

⁹Department of Electrical and Computer Engineering, Rice University, Houston, USA

¹⁰Department of Physiology and Biophysics, University of Washington, Seattle, USA

¹¹UW Computational Neuroscience Center, University of Washington, Seattle, USA

¹²Institute of Molecular Biology and Biotechnology, Foundation for Research and Technology Hellas, Heraklion, Greece

¹³Institute for Bioinformatics and Medical Informatics, University Tübingen, Tübingen, Germany

¹⁴Institute for Computer Science, University Göttingen, Göttingen, Germany

¹⁵Research Exploratory Development Department, Johns Hopkins University Applied Physics Laboratory, Baltimore, USA

*first author

We are now in the era of millimeter-scale electron microscopy (EM) volumes collected at nanometer resolution (Shapson-Coe et al., 2021; Consortium et al., 2021). Dense reconstruction of cellular compartments in these EM volumes has been enabled by recent advances in Machine Learning (ML) (Lee et al., 2017; Wu et al., 2021; Lu et al., 2021; Macrina et al., 2021). Automated segmentation methods can now yield exceptionally accurate reconstructions of cells, but despite this accuracy, laborious post-hoc proofreading is still required to generate large connectomes free of merge and split errors. The elaborate 3-D meshes of neurons produced by these segmentations contain detailed morphological information, from the diameter, shape, and branching patterns of axons and dendrites, down to the fine-scale structure of dendritic spines. However, extracting information about these features can require substantial effort to piece together existing tools into custom workflows. Building on existing open-source software for mesh manipulation, here we present "NEURD", a software package that decomposes each meshed neuron into a compact and extensively-annotated graph representation. With these feature-rich graphs, we implement workflows for state of the art automated post-hoc proofreading of merge errors, cell classification, spine detection, axon-dendritic proximities, and other features that can enable many downstream analyses of neural morphology and connectivity. NEURD can make these new massive and complex datasets more accessible to neuroscience researchers focused on a variety of scientific questions.

EM connectomics | Neural Morphology | Automated Proofreading | Neural Annotation

Correspondence: reimer@bcm.edu

Introduction

To understand the morphological features of individual neurons and the principles governing their connectivity, the use of large-scale electron microscopy and reconstruction of entire neural circuits is becoming increasingly routine. For example, in the past year, the MICrONS Consortium published a millimeter-scale open-source dataset of mouse visual cortex (Consortium et al., 2021) (approximately 80,000 neurons and 500 million synapses; "MICrONS dataset"), and a team at Harvard published a similar reconstructed volume of human temporal lobe (Shapson-Coe et al., 2021) (approximately 15,000 neurons, 130 million synapses; "H01 dataset"). These reconstructions offer opportunities for analysis of neural morphology and synaptic connectivity at a scale that was previously inaccessible. However, effectively using these massive and complex datasets for scientific discovery requires a new ecosystem of software tools.

Here, we describe NEURD — short for "NEURAl Decomposition" — a software package in Python that extracts useful information from 3-D neuronal meshes and implements workflows for a variety of downstream applications including automated proofreading, morphological analysis and connectomic studies. NEURD decomposes the 3-D meshes of neurons from EM reconstructions into a richly-annotated graph representation with many pre-computed features. These graphs characterize the neuron at the level of non-branching segments in the axonal and dendritic arbor, a relatively coarse level of representation that enables a wide variety of queries about the statistics and geometry of neural wiring.

We begin by demonstrating the utility of this framework in an automated proofreading pipeline that is highly effective at correcting merge errors using heuristic rules. We evaluate the performance of this pipeline against human proofreading both at the level of cleaning merge errors from entire axonal and dendritic arbors, and at the level of generating split points for each error that can be executed in a hierarchical segmentation graph (Dorkenwald et al., 2022a,b) to batch-correct large numbers of errors. Because proofreading is performed by a set of interpretable heuristic rules, we show how the feedback from this validation process can be used to tune specific parameters in order to achieve even higher accuracy on certain kinds of edits. Our proofreading workflow is conservative; it is biased towards a higher number of false negative connections but low false positives, resulting in a high-fidelity but sparsely-sampled connectivity graph.

We then show how the pre-computed features extracted by NEURD can enable us to recapitulate and extend a variety of previous observations about neural morphology and geometry, taking advantage of the thousands of reconstructed neurons spanning all layers in these volumes. These diverse features include everything from spine annotations to the branching geometry of axonal and dendritic arbors, enabling powerful queries across multiple spatial scales. Using combinations of these features, we show that it is possible to distinguish excitatory and inhibitory neurons in both the mouse and human datasets with high accuracy, and to classify neurons into excitatory and inhibitory sub-classes (basket, bipolar, etc.), even when using isolated fragments of dendrites.

Finally, we examine the potential of this highly-annotated, sparse but clean connectivity graph to yield novel scientific insights about neural circuit connectivity, including higher order motifs. NEURD includes a fast workflow to identify close axonal-dendritic proximities that can serve as a null distribution for analyses focused on understanding the specificity of synapse formation between different cells, even if their axon and dendritic arbors are incomplete. We use this capability to show that the ratio of synapses to proximities (synaptic conversion rate) varies as expected across cell-type-specific connections, and we show that our automatic proofreading enables us to identify a set of extremely rare multi-synaptic connections with four or more synapses between pairs of excitatory cells. We find that these pairs have more similar functional properties *in vivo*, as predicted by a principle of "like-to-like" connectivity in the mouse visual cortex (Lee et al., 2016; Ko et al., 2011; Ding et al., 2023).

With NEURD, we set out to achieve two goals: The first goal is to extract a large set of useful features that many downstream analyses have in common, and to represent them in a compact, intuitive, and manageable format. Having ready access to these extracted features will be a valuable resource and enable many analyses beyond the ones we consider here. The second goal is to use these features to perform automated proofreading of merge errors, removing the large majority of these confounding errors in the volume and producing a high-fidelity, sparsely-sampled connectome that can be paired with appropriate proximity controls to study principles of neural connectivity. A useful analogy may be drawn with software tools that have supported the widespread adoption of other complex data modalities such as population calcium imaging (CaImAn, Suite2P; Giovannucci et al. 2019; Pachitariu et al. 2017), neuropixel recordings (KiloSort, MountainSort; Pachitariu et al. 2023; Chung et al. 2017), label-free behavioral tracking (DeepLabCut, MoSeq, SLEAP; Mathis et al. 2018; Pereira et al. 2022; Markowitz et al. 2018), and spatial transcriptomics (Giotto, Squidpy; Dries et al. 2021; Palla et al. 2022). Like these packages, the goal of NEURD is to make "big neuroscience data" (in this case, large-scale EM reconstructions) accessible to a larger community, by providing a common pre-processing and feature extraction workflow that can support a wide variety of downstream analyses. As more large-scale EM reconstructions become available, tools like NEURD will become increasingly essential for exploring principles of neural organization across multiple species.

Results

Summary of large-scale dense EM reconstructions. Data collection for the MICrONS and H01 dataset has been described in previous publications (Consortium et al., 2021; Shapson-Coe et al., 2021). The tissue preparation, slicing procedure, and imaging resolution (8nm x 8nm x 30-40nm) was roughly similar in both cases: Both volumes were stained with heavy metals (Yin et al., 2020; Tapia et al., 2012), embedded in resin, and cut on an automatic tape-collecting ultramicrotome (ATUM) (Schalek et al., 2011; Own et al., 2015). However, the imaging and reconstruction workflows for the two datasets were very different. The MICrONS volume was collected with transmission electron microscopy (TEM) (Yin et al., 2020), while the H01 volume was collected with scanning electron microscopy (SEM) (Hayworth et al., 2014), and slice alignment and dense reconstruction for the two datasets were performed via different reconstruction pipelines (Macrina et al., 2021; Consortium et al., 2021; Shapson-Coe et al., 2021). However, despite the differences in source tissue and significant differences in the acquisition workflows for MICrONS and H01, in the end all volumetric reconstructions produce similar 3-D meshes as a common data product downstream of the segmentation process. The capabilities of NEURD are focused at the level of these mesh representations, which capture rich information about the microscale anatomy of neurons that can be useful for a variety of downstream analyses, including comparative analyses of neural circuitry across species, volumes, and reconstruction methods.

Preprocessing of Neuronal Meshes. EM reconstructions yield neural meshes with varying levels of completeness, and with different kinds of merge errors (Fig. 1b-e). Merge errors include multiple whole neurons connected together (Fig. 1c), and disconnected pieces of neurite ("orphan neurites") merged onto different neural compartments (Fig. 1d). Merge errors may also

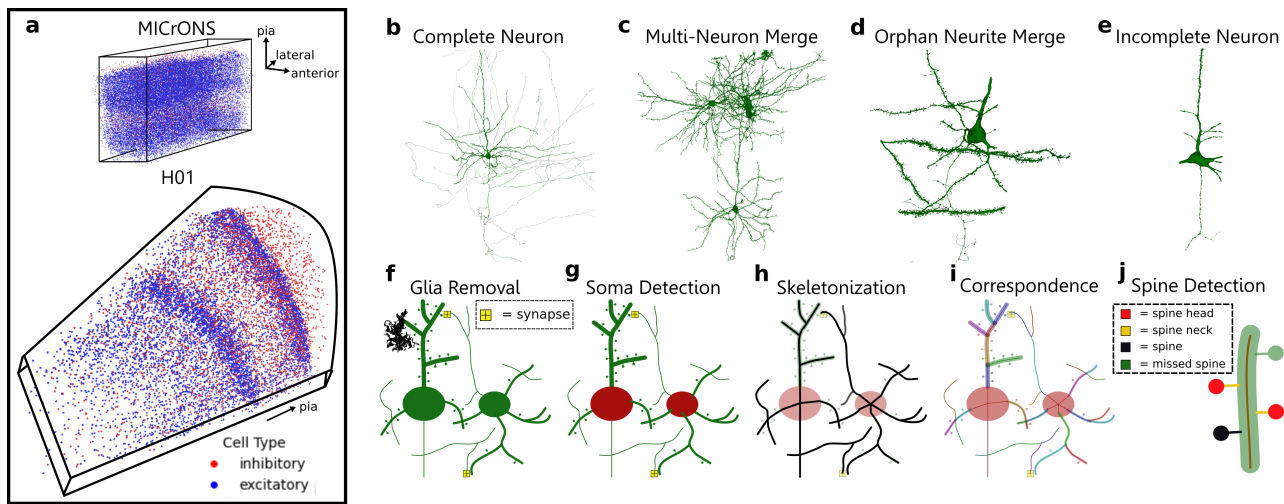


Fig. 1. Working with neuronal meshes from large-volume EM segmentations **a** The portion of the MICrONS volume that we analyzed (65% of the entire reconstruction) is an approximately $1300 \times 820 \times 520 \mu\text{m}^3$ rectangular volume from mouse visual cortex, while H01 is a wedge-shaped volume with a longest dimension of 3mm, a width of 2mm, and a thickness of $150 \mu\text{m}$ from human temporal cortex. Panels **b-e** illustrate the range of accuracy across neural reconstructions. **b** Example of a nearly complete (manually-proofread) single neuron, and **c** a mesh containing two merged neurons from the MICrONS volume. **d** Example of an orphan merge error with a piece of dendrite incorrectly merged onto a neuron mesh, and **e** an incompletely-reconstructed neuron from the H01 volume. The reconstructed meshes are pre-processed through a number of steps including decimation, glia (**f**) and nuclei removal, soma detection (**g**), and skeletonization (**h**). Finally, mesh features are projected onto branches from the skeleton (**i**) and spines are detected (**j**).

113 include glia or pieces of blood vessels merged onto neurons (Fig. 1f). We take advantage of existing tools for mesh processing
 114 (Fabri et al., 2023; Cacciola et al., 2023; Gao et al., 2023; Dorkenwald, 2022) and apply them in an initial workflow that is
 115 agnostic to the identity of the mesh object. The downsampled meshes enter a neuron-specific mesh preprocessing pipeline
 116 (Fig. 1f-g, see Methods): In this pipeline, merged glia and nuclei submeshes are identified and filtered away (Fig. 1f),
 117 the locations of all somas are detected (Fig. 1g), and the mesh representation of axonal and dendritic processes are skeletonized
 118 into a series of connected line segments (Fig. 1h). In a final series of steps, the correspondence between each skeletal segment
 119 and its overlying mesh is reestablished (Fig. 1i), and the mesh segments composing neuronal spines are automatically detected,
 120 taking advantage of the more trivial problem of identifying protrusions on individual non-branching segments (rather than the
 121 entire dendritic arbor; Fig. 1j). Systematic inspection by manual proofreaders confirmed the high accuracy of the soma, axon,
 122 dendrite, glia, and spine annotations generated during the mesh processing workflow (Fig. 8), as well as the labeling of basal
 123 and apical dendritic compartments (Fig. 9).

124 **Graph Decomposition.** We decompose skeletons of axonal and dendritic processes into a directed tree graph (NetworkX
 125 object in Python Hagberg et al. 2008) where the root node is the soma, the other nodes are individual non-branching segments
 126 and edges project downstream in the direction away from the soma; therefore, subgraphs downstream of the soma are a stem.
 127 There may initially be multiple soma nodes in the graph, but these are split apart if more than one soma is detected, and
 128 any cycles in the graph are broken during the decomposition process (see Methods, Fig.2b). Previous work has emphasized
 129 the utility of this kind of graph representation of each neuron, which facilitates flexible queries and analyses of features and
 130 annotations at different scales: individual branches, subgraphs of stems, or all of the stems graphs at once (entire neuron)
 131 (Pastor et al., 2021; Schneider-Mizell et al., 2016).

132 NEURD computes a variety of features of decomposed neurons at the branch (node), stem (subgraph), or whole neuron
 133 (graph) level (Fig. 2a,b). Node features include mesh features of each branch (number of spines, average width, etc.), skeletal
 134 features (length, directed angle of projection, etc.) and synapse features (number of postsyns/presyns, number of spine head
 135 synapses, etc.). Subgraph features are computed from the interactions of node features in a local neighborhood of the decom-
 136 position graph, such as angle between skeleton projections and difference in average width. Graph features can be computed
 137 by averaging node features in the graph, such as the average number of spines per μm of skeletal length or the average number
 138 of postsyns per μm of skeletal length. Alternatively, graph features can be calculated by counting subgraph motifs, such as the
 139 number of subgraphs bordering at least one soma node or the number of subgraphs connected to multiple soma nodes.

140 These multi-scale features make it straightforward to translate neuroscience domain knowledge into neuron or compartment-
 141 level operations and queries. The most important context for this translation is automatic proofreading, where we demonstrate
 142 that these features can be used for accurate cleaning of merge errors including multi-soma merges. NEURD also provides
 143 workflows for common tasks such as annotating cellular compartments and classifying cells based on morphology. In addition
 144 to heuristic approaches for these tasks, the suite of graph features computed during decomposition can be used as inputs to
 145 statistical analyses or as training sets for machine learning. In what follows, we demonstrate the use of these annotated graphs

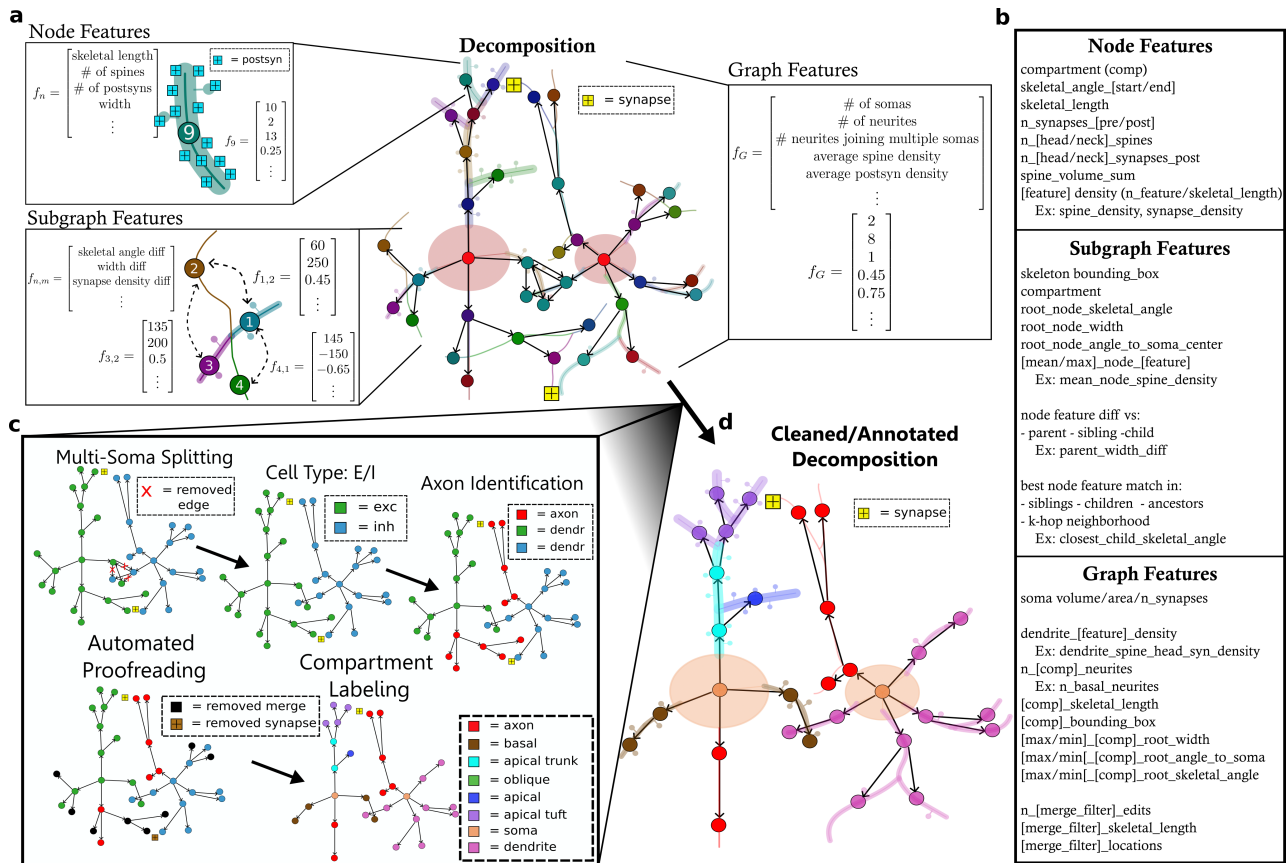


Fig. 2. Decomposition, feature extraction, and graph annotation a Decomposition graph object composed of two neurons merged together. The decomposition compresses the skeleton, mesh and synapses (and any derived features) of a non branching segment into a node in a graph, with directed edges to the downstream segments after a branch point. The soma is a singular node connected to the downstream nodes bordering the soma. **b** NEURD automates computation of features at multiple levels. Node (non-branching segment)-level features include basic mesh characteristics such as the diameter of the neural process or the number of synapses per skeletal length. Subgraph features capture relationships between adjacent nodes like branching angle or width differences. Graph features capture characteristics of the entire neuron and are computed by weighted average or sum of node features, or by counting subgraph motifs. **c** NEURD supports a variety of operations and manipulations on the decomposition objects. Multi-soma splitting is performed if a decomposition graph contains multiple soma nodes, splitting the graph into separate single soma subgraphs based on heuristic rules. The rest of the processing steps are performed on each single soma subgraph separately. The entire decomposition graph is classified as excitatory or inhibitory using a logistic regression model. A heuristic algorithm identifies one subgraph as the axon while labeling all other stems as dendrite. Automated proofreading is performed by a series of heuristic graph rules identifying and removing probable merge errors (see Fig. 3). Finally, a set of heuristic rules are implemented to find subgraphs representing specific neural compartments (apical trunk, apical tuft, etc.). **d** The final product of this workflow is a cleaned and annotated decomposition object with a single soma that can be fed into a variety of downstream analysis pipelines.

146 in the context of automated proofreading (Fig. 3), cell-type classification (Fig. 4), morphological analysis (Fig. 5), and
 147 connectivity analysis (Fig. 6, Fig. 7). Note, these decomposition graphs can easily be converted to SWC files (and NEURD
 148 provides this functionality) in case the user desires to utilize any other neuron processing pipelines requiring that data input
 149 format.

150 **Automated Proofreading.** Node, graph, and subgraph features computed during the decomposition step can be queried to
 151 identify potential sites of merge errors in the reconstruction, enabling rule-based automatic proofreading where all nodes down-
 152 stream of the error are stripped from the mesh. While this approach does not enable automatic extension (i.e. fixing split errors),
 153 it can facilitate many subsequent analyses of individual cell morphology and circuit connectivity.

154 Automated proofreading involves defining graph configurations of nodes and attributes (called graph filters) that typically
 155 indicate merge errors, finding matching subgraphs in each neuron graph, and then labeling either all or some of the nodes in that
 156 match as errors. For example, a graph filter to locate dendritic branches with three or more downstream nodes and with certain
 157 width jumps or skeletal angle differences is useful for identifying errors in the H01 dataset because there are more dendritic
 158 merge errors compared to MICrONS. The graph filters are mostly directed one-hop or zero-hop configurations where one-hop
 159 configurations only consider adjacent downstream nodes, and zero-hop configurations consider only individual node features.
 160 However, proofreading rules can also consider larger-scale features of the graph (such as enforcing the existence of only one
 161 axonal arbor for each cell). Graph filter heuristics and parameters can be tuned for axon and dendrite subgraphs, excitatory
 162 or inhibitory cell types, and specific data sources (H01 vs MICrONS). In a workflow based on this approach, the user defines
 163 which graph filters to apply in which order, and as the errors are identified, the meta-data for each correction is stored (to

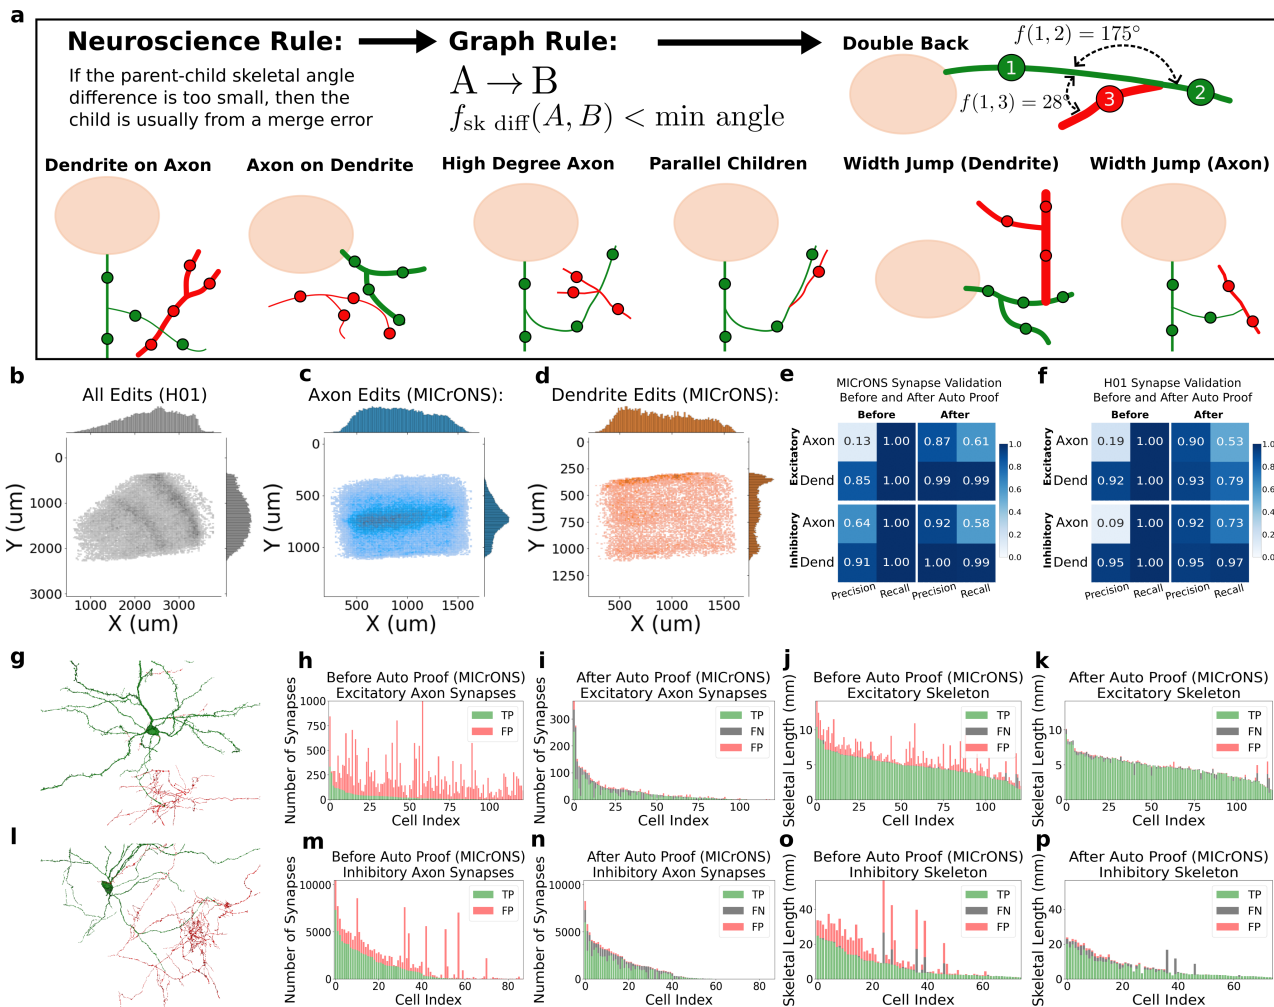


Fig. 3. NEURD graph decomposition enables automated proofreading a Implementing domain knowledge as subgraph rules to automatically identify and remove merge errors. Most of the same rules can be applied across excitatory and inhibitory cells in the MICrONS and H01 volume as-is, or with small changes in parameters. **b** Laminar distribution of all graph rule edits in the H01 dataset illustrating the inhomogeneity of errors across different layers, likely due to differences in neuropil density. Pial surface is to the right and slightly up. **c** In the MICrONS dataset, an increased frequency of axon edits is observed in layer 5 of cortex. Pial surface is up. **d** Dendritic errors in the MICrONS dataset are increased near the top layers of the volume, where fine excitatory apical tufts lead to more frequent merges. **e,f** Validation of MICrONS and H01 neurons quantified by synapse precision and recall compared to manual proofreading. "Before" indicates the state of the raw segmentation prior to any proofreading. Note the high precision of dendrites in both volumes, even prior to automated proofreading. The substantial increases in precision "After" automated proofreading indicate that the cleaned neurons have good fidelity nearly on par with manual cleaning. The reduction in "After" recall indicates that we are losing some valid synapses in the automatic proofreading process (mostly on the axons) but still retaining the majority of correct synapses. For the MICrONS dataset the validation is further separated into neurons with a multi-soma merge and those without in Fig. 15, showing the very high (99%) dendritic recall for single soma segments due to the high quality of initial dendritic reconstruction. **g** Example excitatory neuron from the MICrONS dataset in the 75th percentile of merge error skeletal length; merge errors identified by automatic proofreading are shown in red. **h,i** Number of true positive (TP; green) and false positive (FP; red) axonal synapses from individual excitatory neurons in the validation set before and after automated proofreading, illustrating the large number of false positive synapses in the raw segmentation that are removed using the automated proofreading. **j,k** Skeletal length of true and false excitatory axon skeletons before and after automated proofreading. **l-p** Same as **g-k** but for inhibitory neurons. Note that false-negative axonal skeleton in **j,o** indicates that some axon segments are lost in the mesh or decomposition preprocessing pipeline prior to automatic proofreading.

164 be used for subsequent semi-supervised proofreading or training data for machine learning approaches). Once the algorithm
 165 processes all filters, the mesh, graph nodes, and synapses associated with all error branches are removed.

166 To illustrate this approach, we provide a small set of heuristic proofreading rules implemented as graph filters (Fig. 3a; see
 167 Methods) that yielded good performance on merge error correction in both volumes, but especially in the MICrONS dataset.
 168 An example rule indicating a merge error is "Low Degree Axon" (the parent class of the "Parallel Children" rule), where
 169 an upstream branch node with exactly two downstream nodes has its local graph configuration checked for properties more
 170 typically associated with merge errors (0 degree or 90 degree angle between skeleton projection of downstream nodes, a
 171 synapse at the branching point, very thin mesh segmentation at branching point, etc.).

172 Manual validation of these rules was performed in the context of standard proofreading and multi-soma splitting using the
 173 NeuVue Proofreading Platform (Xenes et al., 2022). We provided the proofreading team at John Hopkins University Applied
 174 Physics Laboratory (APL) with suggested error locations in the MICrONS volume, and experienced proofreaders evaluated

175 each proposed split for accuracy. This validation included identifying split points for multi-soma splitting, axon-on-axon,
176 and axon-on-dendrite merge errors, and enabled us to measure both the accuracy of these proofreading rules and the speed
177 benefits of a semi-supervised approach compared to a fully-manual effort. We were also able to optimize these rules based
178 on proofreader feedback, and this process enabled us to identify specific rules and parameter thresholds that we could apply
179 with high confidence to correct merge errors without human intervention (Fig. 10a-c,e). Specifically, we used feedback from
180 proofreaders to pick parameters that selected a high confidence subset of axon-on-axon merges and axon-on-dendrite merges.
181 Validation of this subset then revealed a 99% and 95% accuracy respectively (Fig. 10e). We applied nearly 150,000 of these
182 high-confidence automatic edits back into the current MICrONS segmentation. Furthermore, using NEURD suggestions in a
183 semi-supervised manner to guide the challenging process of splitting multi-soma segments increased the speed of this process
184 more than three-fold compared to other augmented methods (Fig. 10d, see Methods).

185 It is worth noting the scale of the merge errors removed by automated proofreading. In the MICrONS dataset, we identified
186 hundreds of thousands of merge errors corresponding to dozens of meters of incorrectly-merged axons and dendrites onto the
187 neurons within the 1mm³ volume (Fig. 11). Corrections in the H01 dataset were an order of magnitude smaller due to fewer
188 cells and the less complete initial reconstructions in that volume, but were still substantial. These errors reflect the scale of
189 neural wiring that exists in a tiny volume of cortex, despite the high performance of both automated segmentations. Merge
190 errors were identified throughout both volumes, and were more prevalent for some regions due to sectioning artifacts (Fig.
191 12g-l), or due to intrinsic differences in the morphology of neurons across layers (Fig. 3b-d, Fig. 12a-f, Fig. 13a-g.) For
192 example high- and low-degree axon edits in the MICrONS dataset were frequently made in upper layer 5, potentially due to
193 higher quantities of inhibitory neuropil, while dendritic double-back and width jump errors were more frequently located near
194 the top layers of the volume, due to merges between fine distal apical tufts of excitatory cells (Fig. 12, Fig. 13).

195 We compared the outcome of automatic proofreading (all edits, not just the high confidence subset) to manual proofreading
196 on a test set of cells in the MICrONS (n=122 excitatory and n=88 inhibitory) and H01 (n=49 excitatory and n=18 inhibitory)
197 volumes. The precision of synaptic data was substantially higher after proofreading (for example 0.87 after compared to 0.13
198 before for MICrONS excitatory axons), without a commensurate reduction in recall (summarized in Fig. 3e,f, Fig. 14, Fig.
199 16, and Fig. 15). Note, because our automated proofreading procedure only removes data, recall is measured based on the
200 synapses which were part of the automatic segmentation and which remained after manual proofreading was performed, but
201 does not include synapses which were added in manual proofreading. Note also the high recall for dendrites, which was even
202 higher (99% for MICrONS) when considering only single-soma neurons 15, reflecting the high performance of the initial seg-
203 mentation. Overall recall was lower for axons, indicating a larger number of incorrectly-removed axonal segments compared to
204 dendrites. Performance on the H01 dataset was also reduced compared to MICrONS because the less-extensive reconstruction
205 was associated with fewer merge errors overall. Initial reconstruction quality of dendrites in both volumes was high, and espe-
206 cially in the MICrONS volume, extensive axonal arbors remained after removing merge errors (Fig. 3g,l, Fig. 17). Note that
207 our proofreading heuristics and parameters were tuned to remove incorrect connections at the expense of sometimes removing
208 correct connections, based on the rationale that incomplete, but largely-correct morphology and connectivity is more useful for
209 downstream analyses than more complete neurons with falsely merged segments. In summary, both from the perspective of
210 synapses (Fig. 3h,i,m,n) and skeletons (Fig. 3j,k,o,p), our automated proofreading approach can be applied at scale to yield
211 neural reconstructions (Fig. 17) of similar quality to manually-cleaned cells without any extension.

212 **Cell-type classification.** Densely-reconstructed EM volumes hold great promise for understanding the connectivity between
213 different neural subtypes (Schneider-Mizell et al., 2023, 2021; Dorkenwald et al., 2022a; Weis et al., 2022; Dorkenwald et al.,
214 2022c; Peters and Feldman, 1976; Martin and Whitteridge, 1984). Because EM provides limited access to genetic markers,
215 cell types must be identified by morphological features, although the relationship with molecularly-defined cell classes can be
216 inferred from extensive previous work relating morphological features to transcriptomic classes (Scala et al., 2021; Peng et al.,
217 2021; Gouwens et al., 2020). Compared to light microscopy, the ultra-high resolution of EM reconstructions provides many
218 additional features that can be used to classify cells by morphology. Previous studies have demonstrated that rich information
219 enabling cell-type classification is available even in local nuclear and peri-somatic features (Elabbady et al., 2022; Al-Thelaya
220 et al., 2021), small segments of neural processes (Dorkenwald et al., 2022c), and the shape of postsynaptic regions (Seshamani
221 et al., 2020). NEURD provides an additional rich and interpretable feature set that can be used for cell-type classification via a
222 number of different approaches.

223 As an initial demonstration, we performed Principle Components Analysis (PCA) on the averages of NEURD node features
224 for neurons hand-labeled as inhibitory or excitatory (MICrONS n=3,985 excitatory and n=897 inhibitory; H01 n=5,800 exci-
225 tatory and n=1,755 inhibitory). As expected, spine density (number of spines per micron of skeletal length) and shaft synapse
226 density (number of synapses not on a spine per micron of skeletal length) were very informative features for distinguishing
227 excitatory from inhibitory cells (Azouz et al., 1997). In fact, we found that a logistic regression model trained on just these
228 two spine and synapse features separates excitatory and inhibitory cells with high accuracy, using the same parameters for
229 classification across both the MICrONS and H01 dataset (Fig. 4a,b). We were able to perform this classification step prior to
230 proofreading and axon/dendrite annotation, which enabled these rules to be specific to excitatory or inhibitory neurons.

231 To test whether NEURD graph objects could be used to distinguish even finer cell types, we turned to Graph Convolutional

232 Networks (GCN) (Fig. 4c-e). We trained a simple GCN on the dendritic subgraph of a variety of hand-labeled cell types in the
233 MICrONS volume (n=873 total cells). We focused on the dendrites in this volume because of their high recall from the initial
234 segmentation and the high precision after automated proofreading (Fig. 3f). This classifier produced an embedding space with
235 a continuum of excitatory neurons progressing from the top layers down to the bottom layers, while keeping inhibitory neurons
236 and some morphologically-distinct excitatory neurons (5P-NP and 5P-PT) clearly separated (Fig. 4c). Most of the embedding
237 space was covered by the labeled dataset (Fig. 4c), and cells outside the labeled dataset had soma centroids at expected laminar
238 depths (Fig. 4d); even though no coordinate features are used in the GCN classifier. Thus, the classifier was able to identify
239 these layer-specific classes from distinct morphological features, not simply based on the location of cells in the training set.
240 Within the smaller volume with manually-labeled cells, cell type classification performance on a held-out test set was similarly
241 high for both entire dendritic graphs (n=178 test neurons; Fig. 4e, Fig. 19a-c), and their disconnected stems (n=1023 test stems;
242 Fig. 4f, Fig. 19d-f). Consistent with previous literature classifying cells based on more local features, this demonstrates that
243 the information present in disconnected individual dendritic stems (branching segments connected to the soma) is sufficient to
244 perform fine cell-type classification nearly as well as graphs representing entire neurons.

245 **Morphological Analysis.** The features extracted by NEURD - including features of different compartments (Fig. 5a), the ge-
246 ometry of axonal and dendritic compartments (Fig. 5b), and spine features (Fig. 5c) - provide a rich substrate for morphological
247 analysis (Fig. 20).

248 Extensive work has been done on automatic spine detection in 2-D or 3-D image data using fully-automatic (Xiao et al.,
249 2018; Driscoll et al., 2019; Janoos et al., 2009; Shi et al., 2014; Basu et al., 2018) or semi-automatic (Benavides-Piccione et al.,
250 2013) approaches. NEURD offers an accurate spine detection workflow that achieves high accuracy with a mesh segmentation
251 approach. Importantly, these spine features are then available for queries in combination with all the other morphology and
252 connectivity information available in NEURD graphs. Precision and recall for spines with a skeletal length larger than 700
253 nm was 90% or higher in extensive validation by manual annotators on both volumes (Fig. 8). In addition, NEURD isolates
254 individual spine meshes, then segments the spine head from the neck (when possible), and finally computes statistics about
255 the overall spine mesh and also the individual head and neck submeshes. As expected, the spine head volume and synaptic
256 density volume were the only strongly correlated spine features (Fig. 23; Harris and Stevens 1989; Arellano et al. 2007). The
257 kernel density estimation of UMAP embeddings for feature vectors of spines sampled from the MICrONS and H01 dataset
258 showed a similar embedding structure, with spines that share similar features embedded in similar locations and a somewhat
259 consistent embedding pattern for inhibitory and excitatory spines in the two volumes (Fig. 5d, Fig. 5e). These results build on,
260 and are consistent with, previous work describing the distribution of non-parametric representations of post-synaptic shapes
261 across diverse neural subtypes (Seshamani et al., 2020).

262 We attempted to replicate and extend several other findings observed in previous studies of the MICrONS and H01 datasets
263 regarding the sub-cellular targeting of synaptic inputs. First, we counted synapses onto the axon initial segment (AIS) of
264 neurons at different depths. Replicating a previous report, in the MICrONS volume, superficial L2/3 pyramidal cells received
265 the largest number of AIS synapses, with up to 2-3 times the innervation of the lower cortical layers (Fig. 5f; Schneider-Mizell
266 et al. 2021; Wang et al. 2019; Inan et al. 2013). However, in the H01 dataset, this laminar inhomogeneity in AIS synapses was
267 much less prominent, with more similar numbers of AIS inputs observed across all depths (Fig. 5g). Additionally, like AIS
268 synapses, we found a striking difference in the distribution of somatic synapses across depth between the MICrONS and H01
269 dataset. (Fig. 21g,h). Lastly, the overall frequency of somatic synapses were also distinct across the two volumes, consistent
270 with previous literature describing fewer somatic synapses in the human compared to mouse (Fig. 21d; Wildenberg et al. 2021);
271 however, we found the opposite trend for the AIS, with fewer AIS synapses in the mouse compared to the human (Fig. 21c).

272 In H01, deep layer pyramidal cells were previously observed to have a strong bias in the radial angle of their thickest basal
273 dendrite (Shapson-Coe et al., 2021). The geometry of neural branching is a subgraph-level feature computed by NEURD, and
274 process diameter is a node feature, so we examined the combination of these features to replicate this finding. First we examined
275 the MICrONS volume and did not observe a strong bias in thickest basal, even in deep layers (Fig. 5h). Then, looking at H01,
276 we were able to replicate the pattern of thickest basal dendrite direction preferences in deeper layers (Fig. 5i). However, we
277 also found that this pattern appeared to continue into more superficial layers. We wondered whether the previous emphasis on
278 deep layers might be because the effect was more salient there. Indeed, making use of the precomputed widths for all branches,
279 we found that in deep layers the difference between the thickest and second-thickest dendrite was nearly doubled compared to
280 more superficial layers. Thus, this dendritic orientation bias appears to be a general property of the H01 volume, albeit less
281 salient in more superficial layers.

282 The diversity of precomputed features offered by NEURD enabled us to identify several interesting morphological features
283 that differ across cell types, including many that have been reported previously in other studies. For example, the spindly, non-
284 branching basal dendrites of NP cells (Schneider-Mizell et al., 2023; Weis et al., 2022) are clearly distinct from extensively-
285 branching basal dendrites of L2/3 pyramidal cells. This feature is clear in a histogram of total skeletal length across individual
286 nodes (Fig. 5j), as is the fact that neurogliaform cells are the most highly-branched neurons with the largest number of leaf
287 nodes (Fig. 22a). As an interesting sidenote, we found that across all neurons, dendritic stems with larger numbers of leaf

288 nodes had a larger initial dendritic diameter at their connection to the soma (Fig. 22b-c), potentially reflecting developmental
289 or metabolic constraints.

290 Comparing synapses onto the dendritic shaft with synapses onto dendritic spines provides a rough proxy of inhibitory versus
291 excitatory inputs (Ribak et al., 1981; DeFelipe and Fariñas, 1992; Kwon et al., 2019). In a histogram of shaft to spine synapses,
292 NP cells were again located at the higher end of the distribution, while L4 and L2/3 pyramidal cells had the lowest shaft-synapse
293 to spine-synapse ratio (Fig. 5k), suggesting they receive a relatively larger fraction of excitatory (compared to inhibitory) input.
294 Because NEURD also automatically segments both soma meshes and spine heads and necks, this enables comparison across
295 cell types of features like soma volume and somatic synapses (Fig. 8b,f), spine neck length (Fig. 5l), spine density (Fig. 8a),
296 and the relationship between spine synapse size and spine head volume as in Fig. 5k and Fig. 23.

297 **Connectivity and Proximities.** Next, we examined the connectivity graph in the MICrONS and H01 datasets after automatic
298 proofreading. Proofreading substantially reduced the mean in- and out-degree across both volumes due to the removal of merge
299 errors, resulting in a sparsely-sampled but high-fidelity graph (Fig. 3e,f). A variety of connectivity statistics including number
300 of nodes and edges, mean in and out degrees, and mean shortest path between pairs of neurons along excitatory and inhibitory
301 nodes is provided in Fig. 25. Note that "edges" here are synapses between neurons, and a subset of neurons that were not
302 fully processed by the decomposition pipeline are excluded from all statistics comparing the "raw" graph to the graph after
303 automated proofreading (neurons with manual proofreading are excluded from this comparison). The shortest path statistics
304 are generated by randomly sampling source and target neurons; sample pairs without a path between them are excluded from
305 the calculation.

306 To facilitate the analysis of synapse specificity in sparse connectomes, we implemented a fast workflow for identifying axo-
307 dendrite proximities, regions where the axon of one neuron passes within at least within 5 μm of the dendrite of another neuron
308 (See Methods, Fig. 6a). Note that these proximities would not be meaningful without first cleaning all skeletons of merge errors.
309 The width measurement at every skeletal point can also be used to exclude myelinated axons from this proximity computation
310 using a maximum width threshold. Previous studies have computed proximities from skeletons of simulated models (Udvary
311 et al., 2022), or manually traced data (Mishchenko et al., 2010; Kasthuri et al., 2015), with a similar logic. Proximities are
312 necessary but not sufficient for the formation of a synapse (Peters and Feldman, 1976; Brown and Hestrin, 2009; Mishchenko
313 et al., 2010; Costa and Martin, 2011), and so the "proximity graph" can serve as a valuable null distribution for comparing
314 potential connectivity with synaptic connectivity between neurons: Instead of looking at synapse counts between cells which
315 are dependent on the geometry and completeness of the neuropil, proximities make it possible to calculate "conversion rates" -
316 the fraction of proximities which resulted in actual synaptic connections. NEURD also provides functions to compute presyn
317 skeletal walk - the distance from a synapse to the soma of the presynaptic neuron along the axon, and postsyn skeletal
318 walk - the distance from synapse to soma along the postsynaptic dendrite. Combined with cell typing, compartment labeling
319 and spine annotation, these features enable powerful analyses of neural connectivity conditioned on the cellular identity and
320 subcellular location of synapses on both pre- and post-synaptic partners (Fig. 6b).

321 Conversion rates between neural subtypes in the MICrONS dataset replicated previous results from connectivity measured
322 via slice multi-patching and EM reconstructions, especially the prolific connectivity of basket cells onto both excitatory and
323 inhibitory somas (Fig. 6c; Jiang et al. 2015; Schneider-Mizell et al. 2023; Lee et al. 2013; Freund and Katona 2007) and
324 inhibitory-inhibitory relationships including BC inhibiting other BC, MC avoiding inhibiting other MC, and BPC preferentially
325 inhibiting MC (Fig.27; Pfeffer et al. 2013; Jiang et al. 2015; Lee et al. 2013; Schneider-Mizell et al. 2023).

326 The subcellular targeting of different inputs is apparent in plots of the postsynaptic skeletal walk distance to the soma for
327 synapses arriving at the basal dendrite. As has been previously described (Ribak et al., 1981; Hwang et al., 2021; Megias et al.,
328 2001), inhibitory-onto-excitatory synapses tend to be found closer to the somatic compartment than excitatory-onto-excitatory
329 synapses (Fig. 6d,e). At an even smaller scale, with the spine head, spine neck, or shaft classification propagated to synapses,
330 we can study how excitatory and inhibitory inputs to spines display different scaling relationships between synapse size and
331 spine head volume (Fig. 24). We also show, as expected, that excitatory and inhibitory cells differ in the number and relative
332 sizes of synapses on their target spine heads (Fig. 24; Parnavelas et al. 1977; Megias et al. 2001).

333 Conversion rates for excitatory-to-excitatory proximities were low in both H01 and MICrONS, consistent with previous
334 findings of sparse pyramidal cell connectivity in the cortex (Fig. 6f,g; Campagnola et al. 2022; Jiang et al. 2015; Kasthuri et al.
335 2015; Mishchenko et al. 2010). However, conversion rates were substantially higher for excitatory-to-inhibitory proximities
336 (Fig. 6f,g), especially in H01, and were substantially higher for proximity distances less than 2 microns (unlike excitatory
337 synapses onto excitatory cells, where spines presumably reduce the dependence on distance). Combining the presyn (axonal)
338 skeletal walk features and proximity analyses revealed an interesting similarity in excitatory-onto-inhibitory connectivity be-
339 tween the MICrONS and H01 datasets, with conversion rates peaking in the more proximal axon a few hundred microns from
340 the soma (Fig. 6h,i; Bock et al. 2011; Bopp et al. 2014; Schmidt et al. 2017; Dorkenwald et al. 2022c). Conversion rates were
341 also higher above (more superficial to) the presynaptic soma than below (deeper than) the presynaptic soma for excitatory-onto-
342 inhibitory connections in both volumes (Fig. 26).

343 Large-volume EM connectomics offers tremendous potential opportunities to examine higher-order motifs on a large scale,
344 beyond pairwise connectivity. By comparing the frequency of synaptic motifs with "proximity motifs" defined by neurons

345 that had a chance to participate in a particular synaptic motif based on their axon/dendrite geometry, it is possible to begin to
 346 ask questions about higher-order connectivity even with incompletely-reconstructed neurons because they have been stripped
 347 of merge errors. With this approach we found that more densely-connected triangle motifs were enriched in the MICrONS
 348 volume compared to several controls, including a null graph with the same number of nodes and edges where synaptic degree
 349 distribution is held the same but edges are shuffled (configuration model), a control where the synaptic edges are shuffled only
 350 between neurons with an existing proximity edge, or one where synapses are random shuffled between neurons regardless
 351 of proximity (Fig. 6j). This is consistent with previous findings at smaller spatial scales suggesting that this higher-order
 352 organization is enriched in the cortex (Song et al., 2005; Perin et al., 2011; Milo et al., 2002; Udvary et al., 2022; Schneider-
 353 Mizell et al., 2021). A similar pattern was observed in the H01 dataset, consistent with previous modeling of connections and
 354 proximities there (Udvary et al., 2022). In the H01 volume several of the three-node motifs with larger numbers of connected
 355 edges were missing due to the less complete reconstruction (Fig. 28).

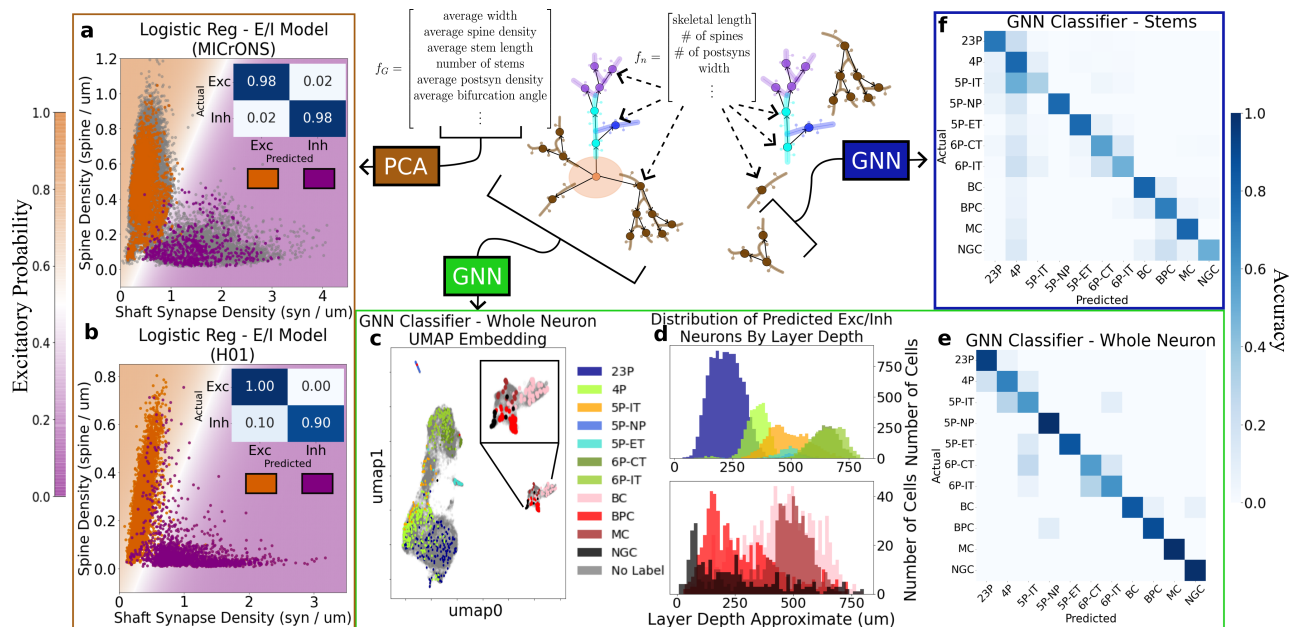


Fig. 4. Graph decomposition enables cell-type classification. **a,b** Two interpretable features identified by PCA were highly informative for excitatory/inhibitory classification: spine density (number of spines per μm of skeletal length) and shaft synapse density (number of synapses not on a spine per μm of skeletal length). Consistent with previous studies (Azouz et al., 1997; Kawaguchi et al., 2006), a logistic regression model trained on just these two features enabled linear discrimination of excitatory and inhibitory cells with high accuracy (same parameters for logistic regression model used for both MICrONS and H01 dataset). **c** A Graph Convolutional Neural Network (GCN) trained on manually-annotated cell types produces an embedding space with a continuum of excitatory neurons progressing from the top layers down to the bottom layers while keeping inhibitory neurons and some excitatory neurons with distinct morphology (5P-NP and 5P-PT) clearly separated. **d** The depth of predicted cell types outside of the training volume are consistent with their expected laminar distribution even though no coordinate features are used in the GCN classifier. **e** Confusion Matrix of the test dataset for the neuron GCN classifier tested on $n=178$ held out neurons: 23P ($n=33$), 4P ($n=51$), 5P-IT ($n=10$), 5P-NP ($n=4$), 5P-PT ($n=7$), 6P-CT ($n=19$), 6P-IT ($n=29$), BC ($n=13$), BPC ($n=9$), MC ($n=1$), NGC ($n=2$). **f** Cell classes could also be determined using a dendritic subgraph of only one stem (branching segment connected to the soma) nearly as well as when using the entire dendritic tree, suggesting that the GCN identifies somewhat local features that enable classification. GCN tested on $n=1023$ test stems: 23P ($n=230$), 4P ($n=301$), 5P-IT ($n=48$), 5P-NP ($n=18$), 5P-PT ($n=45$), 6P-CT ($n=116$), 6P-IT ($n=137$), BC ($n=65$), BPC ($n=30$), MC ($n=19$), NGC ($n=14$).

356 **Functional Connectomics.** A key advantage of the MICrONS dataset is the additional functional characterization of matched
 357 neurons in vivo prior to EM data collection. The relationship between function and synaptic connectivity is covered in detail in
 358 a separate paper (Ding et al., 2023), but here we wanted to provide an illustration of how automated proofreading can enable
 359 functional connectomics analyses that would be impossible otherwise. We identified pairs of excitatory neurons connected by
 360 one, two, three, or four or more synapses. Querying for these rare high-degree connections between pyramidal cells was only
 361 possible after automated proofreading to correct merge errors, especially since inhibitory axon fragments (for example from
 362 basket cells) are much more likely to form multi-synaptic connections and so merge errors onto excitatory cells are a substantial
 363 confound for this analysis. In fact, approximately 97% of the 10,000+ pairs with four or more connections were identified as
 364 merge errors during automatic proofreading. Connections were further restricted to synapses onto postsynaptic spines in order
 365 to guard against possible missed merge errors where an inhibitory axon segment might still be merged to an excitatory neuron.
 366 Examples of these multi-synaptic connections have been highlighted in the H01 data set (Fig. 7a), and rare examples can also
 367 be found in the MICrONS data set (Fig. 7b; see also Chicurel and Harris 1992). We were able to identify $n=11$ of these
 368 pairs in exclusively-automatically proofread neurons (no manual proofreading) where both neurons also had been characterized
 369 functionally (Fig. 7c). The average response correlation was calculated for each group of pairs (see Methods and Ding et al.

2023; Wang and Tolias 2023). We found that neurons with four or more synapses had significantly higher response correlations to visual stimuli than neurons with < 4 synapses connecting them (Fig. 7d), consistent with a Hebbian "fire together/wire together" rule governing high-degree connectivity in the cortex, and this same pattern was also observed for n=12 pairs of manually-proofread neurons. (For autoproofread neurons, two Sample Kolmogorov-Smirnov test and t-test for comparing each multi-synaptic group's null likelihood of being drawn from the same distribution as the one-synapse group: 2 synapses KS test statistic = 0.068 and p = 0.17, t-test p<0.03, 3 synapses KS test statistic = 0.139 and p=0.49, t-test p = 0.60, 4 synapses KS test statistic = 0.508 and p < 10⁻², t-test p<10⁻²).

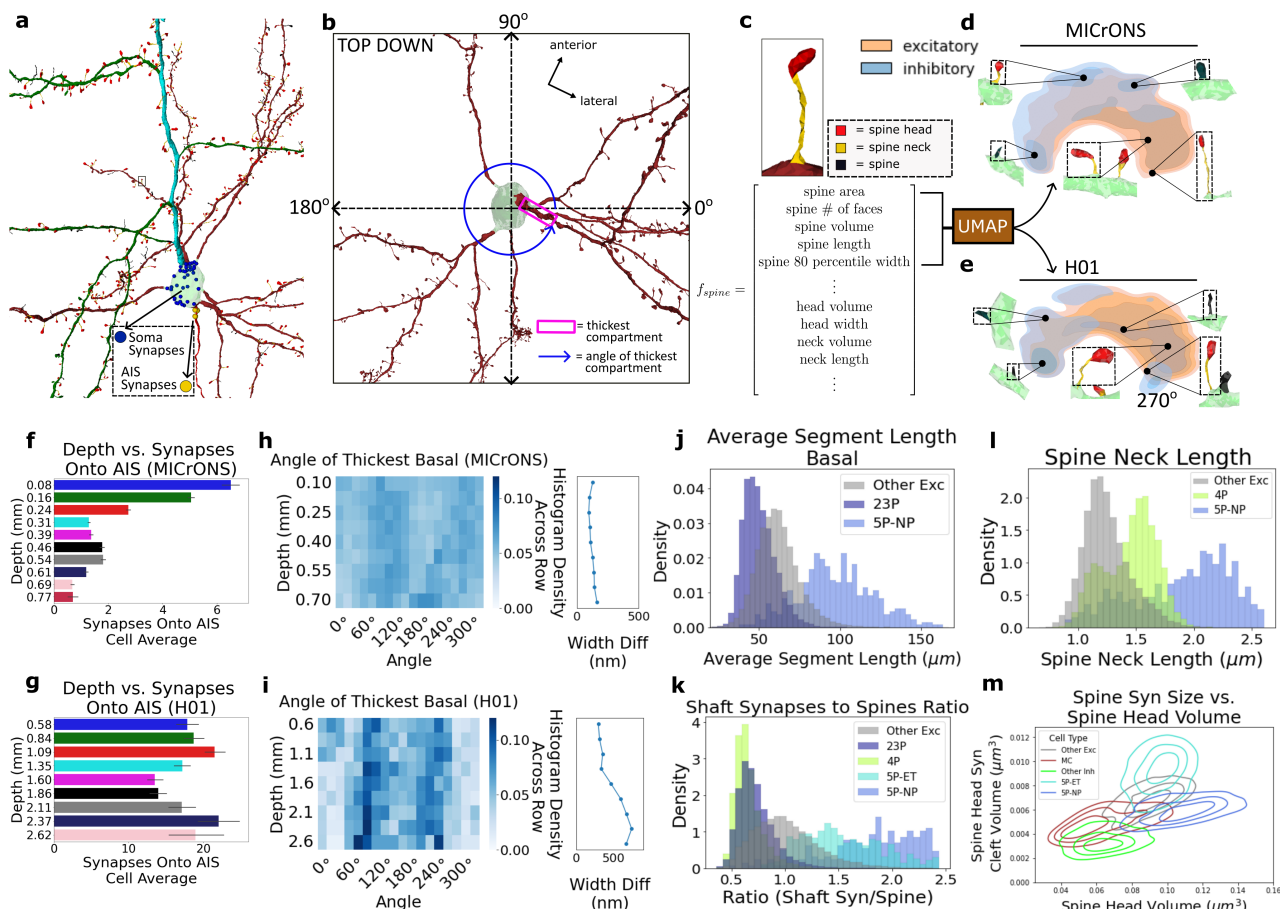


Fig. 5. Morphological analysis enabled by NEURD feature extraction. **a**) Cleaned and annotated neuron mesh; soma synapses in blue, axon initial segment (AIS) synapses in yellow, basal dendrite in brown, apical trunk in aqua, oblique branches in green. Spine heads in red, spine necks in yellow and non-segmented spine in black along dendritic segments. **b**) Top down view of the neuron in **a**). An example of geometric analysis: the thickest basal branch is boxed in pink and the xz angle of these branches are indicated with a blue circular angle marker. **c**) Most spines are annotated with interpretable features such as head volume, spine skeletal length, and spine neck length. Smaller spines (black) are left un-annotated. **d,e**) Kernel density estimation of UMAP embedding of spines sampled from MICrONS and H01 dataset using spine features from panel **c**) (without head or neck features). The embeddings show a similar embedding structure between the two datasets in terms of spine shape and inhibitory/excitatory class, similar to previous work clustering a non-parametric representation of postsynaptic shapes (Seshamani et al., 2020). **f,g**) Average number of synapses onto the axon initial segment (AIS) of cells at different laminar depths (mean +/- std) for the MICrONS and H01 volume. Same depth color scheme used in 21. **h,i**) Histograms showing the distribution of mean skeletal angle of the thickest basal stem. Each row is a normalized histogram for a specific depth bin. The H01 dataset shows more bimodal structure (especially in the deeper layers) than the MICrONS dataset, consistent with a previous report (Shapson-Coe et al., 2021). We find that this pattern is also visible in more superficial layers of H01, but is less obvious because the width difference between the widest and the second-widest basal branch is much larger in deep layers ("Width Diff" plot to right of heat maps). **j**) Distribution of the average skeletal length of non-branching dendritic segments for stems of different cell-type subclasses. **k**) The ratio of non-spine synapses to spine counts varies across cell types. **l**) Distribution of average spine neck length of different excitatory cell type subclasses. As expected, 5P-NP cells have the longest average neck lengths, but 4P cells also display significantly longer necks than other excitatory cell types. **m**) KDE plots of spines on different cell type subclasses, comparing each spine head's volume and the size of the largest synapse on that spine head. These plots reveal differences in both the distribution and scaling of synapse sizes and spine heads across cell types. For example Martinotti cells (red) have larger spine head synapses than other inhibitory spines (light green) and the head volume and synapse size scale at a rate more similar to other excitatory cells (grey).

377 Discussion

378 NEURD is an end-to-end automated pipeline capable of cleaning and annotating 3-D meshes from large electron microscopy
 379 volumes and pre-computing a rich set of morphological and connectomic features that are ready for many kinds of downstream
 380 analyses. Building on existing mesh software packages, NEURD adds a suite of neuroscience-specific mesh functions for soma

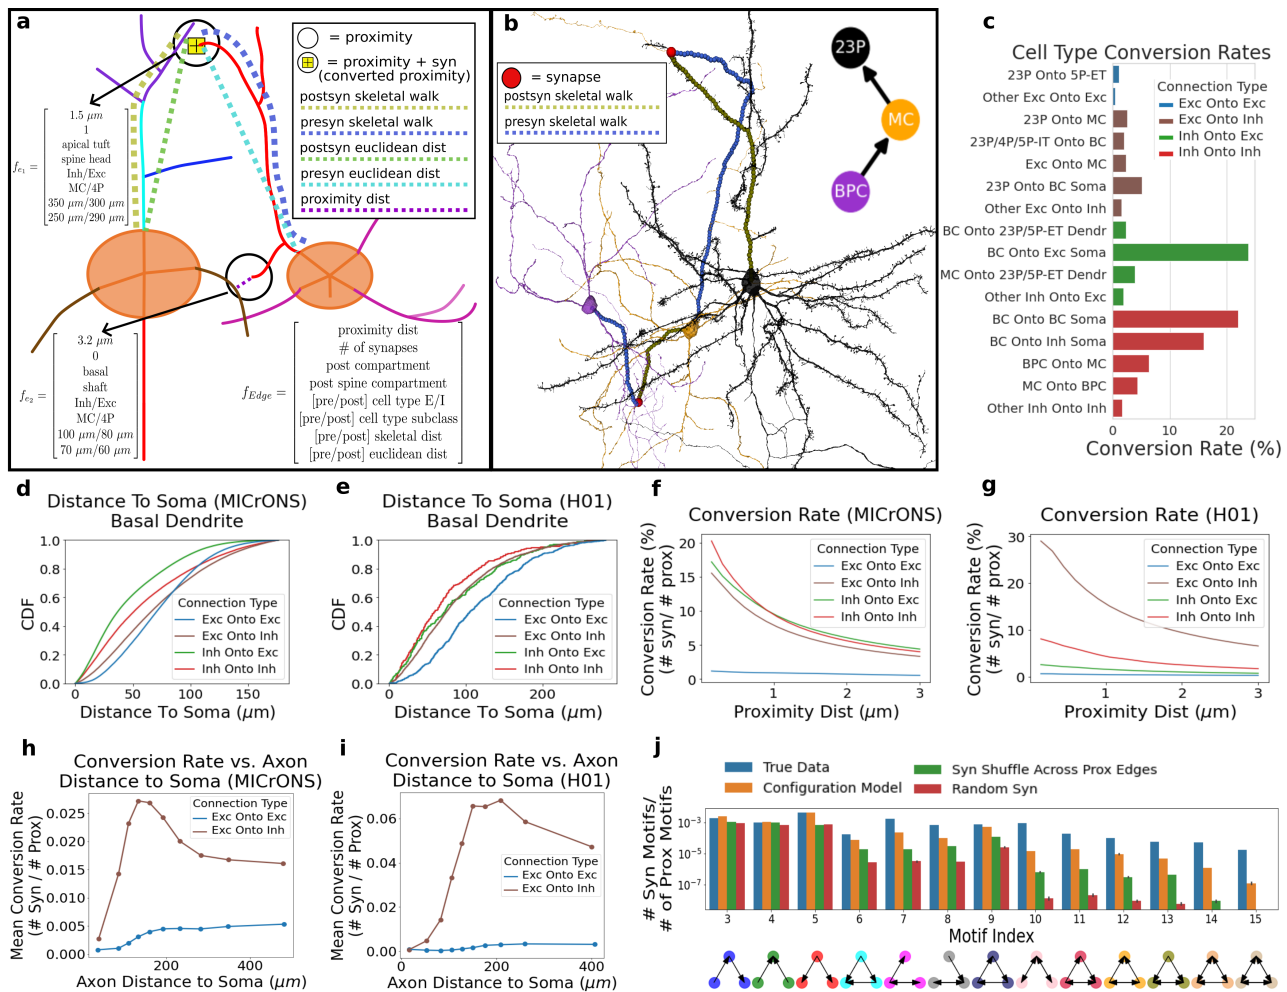


Fig. 6. Connectivity analysis enabled by NEURD **a)** Schematic illustrating two proximities between a pair of neurons where one neuron's axon passes within 5 μm radius of a target neuron's dendrite. Only one proximity actually includes a synapse, thus in this example the "conversion rate" is 50%. Each proximity creates an edge in a proximity graph between the two neurons. Edge features for proximities include the number of synapses, the proximity distance between the closest presyn and postsyn skeletal point, postsyn compartment, pre/postsyn skeletal walk and euclidean distances, etc. **b)** Cell-type specific connections and motifs in the MICrONS dataset can be found by querying the annotated connectivity graph. NEURD allows for visualization of these connection paths and motifs so they can be quickly inspected. **c)** Conversion rates for different cell type subclasses and compartments in the MICrONS dataset are consistent with previous cell type specific connectivity work (cell type labels generated from GNN classifier). **d,e)** Cumulative density function (CDF) of the postsynaptic skeletal walk distance distribution for different exc/inh connection combinations (apical and soma synapses are excluded). In both datasets, excitatory inputs are further along the dendrite from the soma. **f,g)** Conversion Rates (synapses / proximities) for different exc/inh combinations. The x-axis represents the maximum distance that is considered a "proximity". **h,i)** Mean conversion rate as a function of distance from the synapse to the presynaptic cell along the axon. **j)** Proximities enable the computation of null distributions for higher order motifs even with incompletely-reconstructed neurons. The frequency of edge-dense three-node motifs was enriched compared to null graphs (with the same number of nodes and edges) where synaptic degree distribution is held the same but edges are shuffled (configuration model), where the synaptic edges are shuffled only between neurons with an existing proximity edge, or where synapses are randomly shuffled between neurons regardless of proximity (250 random graph samples for each null distribution comparison).

381 identification, spine detection, and spine segmentation that are applicable across multiple data sets, as well as workflows for
 382 skeletonization and mesh correspondence that complement existing tools. The atomic unit (node) in the NEURD decomposition
 383 graph is either a single non-branching segment of axon or dendrite, or a soma - each of these nodes is automatically decorated
 384 with a large number of features. We demonstrate the utility of this highly-annotated coarse-grained graph representation for
 385 many downstream analyses and tasks, including splitting multiple-soma merges, cell typing, automated proofreading of merge
 386 errors, and compartment labeling. We highlight the ability to implement interpretable heuristic rules in these workflows, but
 387 NEURD graphs can also be used as inputs for machine-learning based approaches, as we demonstrate in the context of cell-type
 388 classification. Our hope is that offering many automatically-computed features at the level of coarse-grained NEURD graphs
 389 will make these daunting datasets more accessible for a larger group of researchers. The set of features generated by NEURD
 390 is easily extensible, and the illustrations here have only tapped a tiny bit of their potential for scientific discovery.

391 Several previous studies have proposed post-hoc methods for automated proofreading including merge and split error detec-
 392 tion and correction. Some of these methods make use of Convolutional Neural Networks (CNNs) (either supervised Zung
 393 et al. 2017; Gonda et al. 2021 or unsupervised Rolnick et al. 2017), reinforcement learning methods (Nguyen et al., 2022),

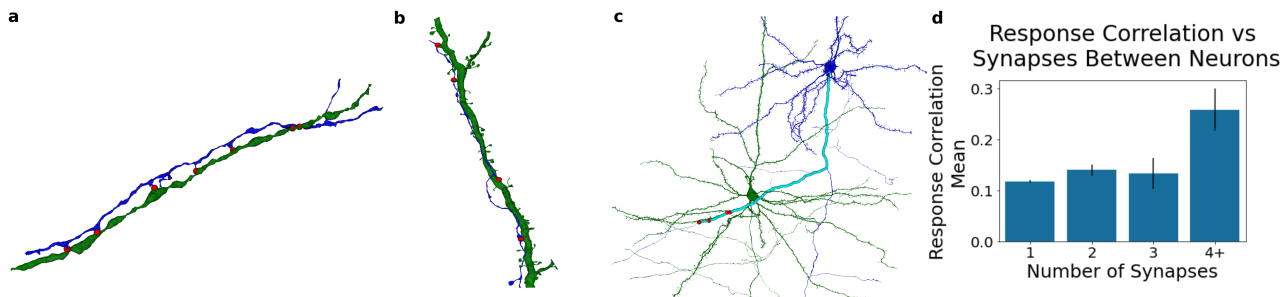


Fig. 7. Functional Connectomics Illustration: High-degree cell pairs **a)** Example multi-synaptic connection ($n=7$ synapses) from an excitatory to inhibitory neuron in the H01 dataset **b)** Example multi-synaptic connection ($n=4$ synapses) from excitatory to excitatory neuron in the MICrONS dataset. **c)** Example of a highly spatially clustered multisynaptic connection ($n=4$ synapses) on a postsynaptic basal dendrite from a neuron cleaned with automated proofreading in the MICrONS dataset (presynaptic skeletal walk shown in aqua, synapses shown in red) **d)** Distribution of response correlation between pairs of functionally matched excitatory neurons in the MICrONS dataset. Response correlation is significantly larger for pairs of neurons with 4 or more synapses connecting them ($n=11$ pairs) compared to those with 1, 2, or 3 synapses ($n=5350, 280, 34$ pairs respectively).

394 or other machine learning approaches (Schubert et al., 2019; Schmidt et al., 2022; Berman et al., 2022). Others make use
395 of heuristic rules applied to neural skeletons (Meirovitch et al., 2016; Sicut et al., 2013), and at least one approach uses both
396 skeleton heuristics and CNNs (Matejek et al., 2019). Although our method is based primarily on heuristic rules, NEURD can
397 support a variety of downstream methods. Multiple proofreading approaches can benefit from NEURD's pre-computed feature
398 set that includes cell types, compartments, spine information, and skeleton geometry. The coarse-scale graph structure facili-
399 tates flexible querying at the level of non-branching segments and the characteristics of the branch points between them. In the
400 context of automated proofreading, having a pre-computed feature set at multiple scales can provide a clearer view of the con-
401 text in which a potential segmentation error occurs. For instance, determining whether a thin, aspiny projection from a dendrite
402 is the true axon or a merge error may be an ambiguous without additional information about the cell type of the neuron, the
403 distance of the candidate axon from the soma, and the spine density of the parent dendritic branch. This information can make
404 the probability of a merge much more clear. We demonstrate how this information can be leveraged together in rules defined
405 as a graph filters with interpretable parameters. We find that many of these graph filters generalize across the two different
406 volumes, either with or without tuning of thresholds. These parameters can also be tuned in interpretable ways for more or less
407 conservative proofreading, or for different downstream applications.

408 Our present implementation does not address some types of errors in automated segmentation. For example, it cannot
409 presently handle merge errors where a single segment needs to be split within its primary trunk as it looks for edits to occur
410 surrounding branch points. Second, because NEURD's present implementation has focused on removing false mergers it
411 is unable to fix incomplete neural processes, though for performing extensions we could potentially leverage its feature set
412 to lower the range of plausible extension candidates. Another limitation is that NEURD currently only takes advantage of
413 the information available in neuronal meshes, which of course only includes features that are visible at the level of the cell
414 membrane. Using the segmentation to loop back to the raw EM data might allow us to capture some of the additional rich
415 ultrastructural features that are available in EM data, for example the density of organelles such as mitochondria, in order to
416 add additional annotations to the NEURD graph.

417 Thanks to the efforts of trained proofreaders at APL, we were able to perform extensive validation of a set of these rules "in
418 situ", which enabled us to speed up semi-supervised splitting of multiple somas, and also enabled us to identify a subset of rules
419 with high accuracy that could be batch-applied automatically to the current segmentation of the MICrONS volume. As far as
420 we know the latter is by far the most extensive set of post-hoc edits generated by an automatic proofreading method that have
421 actually been applied to improve the quality of an EM segmentation. Although EM segmentation methods continue to improve,
422 post-hoc workflows for automated proofreading are likely to continue to be necessary given the amount of neuronal wiring in
423 even a 1mm^3 volume, and having an interpretable and tunable pipeline where we know what errors are easy or difficult for
424 these methods to fix may even inform choices about what to optimize during the segmentation pipeline itself.

425 NEURD graphs are useful for cell typing. Previous work discusses differences in spine and synapse characteristics between
426 cell type classes (Azouz et al., 1997; Kawaguchi et al., 2006). Based on these findings and our exploration with PCA, we find
427 that the precomputed features of spine density and shaft synapse density enable a linear classifier to distinguish these coarse cell
428 classes using the same parameters for both the H01 and MICrONS dataset (Fig.4a,b). More fine-scale subclass classification
429 can be performed with a simple GNN operating directly on the annotated NEURD graph (Fig.4c - f). Even higher performance
430 may be achieved in the future with different classifier architectures. Similar approaches may also enable classification of orphan
431 dendrite or axon segments within the EM volume, greatly expanding the scale of connectivity principles that can be explored
432 in these datasets.

433 Highly-annotated NEURD graphs provide a compact representation of many features that are useful for all kinds of mor-
434 phological analysis, making it possible to extract more of the incredible value that these large datasets offer for looking at

435 morphological characteristics of neurons. For example, a simple query indicates that the percentage of pyramidal cells with
436 axons protruding from dendrites (rather than the soma) is higher in the mouse (17.8%) than the human (8%), which closely
437 replicates the findings of a previous study focused on this question (Wahle et al., 2022). Several of the morphological properties
438 of cell types shown in Fig. 5j-m and Fig. 20 replicate observations from previous studies (Elabbady et al., 2022; Kawaguchi
439 et al., 2006; Villa and Nedivi, 2016). Additionally, using the spine metrics extracted by NEURD, we were able to replicate
440 many of the findings of (Harris and Stevens, 1989; Arellano et al., 2007) concerning synapse size and spine head volume cor-
441 relation, and we also show that these scaling rules and others depend on cell type (Fig.5m, Fig. 23). Looking at how synapses
442 onto the AIS and soma vary across species, we find the expected lower rate of soma synapses on human neurons than mouse
443 (Wildenberg et al., 2021), and replicate the expected distribution of AIS synapses across depth in the mouse (Schneider-Mizell
444 et al., 2021). We also find that the human dataset does not show a similar change over depth in AIS synapses, and that the
445 human AIS is more densely innervated than mouse (Fig.5f-g, Fig. 21c). Finally, we demonstrate the use of a query combining
446 geometric information and branch-level characteristics to replicate the previously-reported bias in the orientation of the thickest
447 basal segment in the H01 dataset (Shapson-Coe et al., 2021). We extend this finding with an observation that this bias is actually
448 consistent across all depths but is just less salient in upper layers because the relative size of the thickest and second-thickest
449 basal dendrites changes smoothly across depth (Fig.5h-i).

450 At the level of the connectome, not only do we have access to a large number of high fidelity synapses and proximities
451 but also precomputed node (neuron) and edge (synapse) features, enabling a variety of connectomics analyses. For instance,
452 using the cell type node labels and the skeletal walk length edge features, we confirmed previous work (Hwang et al., 2021;
453 Megias et al., 2001) describing different distal and proximal preferences for different excitatory and inhibitory connection
454 types (Fig.6d-g). Furthermore, using the proximity controls computed on the cleaned skeletons, we also were able to observe
455 a consistent trend across datasets showing the propensity for forming connections from excitatory to inhibitory neurons peaks
456 around 200 μm away from the soma, potentially consistent with a pattern of surround suppression (Fig.6k,l). Additionally,
457 with the cell type node labels and spine compartment and synapse size edge labels, we confirmed a variety of expected findings
458 about synaptic and spine head size: excitatory to excitatory connections have the largest synapses, synapse size correlates with
459 spine heads for excitatory sources but not inhibitory sources, and spine heads with inhibitory synapses generally are multi-
460 synaptic spines where the inhibitory synapse is typically much smaller than the largest synapse on the spine head (Fig. 24).
461 With spine features (spine length, neck length, etc.) thoroughly documented throughout the volume and immediately available
462 for queries, these properties can be investigated at the level of cell types and compartments to help understand the specificity
463 of these synaptic features. Finally, by using proximities between cleaned skeletons to compute conversion rates as an estimate
464 of connection probability, we replicate a variety of previous findings about the connection probability between both coarse
465 (excitatory/inhibitory) and fine cell types (Jiang et al., 2015; Pfeffer et al., 2013; Schneider-Mizell et al., 2023). A variety of
466 exciting questions about cell type specificity can be addressed with this kind of data.

467 We illustrate the power of combining automated proofreading to generate a clean but incomplete graph, with proximities to
468 serve as a null distribution to account for this incompleteness. Together this powerful approach can begin to reveal principles
469 of pairwise and higher-order connectivity motifs in dense reconstructions. We demonstrate our ability to take advantage of
470 automated proofreading to identify a rare two-node motif (four or more synapses connecting two excitatory neurons) that would
471 be buried in noise in the unproofread volume. This just illustrates one example of the potential for future discovery, especially
472 given the matched functional properties available in the MICrONS dataset. Additionally, with the cleaned connectome, we
473 were able to start counting the number of different triangle motifs (the simplest of higher order motifs) in the datasets, and
474 we observe a general overexpression of these motifs in comparison to proximity controls and some standard null models, as
475 previously reported in (Udvary et al., 2022; Song et al., 2005; Perin et al., 2011; Milo et al., 2002) However, the ability to
476 expand this work to include cell type colorings of these motifs and add proximity based controls will enable investigation of
477 more complicated motif questions, unleashing the power of these extraordinary datasets for a larger community of researchers.

478 Methods

479 **Data Management.** For simplified data management and querying of input neuron reconstruction meshes, NEURD interme-
480 diate decomposition graphs, and all derived statistics and data products, we utilized the DataJoint Python package (Yatsenko
481 et al., 2015, 2018)

482 **Mesh Preprocessing.** NEURD operates on 3-D meshes which are represented in a standard form as lists of vertices and faces
483 in 3-D coordinates. A connected mesh component is a set of faces and vertices in which all faces have at least one adjacent edge
484 to another face. Segmentation algorithms may not output a single connected component as a mesh, but instead may generate
485 several disconnected submeshes, each of which is a subset of faces that is a connected component. NEURD is generally robust to
486 discontinuous meshes, meshes of different resolutions, and several kinds of meshing errors.

487 The resolution of meshes delivered as part of the MICrONS and H01 datasets was sufficiently high that we performed an
488 initial decimation of the mesh (reduced to 25% for MICrONS and 18% for H01) to speed up subsequent computations while
489 retaining all the detail necessary for morphological characterization even of fine axons and spine necks. This decimation was

490 performed using the MeshLab Quadric Edge Collapse Decimation function (Cignoni et al., 2008). Following decimation, we
491 separated this decimated mesh into into connected components. We next applied a Poisson Surface Reconstruction (Cignoni
492 et al., 2008) to each connected component. This can be thought of as "shrink-wrapping" the mesh - it smooths discontinuities
493 on the surface of the mesh and ensures that each connected component is "water-tight" (i.e. no gaps or missing faces). This
494 pre-processing facilitates the subsequent decomposition steps.

495 **Glia, Nuclei Removal.** Glia and nuclei submeshes are identified and filtered away using ambient occlusion functions (Cignoni
496 et al., 2008) to identify regions with a high density of inside faces. Inside faces are mesh faces that are almost fully surrounded
497 by other mesh faces. For example, glia that are merged onto neurons appear as cavities filled with a high density of mesh faces,
498 and are distinct from the hollow reconstructions of most excitatory and inhibitory neurons. Similarly, the mesh representation
499 of the soma surrounds the nucleus mesh and thus nuclei are almost entirely made up of inside faces. Therefore to identify
500 glia and nuclei in the reconstructed meshes, we look for large connected components with high percentages of inside faces as
501 candidates. To determine whether mesh faces are internal or external, we simulate an external "light source" that emits from
502 all angles and we compute the exposure each face receives. This metric is thresholded to classify faces as either inside or
503 outside faces, and submeshes made up entirely of inside faces are candidates for removal. We then apply additional thresholds
504 on the candidate submeshes volume and number of faces to classify them as a glia mesh, nuclei mesh or neither. Finally, for
505 glia we include all floating meshes within the bounding box of the submesh or within a search radius (3000 nm) of any faces
506 of the submesh. This post-hoc mesh agglomeration serves to clean up the areas around glia segmentation, which can be very
507 unconnected and non-standard.

508 **Soma detection.** Soma detection is run on any segment containing at least one detected nucleus (note that nucleus detection
509 was performed previously as part of the segmentation and annotation workflow Consortium et al. 2021; Shapson-Coe et al.
510 2021). To detect the soma, we first perform a temporary heavy decimation of the mesh to remove small features and facilitate
511 detection of the large somatic compartment. We then segment this low-resolution mesh into contiguous submeshes using the
512 CGAL mesh segmentation algorithm (Yaz and Loriot, 2023). This function not only provides the specific faces in each submesh
513 but also an estimate of the width of the submesh as a SDF value (Shape Diameter Function, a measure of diameter at every
514 face of the mesh surface Yaz and Loriot 2023). We then filter all the resulting submeshes for soma candidates by restricting
515 to those within a set size (number of faces), SDF range, bounding box length and volume threshold. We restrict to candidates
516 that are sufficiently spherical to represent the general shape of a soma, but liberal enough to account for somas that are partially
517 reconstructed, for example at the edge of the volume. Once we identify candidate somas in the low-resolution mesh, we return
518 to the original mesh representation (prior to starting soma detection) and apply a final size and width threshold. Given the initial
519 restriction to segments with at least one detected nucleus, if we are not able to detect at least one soma after this process, we
520 relax the thresholds slightly and iterate until a soma is detected or a threshold limit is reached.

521 **Decomposition.** With the glia and nuclei submeshes identified, we filter those away from the original mesh, which may cause
522 splitting into additional connected components. We identify connected components containing at least one soma submesh
523 (note that some segmentations may contain multiple somas prior to soma splitting). Mesh fragments that are not connected to
524 somas may be floating meshes inside the soma (which are filtered away using the same ambient occlusion methods described
525 for glia above), or detached mesh pieces of neural processes that can be stitched to the neuron representation later. For each
526 of the soma-containing meshes we filter away the soma submeshes and identify connected components of these meshes as
527 stems. Any stem submesh must contain at least one set of connected adjacent edges and common vertices shared with a soma
528 submesh ("border vertices"). We construct a connectivity graph where edges only exist between stem and soma nodes if there
529 exists border vertices between the stem and the soma. Through the graph constructed in this manner, limbs that provide paths
530 between multiple somas can be easily identified for subsequent splitting (see below).

531 Each of the stem submeshes is then processed into a skeleton - a 3-D "line-segment" representation that is a set of vertices
532 and edges. We use the Meshparty package for this initial round of skeletonization because of its efficient implementation,
533 and because it provides both a width estimate and a correspondence between the faces of the original mesh and each vertex
534 in the skeleton (Dorkenwald, 2022). This skeleton is then further divided into branches (non-branching subskeletons). The
535 corresponding meshes of branches with an average width greater than a threshold are re-skeletonized with a higher-fidelity
536 method that yields skeletons which pass through the hollow centers of the mesh to provide a better estimation of the location
537 and curvature of the surrounding mesh. This is particularly important for some neural processes, for example wide apical trunks
538 where a skeleton that is not centered within the mesh could be displaced nearly a micron from the actual center of the trunk.
539 This higher-fidelity method is performed using the CGAL Triangulated Surface Mesh Skeletonization algorithm (Gao et al.,
540 2023).

541 For mesh correspondence and width determinations of all skeletons, the NEURD algorithm employs a custom mesh corre-
542 spondence workflow based on the following steps: First, each non-branching segment of the skeleton is divided into smaller
543 pieces, and for each piece a cylindrical search radius is used to identify the mesh correspondence. The closest distance between
544 the skeleton and corresponding mesh faces is computed at multiple points along the skeletal segment, and these are averaged

545 to get a mean radius. Finally, all mesh correspondences of sub-branches are combined into the mesh correspondence of the
546 branch. Concatenating the widths along the sub-branches forms a width array along the branch, with the entire branch width
547 determined from the average of the array. This method results in one face of the original branch possibly corresponding to
548 more than one branch's mesh correspondence, so the algorithm employs a final graph propagation step from unique mesh
549 correspondence faces to allow branches to claim the previously conflicting faces.

550 The procedure described above yields a collection of disconnected non-branching skeletal segments as well as their asso-
551 ciated mesh correspondence. The skeleton of finer-diameter processes is the initial MeshParty [Dorkenwald \(2022\)](#) skeleton
552 which tracks along the mesh surface, while larger-diameter processes have skeletons that track through the center of their vol-
553 ume. These pieces are then all stitched together into a single connectivity graph where each non-branching segment is a node,
554 and the edges between them represent their connectivity. Any conflicts in the mesh correspondence of adjacent nodes at stitch-
555 ing points is again resolved, yielding a smooth and connected mesh representation of the entire stem where each mesh face is
556 associated with a single node (non-branching segment). This entire process is repeated for every stem submesh connected to
557 the soma. Finally, all floating meshes outside the soma are decomposed in the same manner as the stems, and then appended to
558 the existing skeleton if they have any faces within a threshold distance of another node (for example, in the MICrONS dataset
559 the maximum stitch distance was set at 8 μm).

560 The soma(s) and decomposed stems of a neuron are then represented as a NetworkX graph object ([Hagberg et al., 2008](#)).
561 In the ideal scenario there is a single soma root node with multiple stem subgraphs, and each stem subgraph is a directed
562 tree structure representing the connectivity between non-branching segments of the skeleton from the most proximal branch
563 connecting to the soma to the most distal leaves of the axonal or dendritic process. In less ideal cases (which are common),
564 cycles may exist in the skeleton due to self-touches of the axonal or dendritic process (close proximities of neurites that are
565 incorrectly meshed together), and multiple somas may be included in a single segmented object. Handling of these cases is
566 described in "Multi-Soma and Multi-Touch Splitting" below. The soma node contains the soma submesh and SDF values
567 generated during the soma extraction, and each branch node in each stem stores the mesh correspondence, skeleton and width
568 array for that node. Using these raw features, many more features of these branches can be extracted (for example spines), and
569 additional annotations can be added (such as synapses).

570 **Spine Detection.** The non-branching segments produced by the mesh decomposition of each node present an ideal scenario
571 for spine detection. Briefly, we started by using an existing mesh segmentation algorithms ([Fabri et al., 2023](#)) which calculates
572 a local estimate of the volume for each face of the mesh (SDF), applies a Gaussian Mixture Model (GMM) to the distribution
573 of SDF values to enable a soft clustering of faces to k clusters, and finally minimizes an energy function defined by the
574 alpha-expansion graph-cut algorithm to finish with a hard cluster assignment over the mesh. This last step takes a smoothness
575 parameter controlling the likelihood that adjacent faces with concave edges will be more or less likely to be clustered together.
576 We found that setting the number of clusters to 3 and a smoothness of 0.2 was optimal to produce an over-segmentation of
577 the branch mesh which serves as input to the next spine detection step. Then, we convert the branch segmentation into a
578 graph representation (branch segmentation graph) where the nodes are submeshes of the segmentation and edges exist between
579 submeshes with adjacent faces. The dendritic shaft subgraph is determined by establishing the longest contiguous shaft line
580 path in the graph (from most likely node candidates defined by size, width and diameter thresholds), and then spine candidates
581 are identified as subgraphs (not in the shaft path) based on size, volume, and distance from the mesh surface. The final product
582 of this stage is the individual spine submeshes and then subsequently calculated spine statistics (volume, length, area). Based
583 on these statistics, in some cases we perform an additional processing step that divides the submesh of larger spines into a spine
584 head and spine neck. At the completion of this workflow, each mesh face in the node receives a spine label of head, neck, shaft
585 or just "spine" (if no head and neck segmentation could be performed). Finally, the width of each branch is recomputed after
586 removing spines that may have previously confounded that measurement.

587 **Synapse Addition.** Synapses from the reconstruction pipeline are mapped to the closest mesh face of the closest branch. Any
588 annotations of the associated face (for example) spine head, spine neck or shaft can then be propagated to the synapse. In
589 addition, the closest skeletal point on the associated branch object is computed to define an anchor point for the synapse on the
590 neuron's skeleton. This anchor point enables computation of metrics such as skeletal walk distance to the closest spine, closest
591 neurite branch point (upstream or downstream) or skeletal walk distance to the soma.

592 **Multi-Soma and Multi-Touch Splitting.** After the initial decomposition, cycles may exist in the stem graphs due to self-
593 touches in the decomposed mesh (regions where neurites pass very close to each other, resulting in inappropriate connectivity
594 of faces in the mesh representation). Furthermore, stems may include edges with multiple somas if two or more somas are
595 merged together in the same mesh object. This is a challenging problem that requires a general solution since stems can be
596 both multi-touch and multi-soma of any degree. For example, some apical stems of neurons in the MICrONS dataset were
597 initially connected to 9 or more somas due to close mesh proximities with apical tufts of other cells. The aim of this stage in
598 the processing pipeline is to split the stem objects optimally, while attributing the correct portion of the stem mesh and skeleton
599 to the correct neurite.

600 The workflow for splitting both multi-touches and multi-somas proceeds as follows: For every stem identified as having a
601 multi-touch or multi-soma connection, the process first starts by identifying cyclic or soma-to-soma paths. The best edge to cut
602 on the path is then determined using a series of heuristic rules that are applied in the order listed below:

- 603 1. Break any edge on the multi-touch or multi-soma path where the angle between the skeleton vectors of two adjacent
604 branches on the path is larger than some threshold (reflecting the fact that neurite processes generally do not typically
605 abruptly double-back).
- 606 2. Break any edge on the multi-touch or multi-soma path where more than two downstream branches exist at a branch point
607 and the best match for skeletal branch angle and width is not on the multi-touch or multi-soma path.
- 608 3. Break any edge on the multi-touch or multi-soma path where there is a difference in width between two nodes along the
609 path greater than a threshold amount.

610 After an edge is removed based on any of these rules, the process restarts and the graph representation is checked again to
611 see if cyclic or multi-soma paths still exist. The process is repeated until no such paths exist. If no candidate edge is identified
612 by using these rules then, depending on the user settings, the stem may be completely discarded from the neuron object or
613 cut at the very first or last branch.

614 This splitting algorithm is not guaranteed to optimally split all multi-soma or multi-touch paths, but residual errors from
615 a sub-optimal split may be cleaned by further proofreading steps. As with any automated proofreading, the rule and relevant
616 parameters that determined the edit are stored for subsequent evaluation and use.

617 **Excitatory/Inhibitory Classification.** Once each neuron object has a single soma, the NEURD workflow moves on to an
618 initial round of coarse cell classification, determining whether each neuron is excitatory or inhibitory. Performing the classi-
619 fication at this point in the workflow enables the use of subsequent proofreading or annotation algorithms that are excitatory-
620 or inhibitory-specific. For example, axon identification (see below) is much easier if the coarse E/I type of the cell is known
621 beforehand. The cell class is determined via logistic regression on two features: postsynaptic shaft density (number of synapses
622 onto dendritic shafts per micron of skeletal distance on the postsynaptic dendrite) and spine density (number of spines per
623 micron of skeletal distance on the postsynaptic neuron). These two features enabled linearly-separable elliptical clusters for
624 excitatory and inhibitory cells. To enable this classification prior to proofreading, we applied two restrictions to the unproofread
625 graph that reduce potential confounds due to merge errors and ambiguity between axon and dendrite. First, we restrict to larger
626 dendrites using a simple width threshold to not include potential orphan axon merges, and second, we restrict to the proximal
627 dendrite within a limited skeletal walk distance from the soma center. The latter reduces confounds due to dendritic merge
628 errors which are more common at the distal branches (Fig. 3d). When compared to human E/I labels, the classifier results are
629 shown in Fig. 4a,b. These results are also robust against an approximate 10:1 and 1.8:1 excitatory to inhibitory class imbalance
630 in the labeled MICrONS and H01 datasets respectively.

631 **Non-Neuronal Filtering.** While all the neurons in the H01 dataset were hand-checked as neurons and manually annotated
632 for cell types, the MICrONS dataset initially was not. Consequently, segments with nuclei in the MICrONS dataset could
633 include blood vessels, glia cells, or agglomerations of orphan axons without a neuron mesh due to an incorrect nucleus merge,
634 and we did observe this frequently in the version of MICrONS processed in this study, version 374 (but this issue is now
635 largely resolved in the most current data release with tables that indicate which segments are predicted as non-neuronal, using
636 a method independent from NEURD, and with more accurate nuclei merging). Therefore, in order to filter away almost all of
637 these segments without sacrificing a significant amount of valid neurons, we found a suitable filter using cell type classification
638 (predicted by our logistic regression model), number of soma synapses and the mean dendritic branch length. Specifically, the
639 filter excluded the following: all segments with less than 3 soma synapses, excitatory cells with less than 17 soma synapses
640 and a mean dendritic branch length less than 35 μm , and inhibitory cells with less than 17 soma synapses and a mean dendritic
641 branch length less than 28 μm . This filter then removed approximately 14,000 segments from all of our downstream analysis.

642 **Axon Identification.** The goal of this stage in the pipeline is to identify at most one connected component subgraph that
643 represents the axon of the cell. In the absence of merge and split errors, identifying the axon would be a simple process of
644 identifying the subgraph with presynaptic connections, but un-proofread datasets pose a number of challenges to this approach.

- 645 1. Due to partial reconstruction of cells, the true axon may not exist or only the axon initial segment (AIS) may exist. In
646 this case there would be no true presynaptic connections from the cell.
- 647 2. Postsynaptic synapses on dendritic segments may be incorrectly labeled as presynaptic connections if the synapse classi-
648 fier is incorrect.

649 3. Orphan axons may be incorrectly merged onto the cell's dendrite or soma. These frequent merge errors add incorrect
650 presynaptic connections onto dendrites that make identifying the true axon subgraph more difficult. Anecdotally, if
651 the algorithm simply chose the connected component subgraph with the highest presynaptic density, this would almost
652 always be an orphan-axon-onto-dendrite merge error.

653 Our approach to axon identification is thus motivated by the following neuroscience "rules" which we implement as heuristic
654 selection criteria. Note that in this and following sections "up" or "higher" refers to the pial direction, while "down" or "lower"
655 refers to the white matter.

- 656 1. Axons can either project directly from the soma or from a proximal dendritic branch.
- 657 2. The axon is the only compartment (possibly excluding the soma) that forms presynaptic synapses.
- 658 3. The width of axon segments are typically thinner than most dendritic branches
- 659 4. Axon segments do not have spines (although boutons may have similar features)
- 660 5. Axons receive postsynaptic inputs at the AIS, but these typically have low postsynaptic density compared to dendrites.
661 (Note that we found the latter was not necessarily true in the H01 dataset and adjusted accordingly.)
- 662 6. For excitatory cells, the axon typically projects directly from the soma or from dendritic limbs that originate from the
663 deeper half of the soma.
- 664 7. For excitatory cells, the AIS starts at most 14 μm skeletal distance from the soma and the general skeleton vector of the
665 AIS typical projects downwards. The AIS does not split into multiple branches close to the soma.
- 666 8. For inhibitory cells the AIS can start much farther away (up to 80 μm skeletal distance from soma) and can come off the
667 soma or dendritic branch. The inhibitory AIS has very low postsynaptic and presynaptic density.

668 NEURD identifies candidate axonal submeshes based on a combination of these heuristics applied in a cell-type-specific
669 manner. For example, if the neuron being analyzed is excitatory, the search space of potential axonal stems is restricted to only
670 those with a projection angle from the soma greater than 70 degrees relative to the top of the volume. Candidate AIS branches
671 must fall within a maximum and minimum width range, they must have a synapse density below a threshold value, and they
672 must be within a threshold skeletal distance from the soma that dictated by the cell type. If multiple potential candidates exist,
673 the best potential axonal subgraph is selected based on longest skeletal length and closest proximity to the soma. Subgraphs
674 that meet the heuristic criteria for axons but are not chosen as the actual axon of the cell in question retain a label of "axon-like"
675 which facilitates subsequent proofreading.

676 An additional round of skeletonization is performed once the axon is detected. This re-skeletonization better captures fine
677 details of the axonal trajectory and enables auto-proofreading methods to catch more subtle axon-to-axon merge errors.

678 **Automatic Proofreading.** The goal of the automatic proofreading stage is to identify merge errors and remove all down-
679 stream branches (split errors resulting in missing portions of the axon or dendrite are not addressed by this process). NEURD
680 implements a series of heuristic proofreading rules to identify merge errors based on graph filters - configurations of nodes
681 and attributes that typically indicate merge errors. The graph filters are either directed one-hop or zero-hop configurations
682 where one-hop configurations consider a node and its downstream nodes and zero-hop configurations consider only the node's
683 features itself. These graph filters have parameters that can be tuned for axons or dendrites, excitatory or inhibitory cells, or
684 different data sources (H01 vs MICrONS). For example, a graph filter for resolving graph configurations of dendritic branches
685 with 3 or more downstream nodes is useful for the H01 dataset which has more dendritic merge errors than MICrONS. Once the
686 algorithm processes all filters, the mesh, graph nodes and synapses associated with all error branches are removed. Metadata
687 for each correction is stored for subsequent review or for training non-heuristic models.

688 The following graph filters exist for proofreading axon submeshes. Those only used for excitatory cell types are indicated.

- 689 1. **High Degree Branching:** The filter identifies any node (below a potential width threshold to exclude myelinated sections)
690 with more than two downstream nodes. The filter assumes this configuration results from a single or multiple crossing
691 axon(s) merged at an intersection point, adding 2 or more additional downstream nodes. The filter aims to identify the
692 one true downstream node. Possible upstream to downstream node pairings are filtered away if the width, synapse density
693 or skeletal angle differ by threshold amounts. If multiple downstream nodes are viable, the algorithm attempts to find a
694 downstream candidate with the best match of skeletal angle or width. If no clear winner exists, the algorithm can either
695 mark all downstream nodes as errors if the user wishes to be conservative, or can pick the best skeletal-match candidate.
696 There are more rare scenarios where a myelinated axon has 2 collateral projections protruding very close to one another
697 and these would be incorrectly filtered away, but a large majority of these occurrences are simply merge errors.

- 698 2. **Low Degree Branching Filter:** The filter processes any subgraph with one upstream node and exactly two downstream
699 nodes and is only attempted on non-myelinated axon sections (as determined by a width threshold). The method attempts
700 to find one of the following subgraph features within this directed 3 node subgraph, and if a match occurs either the
701 algorithm marks all of the downstream nodes as errors or attempts to determine the correct downstream pairing.
- 702 (a) **Axon Webbing:** An error is detected by an overly-thin mesh at the branching point of an upstream to downstream
703 node. The filter attempts to differentiate between natural branching with cell membrane that forms a "webbing like"
704 appearance as opposed to merge errors where no such thickening occurs.
 - 705 (b) **T-Intersection:** An error is detected by the presence of downstream branches that are thicker than an upstream
706 branch and the downstream branches are aligned and resemble a continuous non-branching axon segment.
 - 707 (c) **Double-Back (excitatory only):** An error is detected when a downstream node "doubles-back" towards the upstream
708 node at an unnatural skeletal angle.
 - 709 (d) **Parallel Children (or Fork Divergence):** An error is detected when the two downstream skeletons are nearly
710 parallel without a natural gap between them.
 - 711 (e) **Synapse At Branching:** An error is detected if a synapse occurs very close to the branch point between upstream
712 and downstream nodes; this usually indicates a merge of an orphan axon to a bouton segment.
- 713 3. **Width Jump:** Processes a subgraph with any number of downstream nodes (only applies to non-myelinated sections).
714 Any downstream node with an absolute width difference between segments above a certain threshold is marked as an
715 error.

716 The following graph filters exist for proofreading dendrite submeshes.

- 717 1. **Axon on Dendrite:** Nodes that were previously labeled "axon-like" during the process of axon identification (see above)
718 that do not end up in the axon submesh are marked as errors.
- 719 2. **High Degree Branching (H01 only, excitatory only, apical trunks excluded):** An error is detected using the same algorithm
720 as described above for the axonal High Degree Branching Filter except it is applied to dendritic nodes below a thresholded
721 width.
- 722 3. **Width Jump:** An error is detected using the same graph filter as described for axons but with larger width difference
723 thresholds
- 724 4. **Double Back:** An error is detected using the same algorithm and parameters as for axons above.

725 **APL Validation of Multi-Soma Splitting and High Confidence Orphan Merge Edits.** Our collaborators at APL (Johns
726 Hopkins University Applied Physics Laboratory) helped extensively validate multiple aspects of the NEURD automated proof-
727 reading workflow. In particular, they provided information about the following:

- 728 1. Validation of specific edits in the context of multi-soma splitting.
- 729 2. Data about the time that proofreaders took to evaluate these split suggestions compared to other methods.
- 730 3. Validation of specific edits focused on axon-onto-dendrite or axon-onto-axon merges

731 This information made it possible to determine whether our suggestions for multi-soma split locations speed up the process
732 (they do, more than three-fold), and if we could identify a set of heuristics and parameters for axon-on-dendrite and axon-
733 on-axon merge correction that could be executed with high confidence on the entire volume without human intervention (we
734 identified two classes of edits with performance >95%, and more than 150,000 of them have been applied to the MICrONS
735 volume to date; see Fig. 10).

736 A key method for both validating and applying automatic edits was the functionality in Neuroglancer (Perlman, 2019) which
737 allows the placement of point annotations to define a split in the PyChunkedGraph segmentation (Dorkenwald et al., 2022a,b).
738 To facilitate the proofreading process, APL created a web-based interface called NeuVue (Xenes et al., 2022) that allows for
739 the efficient queuing, review and execution of split suggestions in Neuroglancer. We built the logic required to translate mesh
740 errors identified by NEURD into split point annotations that can be executed by the NeuVue pipeline. This capability allowed
741 proofreaders at APL to not only evaluate error locations identified by NEURD, but also a proposed set of points that could
742 be subsequently executed in the PyChunkedGraph to correct the error. For a more detailed description of the NeuVue review
743 pipeline see (Xenes et al., 2022).

744 For multi-soma split edits, we generated point annotations for suggested splits that would contribute to separating neurons
745 with between 2 and 6 possible somas in a single segment. APL had both experts and trained student proofreaders review these
746 edits. The classifications for each of the edits was one of the following:

- 747 1. “yes”: same split point annotations the proofreader would have chosen.
- 748 2. “yesConditional”: split point is correct, but point annotations required very minor adjustment.
- 749 3. “errorNearby”: split point is not correct but is very close by, and split point annotations require adjustment
- 750 4. “no”: the correct split location was not at or near the suggested location

751 The expert proofreaders reviewed 5134 unique suggestions with no overlap between proofreaders, while the student proof-
752 readers reviewed 2355 suggestions with some redundancy so the same suggested edit was seen by multiple student proofreaders
753 and a majority vote determined the classification of the edit. The results of reviewing the first approximately 4000 of those ed-
754 its are shown in Fig.10a and the accuracy was determined to be 76.12% when the “yes”, “yesConditional” and “errorNearby”
755 categories were considered true positive classes. Additionally, because each split suggestion had an associated heuristic rule
756 and set of parameters that was used to generate the suggestion, we were able to show that some rules were much higher fidelity
757 than others, and that the parameters could be tuned to achieve a higher classification accuracy (Fig.10b,c).

758 To compare against the performance achieved with NEURD suggestions, multi-soma splits were also performed by expert
759 proofreaders using a tool that highlighted the path along the neural processes connecting two somas. This comparison enabled
760 us to measure if the NEURD suggestions could potentially speed up the soma-splitting process. Because a single segment with
761 a multi-soma merge could contain more than 2 somas and because different merges may require a different amount of work
762 and number of cuts to be applied, we measured the overall time spent reviewing all edits and estimated the additional time
763 that would have been required to make the slight adjustments required in the case of “yesConditional” (+30s) or “errorNearby”
764 (+60s), prior to executing the edit. Note that “no” classifications added time to the review process without any possibility of
765 contributing to an actual edit. Based on these metrics, we divided total time by total number of edits to determine the mean
766 time per edit. We then compared this metric to the mean time per edit when proofreaders used a standard pathfinding tool that
767 displayed the skeletal path connecting multiple somas, and they had to search along this path to identify errors manually. We
768 observed a more than three-fold speed up when using NEURD suggestions (Fig.10d).

769 Finally, outside the context of multi-soma splitting, APL proofreaders evaluated two kinds of merge error corrections that
770 strip orphan axons from both excitatory and inhibitory neurons: axon-on-dendrite, and high degree axon-on-axon. The feedback
771 on each error from proofreaders using the NeuVue pipeline allowed us to determine a subset of parameters that was correlated
772 with high accuracy. Thresholds for the axon-on-dendrite included minimum parent width, distance from the soma, and skeletal
773 length of the error segment. Thresholds for axon-on-axon included a minimum skeletal length, and a branching pattern that
774 resembled a two line segment crossing, where the segments are closer to perpendicular in order to make the correct connectivity
775 more obvious. For the review of orphan merge errors, an additional label was included in the true positive class: “yesPartial”,
776 which indicated that part but not all of the merge was removed by the split point annotations. The feedback from this effort
777 provided our collaboration with enough evidence to then apply nearly 150,000 of these high-confidence automatic edits back
778 into the current dynamic segmentation of the MICrONS dataset (Fig.10e).

779 **Automatic Compartment Labeling.** After automatic proofreading removes as many merge errors as possible, compartment
780 labeling is performed for excitatory cells, classifying graph submeshes as apical trunk, apical tuft, basal, and oblique. NEURD
781 first attempts to identify the apical trunk based on the geometry relative to the soma and total skeletal length. Branches
782 downstream of the end of the apical trunk are classified as the apical tuft, and branches off of the trunk with a skeletal angle
783 close to 90 degrees are labeled as oblique. If the criteria for a defined trunk is not met, then NEURD applies a generic “apical”
784 label. Other dendrites are classified as basal. Additionally, for the MICrONS dataset if the soma center is close enough to the
785 pia as defined by a depth threshold, there can be multiple generic “apical” stems protruding from the top of the soma if they
786 each meet the required width and geometry thresholds.

787 **Connectome-level features computed by NEURD.** At the level of the connectome graph, nodes represent individual single-
788 nucleus neurons and edges represent synaptic connections. In addition to the rich sub-cellular features that NEURD computes
789 for the decomposition graph of each cell, NEURD provides a variety of features at the connectome graph level:

- 790 1. Node Attributes: a wide range of global properties measured for the individual cells (compartment skeletal lengths,
791 synapses, bounding box, spine densities, synapse densities, average width, cell type, etc).
- 792 2. Edges: connections between neurons with a valid presynaptic connection and postsynaptic connection where neither
793 were filtered away in the auto-proofreading stage
- 794 3. Edge Attributes: properties for each of the presynaptic and postsynaptic neurons (compartment, skeletal/euclidean
795 distance to neuron’s soma, size, spine label) and properties of the entire synaptic connection between neurons (eu-
796clidean/skeletal distance from soma of presynaptic neuron to soma of postsynaptic neuron, etc).

797 **GNN Cell Typing.** Using PyTorch geometric software (Fey and Lenssen, 2019) we implemented a Graph Neural Network
798 architecture to build a supervised cell type classifier (including subclasses of excitatory and inhibitory cells) from dendritic
799 graph structure in the NEURD decompositions. We trained this classifier using manual cell types from the Allen Institute
800 for Brain Science (Schneider-Mizell et al., 2023). To create an input graph for the classifier we first removed the soma node
801 and filtered away the axonal subgraph and any dendritic stems with less than 25 μm of total skeletal length. Each node was
802 annotated with the following feature set:

- 803 1. Skeleton features, where theta and phi refer to polar coordinates of the skeleton vector in 3-D (skeleton_length,
804 skeleton_vector_upstream_theta, skeleton_vector_upstream_phi, skeleton_vector_downstream_theta,
805 skeleton_vector_downstream_phi)
- 806 2. Width features (width, width_upstream, width_downstream synapse)
- 807 3. Spine features (n_spines, spine_volume_sum, n_synapses_post, n_synapses_pre, n_synapses_head_postsyn,
808 n_synapses_neck_postsyn, n_synapses_shaft_postsyn, n_synapses_no_head_postsyn,
809 synapse_volume_shaft_postsyn_sum, synapse_volume_head_postsyn_sum, synapse_volume_no_head_postsyn_sum,
810 synapse_volume_neck_postsyn_sum, synapse_volume_postsyn_sum)

811 For the whole neuron classifier, the soma volume and number of soma synapses for the neuron are added to each node's
812 feature vector and also the starting stem angle (2-D angle between the vector from the soma center to the stem's root skeleton
813 point and the vector in the direction of the pia) is added to each node in every stem. For the stem-based classifier, these three
814 soma features are not included, and classification is performed on each stem individually.

815 The GNN architecture used as a 2 layer Graph Convolutional Network (128 hidden units for each layer, ReLU activation
816 function) followed with one linear layer. The aggregation and update steps were implemented using self loops and symmetric
817 normalization as shown here:

$$\mathbf{h}_u^{(k)} = \sigma \left(\mathbf{W}^{(k)} \sum_{v \in \mathcal{N}(u) \cup \{u\}} \frac{\mathbf{h}_v}{\sqrt{|\mathcal{N}(u)| |\mathcal{N}(v)|}} \right)$$

818 where $\mathbf{h}_u^{(k)}$ is the embedding for node u at layer k , $\mathcal{N}(u)$ are the neighbors for node u , $\mathbf{W}^{(k)}$ is the learned weight matrix
819 at layer k and σ is the chosen non-linearity. For graph pooling (to get one learned vector for each graph), the weighted average
820 of all nodes after the final hidden layer was taken (weighted by the skeletal length of the node). A 60%, 20%, 20% split for
821 training, validation and test sets was used for labeled datasets of $n = 873$ whole neurons and $n = 4,114$ stems (Fig.19)

822 **Proximities.** Identifying axon-dendrite proximities makes it possible to determine how often a pair of neurons capitalizes on an
823 opportunity to form a synaptic connection. Proximities are regions where the axon of one neuron passes within a few microns
824 of the dendrite of another neuron. They can be annotated with the same features (dendritic compartment, neural subtype)
825 as synapses, regardless of whether a connection was formed (Fig. 6b). Proximities are identified for all neuron pairs in the
826 volume. To reduce the number of pairwise computations, NEURD first checks whether the bounding box of the presynaptic
827 axon skeleton and postsynaptic dendrite skeleton have any overlap. In order to reduce computation time in the the MICrONS
828 dataset, presynaptic neurons are further restricted to those with at least five axonal synapses, and in the MICrONS volume
829 postsynaptic neurons are restricted to neurons with at least 1 mm of dendritic length (this restriction excludes approximately
830 1% of all MICrONS neurons).

831 The proximity calculation is performed by converting the axonal skeleton of the presynaptic neuron and the postsynaptic
832 skeleton to an array of coordinates without edges (at one micron skeletal walk resolution). A local width and compartment label
833 is associated with every point, and the soma of the postsynaptic neuron is converted to a uniform sampling of the surface mesh
834 face centers for its skeletal representation. The cleaned synaptic connections between the pre and post neuron are retrieved (if
835 there were any) and then the main proximity loop begins.

- 836 1. The closest distance between a presynaptic coordinate and postsynaptic coordinate is computed (where the distances can
837 be adjusted for postsynaptic width by subtracting the local width from the euclidean distance). This minimum distance
838 is the current proximity distance. If the current proximity distance exceeds the thresholded maximum proximity distance
839 (set at 5 μm for both the MICrONS and H01 dataset), then the loop is exited and no more proximities are computed, but
840 otherwise the workflow proceeds.
- 841 2. The following metrics are computed or collected for each proximity: The presynaptic and postsynaptic coordinate of
842 the minimum distance pair, the distance between these coordinates (the proximity distance), the compartment labels and
843 width at the postsynaptic coordinate, the presynaptic and postsynaptic skeletal walk distance, the number of spines and
844 synaptic connections within a three micron radius of the presynaptic and postsynaptic coordinates.

845 3. After these features are collected, the skeleton points within a set radius (10 microns) of the presynaptic proximity
846 coordinate are filtered away from the array of axon presynaptic coordinates.

847 4. All proximity information is saved and the loop continues until the current proximity distance exceeds the threshold.

848 **Functional Connectomics.** We considered pairs of synaptically-connected functionally-matched cells available in the MI-
849 CrONS dataset, restricting to pairs where both neurons met a minimal set of functional quality criteria (test score greater than
850 0.2 and an oracle score greater than 0.3, see [Ding et al. 2023](#); [Wang and Tolia 2023](#)). Synaptic connections were discarded
851 if they were not onto postsynaptic spines (to help guard against possible inhibitory merge errors resulting in increased con-
852 nectivity between neurons). We then divided the pairs into groups based on whether they had 1, 2, 3 or 4+ synapses between
853 them. The final number of functionally matched pyramidal pairs available from automatic proofreading alone were as follows:
854 1 synapse (5350), 2 synapses (280), 3 synapses (34) and 4+ synapses (11). We then investigated how the mean functional
855 response correlation varies as a function of the four different multi-synaptic groups. The response correlation was calculated as
856 detailed in ([Ding et al., 2023](#); [Wang and Tolia, 2023](#)) through the *in silico* response correlation of their model.

857 AUTHOR CONTRIBUTIONS

858 We use the CRediT system for author roles. Conceptualization: BC, JR. Methodology: BC, SP, JR. Software: BC. Validation: BC, DX, LMK, PKR, VAR, CAB, BW, FC.
859 Formal analysis: BC. Investigation: BC. Resources: BC, SP, ZD, PGF, EW, CP, AK, SP, JAB, ALB, DB, JB, DJB, MAC, EC, SD, LE, AH, ZJ, CJ, DK, NK, SK, KL, RL, TM,
860 GM, EM, SSM, SM, BN, SP, CMSM, WS, MT, RT, NLT, WW, JW, SCY, WY, EF, EYW, FHS, HSS, FC, NMdC, RCR, XP, AST, JR. Data Curation: BC. Writing - Original Draft:
861 BC, JR. Writing - Review Editing: BC, SD, CMSM, FC, NMdC, AST, JR. Visualization: BC, JR. Supervision: JR. Project administration: JR. Funding acquisition: HSS, NMdC,
862 RCR, AST, JR.

863 ACKNOWLEDGEMENTS

864 The authors thank David Markowitz, the IARPA MICrONS Program Manager, who coordinated this work during all three phases of the MICrONS program. We thank
865 IARPA program managers Jacob Vogelstein and David Markowitz for co-developing the MICrONS program. We thank Jennifer Wang, IARPA SETA for her assistance.
866 The work was supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/ Interior Business Center (DoI/IBC) contract numbers
867 D16PC00003, D16PC00004, D16PC0005, and 2017-17032700004-005. The views and conclusions contained herein are those of the authors and should not be interpreted
868 as necessarily representing the official policies or endorsements, either expressed or implied, of the ODNI, IARPA, or the U.S. Government. The U.S. Government is
869 authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. XP acknowledges support from NSF CAREER
870 grant IOS-1552868. XP and AST acknowledge support from NSF NeuroNex grant 1707400. AST also acknowledges support from the National Institute of Mental Health
871 and National Institute of Neurological Disorders And Stroke under Award Number U19MH114830 and National Eye Institute award numbers R01 EY026927 and Core Grant
872 for Vision Research T32-EY-002520-37. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily
873 representing the official policies or endorsements, either expressed or implied, of IARPA, DoI/IBC, or the U.S. Government.

874 **Software.** Experiments and analysis are carried out with custom built data pipelines. The data pipeline is developed in Python
875 with the following tools: DataJoint used for storing and managing data. Meshparty, CGAL and MeshLab were used for used
876 for mesh processing (CGAL and MeshLab required python wrappers). DotMotif and NetSci was used to help query graph
877 motifs. Numpy, Pandas, SciPy, Scikit-learn, and PyTorch were used for model training and statistical analysis. Matplotlib,
878 Seaborn, Ipyvolume, and Neuroglancer were used for graphical visualization. Jupyter, Docker, and Kubernetes were used for
879 code development and deployment.

880 Data availability.

881 All MICrONS data have already been released on BossDB (<https://bosddb.org/project/microns-minnie>, please also see
882 <https://www.microns-explorer.org/cortical-mm3>) for details.

883 Code availability.

884 All code will be available at <https://github.com/reimerlab/NEURD.git>

885 **Conflict of Interest.** XP is a co founder of Upload AI, LLC, a company in which he has financial interests. AST is co founder
886 of Vathes Inc., and UploadAI LLC companies in which he has financial interests. JR is co founder of Vathes Inc., and UploadAI
887 LLC companies in which he has financial interests.

888 Bibliography

889 Alexander Shapson-Coe, Michal Januszewski, Daniel R. Berger, Art Pope, Yuelong Wu, Tim Blakely, Richard L. Schalek, Peter Li, Shuohong Wang, Jeremy Maitin-Shepard, Neha Karlupia, Sven
890 Dorkenwald, Evelina Sjostedt, Laramie Leavitt, Dongil Lee, Luke Bailey, Angerica Fitzmaurice, Rohin Kar, Benjamin Field, Hank Wu, Julian Wagner-Carena, David Aley, Joanna Lau, Zudi Lin, Donglai
891 Wei, Hanspeter Pfister, Adi Peleg, Viren Jain, and Jeff W. Lichtman. A connectomic study of a petascale fragment of human cerebral cortex. Technical report, bioRxiv, May 2021.
892 MICrONS Consortium, J. Alexander Bae, Mahaly Baptiste, Agnes L. Bodor, Derrick Brittain, JoAnn Buchanan, Daniel J. Bumbarger, Manuel A. Castro, Brendan Celli, Erick Cobos, Forrest Collman,
893 Nuno Maçarico da Costa, Sven Dorkenwald, Leila Elabbady, Paul G. Fahey, Tim Fliss, Emmanouil Froudarakis, Jay Gager, Clare Gamlin, Akhilesh Halageri, James Hebditch, Zhen Jia, Chris Jordan,
894 Daniel Kapner, Nico Kemnitz, Sam Kinn, Selden Koolman, Kai Kuehner, Kisuk Lee, Kai Li, Ran Lu, Thomas Macrina, Gayathri Mahalingam, Sarah McReynolds, Elanine Miranda, Eric Mitchell,
895 Shanka Subhra Mondal, Merlin Moore, Shang Mu, Taliah Muhammad, Barak Nehoran, Oluwaseun Ogedengbe, Christos Papadopoulos, Stelios Papadopoulos, Saumil Patel, Xaq Pitkow, Sergiy
896 Popovych, Anthony Ramos, R. Clay Reid, Jacob Reimer, Casey M. Schneider-Mizell, H. Sebastian Seung, Ben Silverman, William Silversmith, Amy Sterling, Fabian H. Sinz, Cameron L. Smith,
897 Shelby Suckow, Marc Takeno, Zheng H. Tan, Andreas S. Tolia, Russel Torres, Nicholas L. Turner, Edgar Y. Walker, Tianyu Wang, Grace Williams, Sarah Williams, Kyle Willie, Ryan Willie, William
898 Wong, Jingpeng Wu, Chris Xu, Runzhe Yang, Dimitri Yatsenko, Fei Ye, Wenjing Yin, and Szi-chieh Yu. Functional connectomics spanning multiple areas of mouse visual cortex, August 2021. Pages:
899 2021.07.28.454025 Section: New Results.
900 Kisuk Lee, Jonathan Zung, Peter Li, Viren Jain, and H. Sebastian Seung. Superhuman Accuracy on the SNEMI3D Connectomics Challenge, May 2017. arXiv:1706.00120 [cs].
901 Jingpeng Wu, William M. Silversmith, Kisuk Lee, and H. Sebastian Seung. Chunkflow: hybrid cloud processing of large 3D images by convolutional nets. *Nature Methods*, 18(4):328–330, April 2021.
902 ISSN 1548-7105. doi: 10.1038/s41592-021-01088-5. Number: 4 Publisher: Nature Publishing Group.
903 Ran Lu, Aleksandar Zlateski, and H. Sebastian Seung. Large-scale image segmentation based on distributed clustering algorithms, June 2021. arXiv:2106.10795 [cs].

- 904 Thomas Macrina, Kisuk Lee, Ran Lu, Nicholas L. Turner, Jingpeng Wu, Sergiy Popovych, William Silversmith, Nico Kemnitz, J. Alexander Bae, Manuel A. Castro, Sven Dorkenwald, Akhilesh Halageri,
905 Zhen Jia, Chris Jordan, Kai Li, Eric Mitchell, Shanka Subhra Mondal, Shang Mu, Barak Nehoran, William Wong, Szi-chieh Yu, Agnes L. Bodor, Derrick Brittain, JoAnn Buchanan, Daniel J. Bumbarger,
906 Erick Cobos, Forrest Collman, Leila Elabbady, Paul G. Fahey, Emmanouil Froudarakis, Daniel Kapner, Sam Kinn, Gayathri Mahalingam, Stelios Papadopoulos, Saumil Patel, Casey M. Schneider-
907 Mizell, Fabian H. Sinz, Marc Takeno, Russel Torres, Wenjing Yin, Xaq Pitkow, Jacob Reimer, Andreas S. Tolias, R. Clay Reid, Nuno Maçarico da Costa, and H. Sebastian Seung. Petascale neural
908 circuit reconstruction: automated methods, August 2021. Pages: 2021.08.04.455162 Section: New Results.
- 909 Sven Dorkenwald, Nicholas L. Turner, Thomas Macrina, Kisuk Lee, Ran Lu, Jingpeng Wu, Agnes L. Bodor, Adam A. Bleckert, Derrick Brittain, Nico Kemnitz, William M Silversmith, Dodam Ih, Jonathan
910 Zung, Aleksandar Zlateski, Ignacio Tartavull, Szi-Chieh Yu, Sergiy Popovych, William Wong, Manuel Castro, Chris S Jordan, Alyssa M Wilson, Emmanouil Froudarakis, JoAnn Buchanan, Marc M
911 Takeno, Russel Torres, Gayathri Mahalingam, Forrest Collman, Casey M Schneider-Mizell, Daniel J Bumbarger, Yang Li, Lynne Becker, Shelby Suckow, Jacob Reimer, Andreas S Tolias, Nuno
912 Macarico da Costa, R Clay Reid, and H Sebastian Seung. Binary and analog variation of synapses between cortical pyramidal neurons. *eLife*, 11:e76120, November 2022a. ISSN 2050-084X. doi:
913 10.7554/eLife.76120. Publisher: eLife Sciences Publications, Ltd.
- 914 Sven Dorkenwald, Claire E. McKellar, Thomas Macrina, Nico Kemnitz, Kisuk Lee, Ran Lu, Jingpeng Wu, Sergiy Popovych, Eric Mitchell, Barak Nehoran, Zhen Jia, J. Alexander Bae, Shang Mu, Dodam
915 Ih, Manuel Castro, Oluwaseun Ogedengbe, Akhilesh Halageri, Kai Kuehner, Amy R. Sterling, Zoe Ashwood, Jonathan Zung, Derrick Brittain, Forrest Collman, Casey Schneider-Mizell, Chris Jordan,
916 William Silversmith, Christa Baker, David Deutsch, Lucas Encarnacion-Rivera, Sandeep Kumar, Austin Burke, Doug Bland, Jay Gager, James Hebditch, Selden Koolman, Merlin Moore, Sarah
917 Morejohn, Ben Silverman, Kyle Willie, Ryan Willie, Szi-chieh Yu, Mala Murthy, and H. Sebastian Seung. FlyWire: online community for whole-brain connectomics. *Nature Methods*, 19(1):119–128,
918 January 2022b. ISSN 1548-7105. doi: 10.1038/s41592-021-01330-0. Number: 1 Publisher: Nature Publishing Group.
- 919 Wei-Chung Allen Lee, Vincent Bonin, Michael Reed, Brett J. Graham, Greg Hood, Katie Glatfelder, and R. Clay Reid. Anatomy and function of an excitatory network in the visual cortex. *Nature*, 532
920 (7599):370–374, April 2016. ISSN 1476-4687. doi: 10.1038/nature17192. Number: 7599 Publisher: Nature Publishing Group.
- 921 Ho Ko, Sonja B. Hofer, Bruno Pichler, Katherine A. Buchanan, P. Jesper Sjöström, and Thomas D. Mrsic-Flogel. Functional specificity of local synaptic connections in neocortical networks. *Nature*, 473
922 (7345):87–91, May 2011. ISSN 1476-4687. doi: 10.1038/nature09880. Number: 7345 Publisher: Nature Publishing Group.
- 923 Zhuokun Ding, Paul Fahey, Stelios Papadopoulos, Eric Wang, Alexander Kunin, Brendan Celli, Christos Papadopoulos, Andersen Chang, Jiakun Fu, Zhiwei Ding, Saumil Patel, MICrONS Consortium,
924 Emmanouil Froudarakis, Sebastian Seung, Forrest Collman, Nuno Maçarico da Costa, Clay Reid, Edgar Walker, Xaq Pitkow, Jacob Reimer, and Andreas Tolias. Functional
925 connectomics reveals general wiring rule in mouse visual cortex, March 2023.
- 926 Andrea Giovannucci, Johannes Friedrich, Pat Gunn, Jérémie Kalfon, Brandon L Brown, Sue Ann Koay, Jiannis Taxisidis, Farzaneh Najafi, Jeffrey L Gauthier, Pengcheng Zhou, Baljit S Khakh, David W
927 Tank, Dmitri B Chklovskii, and Eftychios A Pnevmatikakis. CalmAn: an open source tool for scalable calcium imaging data analysis. *eLife*, 8:e38173, January 2019. ISSN 2050-084X. doi:
928 10.7554/eLife.38173. Publisher: eLife Sciences Publications, Ltd.
- 929 Marius Pachitariu, Carsen Stringer, Mario Dipoppa, Sylvia Schröder, L. Federico Rossi, Henry Dalgleish, Matteo Carandini, and Kenneth D. Harris. Suite2p: beyond 10,000 neurons with standard
930 two-photon microscopy, July 2017. Pages: 061507 Section: New Results.
- 931 Marius Pachitariu, Shashwat Sridhar, and Carsen Stringer. Solving the spike sorting problem with Kilosort, January 2023. Pages: 2023.01.07.523036 Section: New Results.
- 932 Jason E. Chung, Jeremy F. Magland, Alex H. Barnett, Vanessa M. Tolosa, Angela C. Tooker, Kye Y. Lee, Kedar G. Shah, Sarah H. Felix, Loren M. Frank, and Leslie F. Greengard. A Fully Automated
933 Approach to Spike Sorting. *Neuron*, 95(6):1381–1394.e6, September 2017. ISSN 0896-6273. doi: 10.1016/j.neuron.2017.08.030.
- 934 Alexander Mathis, Pranav Mamidanna, Kevin M. Cury, Taiga Abe, Venkatesh N. Murthy, Mackenzie Weygant Mathis, and Matthias Bethge. DeepLabCut: markerless pose estimation of user-defined
935 body parts with deep learning. *Nature Neuroscience*, 21(9):1281–1289, September 2018. ISSN 1546-1726. doi: 10.1038/s41593-018-0209-y. Number: 9 Publisher: Nature Publishing Group.
- 936 Talmo D. Pereira, Nathaniel Tabris, Arie Matsliah, David M. Turner, Junyu Li, Shruthi Ravindranath, Eleni S. Papadopyannis, Edna Normand, David S. Deutsch, Z. Yan Wang, Grace C. McKenzie-Smith,
937 Catalin C. Mitelet, Marielisa Diez Castro, John D’Uva, Mikhail Kislin, Dan H. Sanes, Sarah D. Kocher, Samuel S.-H. Wang, Annegret L. Falkner, Joshua W. Shaevitz, and Mala Murthy. SLEAP: A deep
938 learning system for multi-animal pose tracking. *Nature Methods*, 19(4):486–495, April 2022. ISSN 1548-7105. doi: 10.1038/s41592-022-01426-1. Number: 4 Publisher: Nature Publishing Group.
- 939 Jeffrey E. Markowitz, Winthrop J. Gillis, Celia C. Beron, Shay Q. Neufeld, Keiramarie Robertson, Neha D. Bhagat, Ralph E. Peterson, Emalee Peterson, Minsuk Hyun, Scott W. Linderman, Bernardo L.
940 Sabatini, and Sandeep Robert Datta. The Striatum Organizes 3D Behavior via Moment-to-Moment Action Selection. *Cell*, 174(1):44–58.e17, June 2018. ISSN 0092-8674, 1097-4172. doi:
941 10.1016/j.cell.2018.04.019. Publisher: Elsevier.
- 942 Ruben Dries, Qian Zhu, Rui Dong, Chee-Huat Linus Eng, Huipeng Li, Kan Liu, Yuntian Fu, Tianxiao Zhao, Arpan Sarkar, Feng Bao, Rani E. George, Nico Pierson, Long Cai, and Guo-Cheng Yuan.
943 Giotto: a toolbox for integrative analysis and visualization of spatial expression data. *Genome Biology*, 22(1):78, March 2021. ISSN 1474-760X. doi: 10.1186/s13059-021-02286-2.
- 944 Giovanni Palla, Hannah Spitzer, Michal Klein, David Fischer, Anna Christina Schaar, Louis Benedikt Kuemmerle, Sergei Rybakov, Ignacio L. Ibarra, Olle Holmberg, Isaac Virshup, Mohammad Lotfollahi,
945 Sabrina Richter, and Fabian J. Theis. Squidpy: a scalable framework for spatial omics analysis. *Nature Methods*, 19(2):171–178, February 2022. ISSN 1548-7105. doi: 10.1038/s41592-021-01358-2.
946 Number: 2 Publisher: Nature Publishing Group.
- 947 Wenjing Yin, Derrick Brittain, Jay Borseth, Marie E. Scott, Derric Williams, Jedediah Perkins, Christopher S. Own, Matthew Murfitt, Russel M. Torres, Daniel Kapner, Gayathri Mahalingam, Adam Bleckert,
948 Daniel Castelli, David Reid, Wei-Chung Allen Lee, Brett J. Graham, Marc Takeno, Daniel J. Bumbarger, Colin Farrell, R. Clay Reid, and Nuno Macarico da Costa. A petascale automated imaging
949 pipeline for mapping neuronal circuits with high-throughput transmission electron microscopy. *Nature Communications*, 11(1):4949, October 2020. ISSN 2041-1723. doi: 10.1038/s41467-020-18659-3.
950 Number: 1 Publisher: Nature Publishing Group.
- 951 Juan Carlos Tapia, Narayanan Kasthuri, Kenneth J. Hayworth, Richard Schalek, Jeff W. Lichtman, Stephen J. Smith, and JoAnn Buchanan. High-contrast en bloc staining of neuronal tissue for field
952 emission scanning electron microscopy. *Nature Protocols*, 7(2):193–206, February 2012. ISSN 1750-2799. doi: 10.1038/nprot.2011.439. Number: 2 Publisher: Nature Publishing Group.
- 953 R. Schalek, N. Kasthuri, K. Hayworth, D. Berger, J. Tapia, J. Morgan, S. Turaga, E. Fagerholm, H. Seung, and J. Lichtman. Development of High-Throughput, High-Resolution 3D Reconstruction of
954 Large-Volume Biological Tissue Using Automated Tape Collection Ultramicrotomy and Scanning Electron Microscopy. *Microscopy and Microanalysis*, 17(S2):966–967, July 2011. ISSN 1435-8115,
955 1431-9276. doi: 10.1017/S1431927611005708. Publisher: Cambridge University Press.
- 956 Christopher S Own, Matthew F Murfitt, Lawrence S Own, Derrick Brittain, Nuno da Costa, R Clay Reid, David G C Hildebrand, Brett Graham, and Wei-Chung Allen Lee. Reel-to-Reel Electron Microscopy:
957 Latency-Free Continuous Imaging of Large Sample Volumes. *Microscopy and Microanalysis*, 21(S3):157–158, August 2015. ISSN 1431-9276. doi: 10.1017/S1431927615001580.
- 958 Kenneth J. Hayworth, Josh L. Morgan, Richard Schalek, Daniel R. Berger, David G. C. Hildebrand, and Jeff W. Lichtman. Imaging ATUM ultrathin section libraries with WaterMapper: a multi-scale
959 approach to EM reconstruction of neural circuits. *Frontiers in Neural Circuits*, 8, 2014. ISSN 1662-5110.
- 960 Andreas Fabri, Fernando Cacciola, Philipp Moeller, and Ron Wein. CGAL and the Boost Graph Library. In *CGAL User and Reference Manual*. CGAL Editorial Board, 5.5.2 edition, 2023.
- 961 Fernando Cacciola, Mael Rouxel-Labbé, Baskin Şenbaşlar, and Julian Komaromy. Triangulated Surface Mesh Simplification. In *CGAL User and Reference Manual*. CGAL Editorial Board, 5.5.2 edition,
962 2023.
- 963 Xiang Gao, Sébastien Lorient, and Andrea Tagliasacchi. Triangulated Surface Mesh Skeletonization. In *CGAL User and Reference Manual*. CGAL Editorial Board, 5.5.2 edition, 2023.
- 964 Sven Dorkenwald. sdorkenw/MeshParty, March 2022. original-date: 2018-09-11T23:38:46Z.
- 965 Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. Exploring Network Structure, Dynamics, and Function using NetworkX. In Gaël Varoquaux, Travis Vaught, and Jarrod Millman, editors, *Proceedings
966 of the 7th Python in Science Conference*, pages 11 – 15, Pasadena, CA USA, 2008.
- 967 Luis Pastor, Sofia Bayona, Juan P. Brito, María Cuevas, Isabel Feraud, Sergio Emilio Galindo, Juan José García-Cantero, Francisco González de Quevedo, Susana Mata, Oscar David Robles, Angel
968 Rodríguez, Pablo Toharia, and Ana Zdravkovic. A Unified Framework for Neuroscience Morphological Data Visualization. *Applied Sciences*, 11(10):4652, January 2021. ISSN 2076-3417. doi:
969 10.3390/app11104652. Number: 10 Publisher: Multidisciplinary Digital Publishing Institute.
- 970 Casey M Schneider-Mizell, Stephan Gerhard, Mark Longair, Tom Kazimiers, Feng Li, Maarten F Zwart, Andrew Champion, Frank M Midgley, Richard D Fetter, Stephan Saalfeld, and Albert Cardona.
971 Quantitative neuroanatomy for connectomics in *Drosophila*. *eLife*, 5:e12059, March 2016. ISSN 2050-084X. doi: 10.7554/eLife.12059. Publisher: eLife Sciences Publications, Ltd.
- 972 Daniel Xenos, Lindsey M. Kitchell, Patricia K. Rivlin, Rachel Brodsky, Hannah Gooden, Justin Joyce, Diego Luna, Raphael Norman-Tenazas, Devin Ramsden, Kevin Romero, Victoria Rose, Marisel
973 Villafane-Delgado, William Gray-Roncal, and Brock Wester. NeuVue: A Framework and Workflows for High-Throughput Electron Microscopy Connectomics Proofreading, July 2022. Pages:
974 2022.07.18.500521 Section: New Results.
- 975 Casey M. Schneider-Mizell, Agnes Bodor, Derrick Brittain, JoAnn Buchanan, Daniel J. Bumbarger, Leila Elabbady, Daniel Kapner, Sam Kinn, Gayathri Mahalingam, Sharmishta Seshamani, Shelby
976 Suckow, Marc Takeno, Russel Torres, Wenjing Yin, Sven Dorkenwald, J. Alexander Bae, Manuel A. Castro, Paul G. Fahey, Emmanouil Froudakis, Akhilesh Halageri, Zhen Jia, Chris Jordan, Nico
977 Kemnitz, Kisuk Lee, Kai Li, Ran Lu, Thomas Macrina, Eric Mitchell, Shanka Subhra Mondal, Shang Mu, Barak Nehoran, Stelios Papadopoulos, Saumil Patel, Xaq Pitkow, Sergiy Popovych, William
978 Silversmith, Fabian H. Sinz, Nicholas L. Turner, William Wong, Jingpeng Wu, Szi-chieh Yu, MICrONS Consortium, Jacob Reimer, Andreas S. Tolias, H. Sebastian Seung, R. Clay Reid, Forrest
979 Collman, and Nuno Maçarico da Costa. Cell-type-specific inhibitory circuitry from a connectomic census of mouse visual cortex, February 2023. Pages: 2023.01.23.525290 Section: New Results.
- 980 Casey M Schneider-Mizell, Agnes L Bodor, Forrest Collman, Derrick Brittain, Adam Bleckert, Sven Dorkenwald, Nicholas L Turner, Thomas Macrina, Kisuk Lee, Ran Lu, Jingpeng Wu, Jun Zhuang,
981 Anirban Nandi, Brian Hu, JoAnn Buchanan, Marc M Takeno, Russel Torres, Gayathri Mahalingam, Daniel J Bumbarger, Yang Li, Thomas Chartrand, Nico Kemnitz, William M Silversmith, Dodam
982 Ih, Jonathan Zung, Aleksandar Zlateski, Ignacio Tartavull, Sergiy Popovych, William Wong, Manuel Castro, Chris S Jordan, Emmanouil Froudarakis, Lynne Becker, Shelby Suckow, Jacob Reimer,
983 Andreas S Tolias, Costas A Anastasiou, H Sebastian Seung, R Clay Reid, and Nuno Maçarico da Costa. Structure and function of axo-axonic inhibition. *eLife*, 10:e73783, December 2021. ISSN
984 2050-084X. doi: 10.7554/eLife.73783. Publisher: eLife Sciences Publications, Ltd.
- 985 Marissa A. Weis, Stelios Papadopoulos, Laura Hansel, Timo Lüddecke, Brendan Celli, Paul G. Fahey, J. Alexander Bae, Agnes L. Bodor, Derrick Brittain, JoAnn Buchanan, Daniel J. Bumbarger,
986 Manuel A. Castro, Erick Cobos, Forrest Collman, Nuno Maçarico da Costa, Sven Dorkenwald, Leila Elabbady, Emmanouil Froudarakis, Akhilesh Halageri, Zhen Jia, Chris Jordan, Dan Kapner, Nico
987 Kemnitz, Sam Kinn, Kisuk Lee, Kai Li, Ran Lu, Thomas Macrina, Gayathri Mahalingam, Eric Mitchell, Shanka Subhra Mondal, Shang Mu, Barak Nehoran, Saumil Patel, Xaq Pitkow, Sergiy Popovych,
988 R. Clay Reid, Casey M. Schneider-Mizell, H. Sebastian Seung, William Silversmith, Fabian H. Sinz, Marc Takeno, Russel Torres, Nicholas L. Turner, William Wong, Jingpeng Wu, Wenjing Yin,
989 Szi-chieh Yu, Jacob Reimer, Andreas S. Tolias, and Alexander S. Ecker. Large-scale unsupervised discovery of excitatory morphological cell types in mouse visual cortex, December 2022. Pages:

- 990 2022.12.22.521541 Section: New Results.
- 991 Sven Dorkenwald, Peter H. Li, Michał Januszewski, Daniel R. Berger, Jeremy Maitin-Shepard, Agnes L. Bodor, Forrest Collman, Casey M. Schneider-Mizell, Nuno Maçarico da Costa, Jeff W. Lichtman,
992 and Viren Jain. Multi-Layered Maps of Neuropil with Segmentation-Guided Contrastive Learning, November 2022c. Pages: 2022.03.29.486320 Section: New Results.
- 993 A. Peters and M. L. Feldman. The projection of the lateral geniculate nucleus to area 17 of the rat cerebral cortex. I. General description. *Journal of Neurocytology*, 5(1):63–84, February 1976. ISSN
994 0300-4864. doi: 10.1007/BF01176183.
- 995 K A Martin and D Whitteridge. Form, function and intracortical projections of spiny neurones in the striate visual cortex of the cat. *The Journal of Physiology*, 353:463–504, August 1984. ISSN 0022-3751.
- 996 Federico Scala, Dmitry Kobak, Matteo Bernabucci, Yves Bernaerts, Cathryn René Cadwell, Jesus Ramon Castro, Leonard Hartmanis, Xiaolong Jiang, Sophie Laturnus, Elanire Miranda, Shalaka
997 Mulherkar, Zheng Huan Tan, Zizhen Yao, Hongkui Zeng, Rickard Sandberg, Philipp Berens, and Andreas S. Tolias. Phenotypic variation of transcriptomic cell types in mouse motor cortex. *Nature*,
998 598(7879):144–150, October 2021. ISSN 1476-4687. doi: 10.1038/s41586-020-2907-3. Number: 7879 Publisher: Nature Publishing Group.
- 999 Hanchuan Peng, Peng Xie, Lijuan Liu, Xiuli Kuang, Yimin Wang, Lei Qu, Hui Gong, Shengdian Jiang, Anan Li, Zongcai Ruan, Liya Ding, Zizhen Yao, Chao Chen, Mengya Chen, Tanya L. Daigle, Rachel
1000 Dalley, Zhangcan Ding, Yanjun Duan, Aaron Feiner, Ping He, Chris Hill, Karla E. Hirokawa, Guodong Hong, Lei Huang, Sara Kebede, Hsien-Chi Kuo, Rachael Larsen, Phil Lesnar, Longfei Li, Qi Li,
1001 Xiangning Li, Yaoyao Li, Yuanyuan Li, An Liu, Donghuan Lu, Stephanie Mok, Lydia Ng, Thuc Nghi Nguyen, Qiang Ouyang, Jintao Pan, Elise Shen, Yuanyuan Song, Susan M. Sunkin, Bosiljka Tasic,
1002 Matthew B. Veldman, Wayne Wakeman, Wan Wan, Peng Wang, Quanxin Wang, Tao Wang, Yaping Wang, Feng Xiong, Wei Xiong, Wenjie Xu, Min Ye, Lulu Yin, Yang Yu, Jia Yuan, Jing Yuan, Zhixi
1003 Yun, Shaoqun Zeng, Shichen Zhang, Sujun Zhao, Zijun Zhao, Zhi Zhou, Z. Josh Huang, Luke Esposito, Michael J. Hawrylycz, Staci A. Sorensen, X. William Yang, Yefeng Zheng, Zhongze Gu, Wei
1004 Xie, Christof Koch, Qingming Luo, Julie A. Harris, Yun Wang, and Hongkui Zeng. Morphological diversity of single neurons in molecularly defined cell types. *Nature*, 598(7879):174–181, October
1005 2021. ISSN 1476-4687. doi: 10.1038/s41586-021-03941-1.
- 1006 Nathan W. Gouwens, Staci A. Sorensen, Fahimeh Batfzadeh, Agata Budzillo, Brian R. Lee, Tim Jarsky, Lauren Afifler, Katherine Baker, Eliza Barkan, Kyla Berry, Darren Bertagnolli, Kris Bickley, Jasmine
1007 Bomben, Thomas Braun, Krissy Brouner, Tamara Casper, Kirsten Crichton, Tanya L. Daigle, Rachel Dalley, Rebecca A. de Frates, Nick Dee, Tsega Desta, Samuel Dingman Lee, Nadezhda Dotson,
1008 Tom Egdorf, Lauren Ellingwood, Rachel Enstrom, Luke Esposito, Colin Farrell, David Feng, Olivia Fong, Rohan Gala, Clare Gamlin, Amanda Gary, Alexandra Glandon, Jeff Goldy, Melissa Gorham,
1009 Lucas Grayback, Hong Gu, Kristen Hadley, Michael J. Hawrylycz, Alex M. Henry, DiJon Hill, Madie Hupp, Sara Kebede, Tae Kyung Kim, Lisa Kim, Matthew Kroll, Changkyu Lee, Katherine E. Link,
1010 Matthew Mallory, Rusty Mann, Michelle Maxwell, Medea McGraw, Delissa McMillen, Alice Mukora, Lindsay Ng, Lydia Ng, Kiet Ngo, Philip R. Nicovich, Aaron Oldre, Daniel Park, Hanchuan Peng,
1011 Osnat Penn, Thanh Pham, Alice Pom, Zoran Popović, Lydia Potekhina, Ramkumar Rajanbabu, Shea Ransford, David Reid, Christine Rimorin, Miranda Robertson, Kara Ronellenfitch, Augustin Ruiz,
1012 David Sandman, Kimberly Smith, Josef Sulc, Susan M. Sunkin, Aaron Szafer, Michael Tieu, Amy Torkelson, Jessica Trinh, Herman Tung, Wayne Wakeman, Katelyn Ward, Grace Williams, Zhi Zhou,
1013 Jonathan T. Ting, Anton Arkhipov, Uygur Sümbül, Ed S. Lein, Christof Koch, Zizhen Yao, Bosiljka Tasic, Jim Berg, Gabe J. Murphy, and Hongkui Zeng. Integrated Morphoelastic and Transcriptomic
1014 Classification of Cortical GABAergic Cells. *Cell*, 183(4):935–953.e19, November 2020. ISSN 1097-4172. doi: 10.1016/j.cell.2020.09.057.
- 1015 Leila Elabbady, Sharmishta Seshamani, Shang Mu, Gayathri Mahalingam, Casey Schneider-Mizell, Agnes Bodor, J. Alexander Bae, Derrick Brittain, JoAnn Buchanan, Daniel J. Bumbarger, Manuel A.
1016 Castro, Erick Cobos, Sven Dorkenwald, Paul G. Fahey, Emmanouil Froudarakis, Akhilesh Halgeri, Zhen Jia, Chris Jordan, Dan Kapner, Nico Kemnitz, Sam Kinn, Kisuk Lee, Kai Li, Ran Lu,
1017 Thomas Macrina, Eric Mitchell, Shanka Subhra Mondal, Barak Nehoran, Stelios Papadopoulos, Saamil Patel, Xaq Pitkow, Sergiy Popovych, Jacob Reimer, William Silversmith, Fabian H. Sinz, Marc
1018 Takeno, Russel Torres, Nicholas Turner, William Wong, Jिंगpeng Wu, Wenjing Yin, Szi-chieh Yu, Andreas Tolias, H. Sebastian Yin, R. Clay Reid, Nuno Maçarico Da Costa, and Forrest Collman.
1019 Quantitative Census of Local Somatic Features in Mouse Visual Cortex, July 2022. Pages: 2022.07.20.499976 Section: New Results.
- 1020 Khaled Al-Thelaya, Marco Agus, Nauman Ullah Gilal, Yin Yang, Giovanni Pintore, Enrico Gobetti, Corrado Galic, Pierre J. Magistretti, William Mifsud, and Jens Schneider. InShaDe: Invariant Shape
1021 Descriptors for visual 2D and 3D cellular and nuclear shape analysis and classification. *Computers & Graphics*, 98:105–125, August 2021. ISSN 0097-8493. doi: 10.1016/j.cag.2021.04.037.
- 1022 Sharmishta Seshamani, Leila Elabbady, Casey Schneider-Mizell, Gayathri Mahalingam, Sven Dorkenwald, Agnes Bodor, Thomas Macrina, Daniel Bumbarger, JoAnn Buchanan, Marc Takeno, Wenjing
1023 Yin, Derrick Brittain, Russel Torres, Daniel Kapner, Kisuk Lee, Ran Lu, Jिंगpeng Wu, Nuno daCosta, Clay Reid, and Forrest Collman. Automated Neuron Shape Analysis from Electron Microscopy,
1024 May 2020. arXiv:2006.01010 [cs].
- 1025 R. Azouz, C. M. Gray, L. G. Nowak, and D. A. McCormick. Physiological properties of inhibitory interneurons in cat striate cortex. *Cerebral Cortex (New York, N.Y.: 1991)*, 7(6):534–545, September 1997.
1026 ISSN 1047-3211. doi: 10.1093/cercor/7.6.534.
- 1027 Xuerong Xiao, Maja Djurisic, Assaf Hoogi, Richard W. Sapp, Carla J. Shatz, and Daniel L. Rubin. Automated dendritic spine detection using convolutional neural networks on maximum intensity projected
1028 microscopic volumes. *Journal of neuroscience methods*, 309:25–34, November 2018. ISSN 0165-0270. doi: 10.1016/j.jneumeth.2018.08.019.
- 1029 Meghan K. Driscoll, Erik S. Wolf, Andrew R. Jamieson, Kevin M. Dean, Tadamoto Isogai, Reto Fiolka, and Gaudenz Danuser. Robust and automated detection of subcellular morphological motifs in 3D
1030 microscopy images. *Nature Methods*, 16(10):1037–1044, October 2019. ISSN 1548-7105. doi: 10.1038/s41592-019-0539-z. Number: 10 Publisher: Nature Publishing Group.
- 1031 Firdaus Janoos, Kishore Mosaliganti, Xiaoyin Xu, Raghu Machiraju, Kun Huang, and Stephen T. C. Wong. Robust 3D reconstruction and identification of dendritic spines from optical microscopy imaging.
1032 *Medical Image Analysis*, 13(1):167–179, February 2009. ISSN 1361-8423. doi: 10.1016/j.media.2008.06.019.
- 1033 Peng Shi, Yue Huang, and Jinsheng Hong. Automated three-dimensional reconstruction and morphological analysis of dendritic spines based on semi-supervised learning. *Biomedical Optics Express*,
1034 5(5):1541–1553, April 2014. ISSN 2156-7085. doi: 10.1364/BOE.5.001541.
- 1035 Subhadip Basu, Punam Kumar Saha, Matylda Roszkowska, Marta Magnowska, Ewa Baczyńska, Nirmal Das, Dariusz Plewczynski, and Jakob Włodarczyk. Quantitative 3-D morphometric analysis of
1036 individual dendritic spines. *Scientific Reports*, 8(1):3545, February 2018. ISSN 2045-2322. doi: 10.1038/s41598-018-21753-8. Number: 1 Publisher: Nature Publishing Group.
- 1037 Ruth Benavides-Piccone, Isabel Fernaud-Espinosa, Victor Robles, Rafael Yuste, and Javier DeFelipe. Age-Based Comparison of Human Dendritic Spine Structure Using Complete Three-Dimensional
1038 Reconstructions. *Cerebral Cortex*, 23(8):1798–1810, August 2013. ISSN 1047-3211. doi: 10.1093/cercor/bhs154.
- 1039 K. M. Harris and J. K. Stevens. Dendritic spines of CA 1 pyramidal cells in the rat hippocampus: serial electron microscopy with reference to their biophysical characteristics. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 9(8):2982–2997, August 1989. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.09-08-02982.1989.
- 1040 Jon I. Arellano, Ruth Benavides-Piccone, Javier DeFelipe, and Rafael Yuste. Ultrastructure of Dendritic Spines: Correlation Between Synaptic and Spine Morphologies. *Frontiers in Neuroscience*, 1(1):
1041 131–143, October 2007. ISSN 1662-4548. doi: 10.3389/neuro.01.1.1.010.2007.
- 1042 Xiaojun Wang, Jason Tucciarone, Siqi Jiang, Fangfang Yin, Bor-Shuen Wang, Dingkang Wang, Yao Jia, Xueyan Jia, Yuxin Li, Tao Yang, Zhengchao Xu, Masood A. Akram, Yusu Wang, Shaoqun
1043 Zeng, Giorgio A. Ascoli, Partha Mitra, Hui Gong, Qingming Luo, and Z. Josh Huang. Genetic Single Neuron Anatomy Reveals Fine Granularity of Cortical Axo-Axonic Cells. *Cell Reports*, 26(11):
1044 3145–3159.e5, March 2019. ISSN 2211-1247. doi: 10.1016/j.celrep.2019.02.040.
- 1045 Melis Inan, Lidia Blázquez-Llorca, Angel Merchán-Pérez, Stewart A. Anderson, Javier DeFelipe, and Rafael Yuste. Dense and Overlapping Innervation of Pyramidal Neurons by Chandelier Cells. *The Journal of Neuroscience*, 33(5):1907–1914, January 2013. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.4049-12.2013.
- 1046 Gregg A. Wildenberg, Matt R. Rosen, Jack Lundell, Dawn Paukner, David J. Freedman, and Narayanan Kasthuri. Primate neuronal connections are sparse in cortex as compared to mouse. *Cell Reports*,
1047 36(11):109709, September 2021. ISSN 2211-1247. doi: 10.1016/j.celrep.2021.109709.
- 1048 Charles E. Ribak, James E. Vaughn, and Roert P. Barber. Immunocytochemical localization of GABAergic neurones at the electron microscopical level. *The Histochemical Journal*, 13(4):555–582, July
1049 1981. ISSN 1573-6865. doi: 10.1007/BF01002711.
- 1050 Javier DeFelipe and Isabel Fariñas. The pyramidal neuron of the cerebral cortex: Morphological and chemical characteristics of the synaptic inputs. *Progress in Neurobiology*, 39(6):563–607, December
1051 1992. ISSN 0301-0082. doi: 10.1016/0301-0082(92)90015-7.
- 1052 Taekyung Kwon, Angel Merchán-Pérez, Emiliano M Rial Verde, José-Rodrigo Rodríguez, Javier DeFelipe, and Rafael Yuste. Ultrastructural, Molecular and Functional Mapping of GABAergic Synapses
1053 on Dendritic Spines and Shafts of Neocortical Pyramidal Neurons. *Cerebral Cortex*, 29(7):2771–2781, July 2019. ISSN 1047-3211. doi: 10.1093/cercor/bhy143.
- 1054 Daniel Udvardy, Philipp Harth, Jakob H. Macke, Hans-Christian Hege, Christiaan P. J. de Kock, Bert Sakmann, and Marcel Oberlaender. The impact of neuron morphology on cortical network architecture.
1055 *Cell Reports*, 39(2):110677, April 2022. ISSN 2211-1247. doi: 10.1016/j.celrep.2022.110677.
- 1056 Yuriy Mishchenko, Tao Hu, Josef Spacek, John Mendenhall, Kristen M. Harris, and Dmitri B. Chklovskii. Ultrastructural analysis of hippocampal neuropil from the connectomics perspective. *Neuron*, 67
1057 (6):1009–1020, September 2010. ISSN 0896-6273. doi: 10.1016/j.neuron.2010.08.014.
- 1058 Narayanan Kasthuri, Kenneth Jeffrey Hayworth, Daniel Raimund Berger, Richard Lee Schalek, José Angel Conchello, Seymour Knowles-Barley, Dongli Lee, Amelio Vázquez-Reina, Verena Kaynig,
1059 Thouis Raymond Jones, Mike Roberts, Josh Lyskowski Morgan, Juan Carlos Tapia, H. Sebastian Seung, William Gray Roncal, Joshua Tzvi Vogelstein, Randal Burns, Daniel Lewis Sussman,
1060 Carey Eldin Priebe, Hanspeter Pfister, and Jeff William Lichtman. Saturated Reconstruction of a Volume of Neocortex. *Cell*, 162(3):648–661, July 2015. ISSN 0092-8674. doi: 10.1016/j.cell.2015.06.
1061 054.
- 1062 Solange P. Brown and Shaul Hestrin. Intracortical circuits of pyramidal neurons reflect their long-range axonal targets. *Nature*, 457(7233):1133–1136, February 2009. ISSN 1476-4687. doi: 10.1038/
1063 nature07658. Number: 7233 Publisher: Nature Publishing Group.
- 1064 Nuno Maçarico da Costa and Kevan A. C. Martin. How Thalamus Connects to Spiny Stellate Cells in the Cat’s Visual Cortex. *Journal of Neuroscience*, 31(8):2925–2937, February 2011. ISSN 0270-6474,
1065 1529-2401. doi: 10.1523/JNEUROSCI.5961-10.2011. Publisher: Society for Neuroscience Section: Articles.
- 1066 Xiaolong Jiang, Shan Shen, Cathryn R. Cadwell, Philipp Berens, Fabian Sinz, Alexander S. Ecker, Saamil Patel, and Andreas S. Tolias. Principles of connectivity among morphologically defined cell
1067 types in adult neocortex. *Science (New York, N.Y.)*, 350(6264):aac9462, November 2015. ISSN 1095-9203. doi: 10.1126/science.aac9462.
- 1068 Soohyun Lee, Illya Kruglikov, Z. Josh Huang, Gord Fishell, and Bernardo Rudy. A disinhibitory circuit mediates motor integration in the somatosensory cortex. *Nature Neuroscience*, 16(11):1662–1670,
1069 November 2013. ISSN 1546-1726. doi: 10.1038/nn.3544. Number: 11 Publisher: Nature Publishing Group.
- 1070 Tamás F. Freund and István Katona. Perisomatic Inhibition. *Neuron*, 56(1):33–42, October 2007. ISSN 0896-6273. doi: 10.1016/j.neuron.2007.09.012.
- 1071 Carsten K. Pfeffer, Mingshan Xue, Miao He, Z. Josh Huang, and Massimo Scanziani. Inhibition of inhibition in visual cortex: the logic of connections between molecularly distinct interneurons. *Nature Neuroscience*, 16(8):1068–1076, August 2013. ISSN 1546-1726. doi: 10.1038/nn.3446. Number: 8 Publisher: Nature Publishing Group.
- 1072 Yang-Sun Hwang, Catherine MacLachlan, Jérôme Blanc, Anaëlle Dubois, Carl C. H. Petersen, Graham Knott, and Seung-Hee Lee. 3D Ultrastructure of Synaptic Inputs to Distinct GABAergic Neurons in

- 1076 the Mouse Primary Visual Cortex. *Cerebral Cortex (New York, N.Y.: 1991)*, 31(5):2610–2624, March 2021. ISSN 1460-2199. doi: 10.1093/cercor/bhaa378.
- 1077 M Megias, Zs Emri, T. F Freund, and A. I Gulyás. Total number and distribution of inhibitory and excitatory synapses on hippocampal CA1 pyramidal cells. *Neuroscience*, 102(3):527–540, February 2001. ISSN 0306-4522. doi: 10.1016/S0306-4522(00)00496-6.
- 1078 J. G. Parnavelas, K. Sullivan, A. R. Lieberman, and K. E. Webster. Neurons and their synaptic organization in the visual cortex of the rat. *Cell and Tissue Research*, 183(4):499–517, October 1977. ISSN 1081-1432-0878. doi: 10.1007/BF00225663.
- 1080 Luke Campagnola, Stephanie C. Seeman, Thomas Chartrand, Lisa Kim, Alex Hoggarth, Clare Gamlin, Shinya Ito, Jessica Trinh, Pasha Davoudian, Cristina Radaelli, Mean-Hwan Kim, Travis Hage, Thomas Braun, Lauren Alfiler, Julia Andrade, Phillip Bohn, Rachel Dalley, Alex Henry, Sara Kebede, Alice Mukora, David Sandman, Grace Williams, Rachael Larsen, Corinne Teeter, Tanya L. Daigle, Kyla Berry, Nadia Dotson, Rachel Enstrom, Melissa Gorham, Madie Hupp, Samuel Dingman Lee, Kiet Ngo, Philip R. Nicovich, Lydia Potekhina, Shea Ransford, Amanda Gary, Jeff Goldy, Delissa McMillen, Tranthanh Pham, Michael Tieu, La'Akea Siverts, Miranda Walker, Colin Farrell, Martin Schroedter, Cliff Slaughterbeck, Charles Cobb, Richard Ellenbogen, Ryder P. Gwinn, C. Dirk Keene, Andrew L. Ko, Jeffrey G. Ojemann, Daniel L. Silbergeld, Daniel Carey, Tamara Casper, Kirsten Crichton, Michael Clark, Nick Dee, Lauren Ellingwood, Jessica Gloe, Matthew Kroll, Josef Sulc, Herman Tung, Katherine Wadhvani, Krissy Brouner, Tom Egdorf, Michelle Maxwell, Medea McGraw, Christina Alice Pom, Augustin Ruiz, Jasmine Bomben, David Feng, Nika Hejazinia, Shu Shi, Aaron Szafer, Wayne Wakeman, John Phillips, Amy Bernard, Luke Esposito, Florence D. D'Orazi, Susan Sunkin, Kimberly Smith, Bosiljka Tasic, Anton Arkhipov, Staci Sorensen, Ed Lein, Christof Koch, Gabe Murphy, Hongkui Zeng, and Tim Jarsky. Local connectivity and synaptic dynamics in mouse and human neocortex. *Science*, 375(6585):eabj5861, March 2022. doi: 10.1126/science.abj5861. Publisher: American Association for the Advancement of Science.
- 1089 Davi D. Bock, Wei-Chung Allen Lee, Aaron M. Kerlin, Mark L. Andermann, Greg Hood, Arthur W. Wetzel, Sergey Yurgenson, Edward R. Soucy, Hyon Suk Kim, and R. Clay Reid. Network anatomy and in vivo physiology of visual cortical neurons. *Nature*, 471(7337):177–182, March 2011. ISSN 1476-4687. doi: 10.1038/nature09802. Number: 7337 Publisher: Nature Publishing Group.
- 1092 Rita Bopp, Nuno Maçarico da Costa, Björn M. Kampa, Kevan A. C. Martin, and Morgane M. Roth. Pyramidal Cells Make Specific Connections onto Smooth (GABAergic) Neurons in Mouse Visual Cortex. *PLoS Biology*, 12(8):e1001932, August 2014. ISSN 1545-7885. doi: 10.1371/journal.pbio.1001932. Publisher: Public Library of Science.
- 1093 Helene Schmidt, Anjali Gour, Jakob Straehle, Kevin M. Boergens, Michael Brecht, and Moritz Helmstaedter. Axonal synapse sorting in medial entorhinal cortex. *Nature*, 549(7673):469–475, September 2017. ISSN 1476-4687. doi: 10.1038/nature24005.
- 1095 Sen Song, Per Jesper Sjöström, Markus Reigl, Sacha Nelson, and Dmitri B. Chklovskii. Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits. *PLOS Biology*, 3(3):e68, March 2005. ISSN 1545-7885. doi: 10.1371/journal.pbio.0030068. Publisher: Public Library of Science.
- 1097 Rodrigo Perin, Thomas K. Berger, and Henry Markram. A synaptic organizing principle for cortical neuronal groups. *Proceedings of the National Academy of Sciences*, 108(13):5419–5424, March 2011. doi: 10.1073/pnas.1016051108. Publisher: Proceedings of the National Academy of Sciences.
- 1100 R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon. Network Motifs: Simple Building Blocks of Complex Networks. *Science*, 298(5594):824–827, October 2002. doi: 10.1126/science.298.5594.824. Publisher: American Association for the Advancement of Science.
- 1101 Yasuo Kawaguchi, Fuyuki Karube, and Yoshiyuki Kubota. Dendritic branch typing and spine expression patterns in cortical nonpyramidal cells. *Cerebral Cortex (New York, N.Y.: 1991)*, 16(5):696–711, May 2006. ISSN 1047-3211. doi: 10.1093/cercor/bhj015.
- 1104 M. E. Chicurel and K. M. Harris. Three-dimensional analysis of the structure and composition of CA3 branched dendritic spines and their synaptic relationships with mossy fiber boutons in the rat hippocampus. *The Journal of Comparative Neurology*, 325(2):169–182, November 1992. ISSN 0021-9967. doi: 10.1002/cne.903250204.
- 1106 Eric Wang and Andreas Tolias. Foundation model of the cortex, March 2023.
- 1107 Jonathan Zung, Ignacio Tartavull, Kisuk Lee, and H. Sebastian Seung. An Error Detection and Correction Framework for Connectomics. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- 1109 Felix Gonda, Xueying Wang, Johanna Beyer, Markus Hadwiger, Jeff W. Lichtman, and Hanspeter Pfister. VICE: Visual Identification and Correction of Neural Circuit Errors, May 2021. arXiv:2105.06861 [cs].
- 1110 David Rolnick, Yaron Meirovitch, Toufiq Parag, Hanspeter Pfister, Viren Jain, Jeff W. Lichtman, Edward S. Boyden, and Nir Shavit. Morphological Error Detection in 3D Segmentations, May 2017. arXiv:1705.10882 [cs, q-bio, stat].
- 1111 Khoa Tuan Nguyen, Ganhhee Jang, Tran Anh Tuan, and Won-ki Jeong. RLCorrector: Reinforced Proofreading for Cell-level Microscopy Image Segmentation, March 2022. arXiv:2106.05487 [cs].
- 1114 Philipp J. Schubert, Sven Dorkenwald, Michal Januszewski, Viren Jain, and Joergen Kornfeld. Learning cellular morphology with neural networks. *Nature Communications*, 10(1):2736, June 2019. ISSN 2041-1723. doi: 10.1038/s41467-019-10836-3. Number: 1 Publisher: Nature Publishing Group.
- 1115 Martin Schmidt, Alessandro Motta, Meike Sievers, and Moritz Helmstaedter. RoboEM: automated 3D flight tracing for synaptic-resolution connectomics, September 2022. Pages: 2022.09.08.507122 Section: New Results.
- 1118 Jules Berman, Dmitri B. Chklovskii, and Jingpeng Wu. Bridging the Gap: Point Clouds for Merging Neurons in Connectomics. In Ender Konukoglu, Bjoern Menze, Archana Venkataraman, Christian Baumgartner, Qi Dou, and Shadi Albarqouni, editors, *Proceedings of The 5th International Conference on Medical Imaging with Deep Learning*, volume 172 of *Proceedings of Machine Learning Research*, pages 150–159. PMLR, July 2022.
- 1121 Yaron Meirovitch, Alexander Matveev, Hayk Saribekyan, David Budden, David Rolnick, Gergely Odor, Seymour Knowles-Barley, Thouis Raymond Jones, Hanspeter Pfister, Jeff William Lichtman, and Nir Shavit. A Multi-Pass Approach to Large-Scale Connectomics, December 2016. arXiv:1612.02120 [cs, q-bio].
- 1122 Ronell B. Sicat, Markus Hadwiger, and Niloy J. Mitra. Graph Abstraction for Simplified Proofreading of Slice-based Volume Segmentation. 2013. ISSN 1017-4656. doi: 10.2312/conf/EG2013/short/077-080. Accepted: 2014-01-26T14:58:25Z Publisher: The Eurographics Association.
- 1124 Brian Matejek, Daniel Haehn, Haidong Zhu, Donglai Wei, Toufiq Parag, and Hanspeter Pfister. Biologically-Constrained Graphs for Global Connectomics Reconstruction. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2084–2093, June 2019. doi: 10.1109/CVPR.2019.00219. ISSN: 2575-7075.
- 1127 Petra Wahle, Eric Sobierajski, Ina Gasterstädt, Nadja Lehmann, Susanna Weber, Joachim HR Lübke, Maren Engelhardt, Claudia Distler, and Gundela Meyer. Neocortical pyramidal neurons with axons emerging from dendrites are frequent in non-primates, but rare in monkey and human. *eLife*, 11:e76101, April 2022. ISSN 2050-084X. doi: 10.7554/eLife.76101. Publisher: eLife Sciences Publications, Ltd.
- 1129 Katherine L. Villa and Ely Nedivi. Excitatory and Inhibitory Synaptic Placement and Functional Implications. In Kazuo Emoto, Rachel Wong, Eric Huang, and Casper Hoogenraad, editors, *Dendrites*, pages 467–487. Springer Japan, Tokyo, 2016. ISBN 978-4-431-56048-7 978-4-431-56050-0. doi: 10.1007/978-4-431-56050-0_18.
- 1132 Dimitri Yatsenko, Jacob Reimer, Alexander S. Ecker, Edgar Y. Walker, Fabian Sinz, Philipp Berens, Andreas Hoenselaer, R. James Cotton, Athanasios S. Siapas, and Andreas S. Tolias. DataJoint: managing big scientific data using MATLAB or Python, November 2015. Pages: 031658 Section: New Results.
- 1133 Dimitri Yatsenko, Edgar Y. Walker, and Andreas S. Tolias. DataJoint: A Simpler Relational Data Model. *arXiv:1807.11104 [cs]*, July 2018. arXiv: 1807.11104.
- 1134 Paolo Cignoni, Massimiliano Corsini, and Guido Ranzuglia. MeshLab: an Open-Source 3D Mesh Processing System. *ERCIM News*, 2008(73), 2008.
- 1136 Ilker O. Yaz and Sébastien Lorient. Triangulated Surface Mesh Segmentation. In *CGAL User and Reference Manual*. CGAL Editorial Board, 5.5.2 edition, 2023.
- 1137 Eric Perlman. Visualizing and Interacting with Large Imaging Data. *Microscopy and Microanalysis*, 25(S2):1374–1375, August 2019. ISSN 1431-9276, 1435-8115. doi: 10.1017/S1431927619007608. Publisher: Cambridge University Press.
- 1138 Matthias Fey and Jan Eric Lenssen. Fast Graph Representation Learning with PyTorch Geometric, May 2019.

1140 **Supplementary Figures**

a MICrONS Mesh Processing Validation					b H01 Mesh Processing Validation				
category	precision	recall	f1	# datapoints	category	precision	recall	f1	# datapoints
soma	1.000	1.000	1.000	195	soma	0.990	0.990	0.990	102
axon	0.995	0.968	0.981	192	glia	1.000	1.000	1.000	100
glia	1.000	1.000	1.000	188	axon	0.942	0.875	0.907	76
spine	0.896	0.870	0.883	22226	spine	0.831	0.765	0.796	7945
spine (> 0.7 μ m)	0.942	0.978	0.960	18030	spine (> 0.7 μ m)	0.892	0.953	0.922	5851

Fig. 8. Mesh Processing Pipeline Validation **a)** (MICrONS) Validation scores of automatic submesh (compartment) identification in comparison to human labels. This data set was produced by randomly presenting processed mesh segments to human annotators who evaluated automatic labels as true positive, true negative, false positive, or false negative (TP,TN,FP,or FN) for each structure. Only glia merges larger than the volume of a 5 μ m sphere were considered glia merges in this processing step. The “spine (> 0.7 μ m)” row reports the agreement between the automatic spine detection and human spine labeling for spines with a skeletal length greater than 0.7 μ m. Below this threshold there was disagreement even among human proofreaders about whether small protrusions should be classified as a spine or not. **b)** Identical validation scores for H01 dataset.

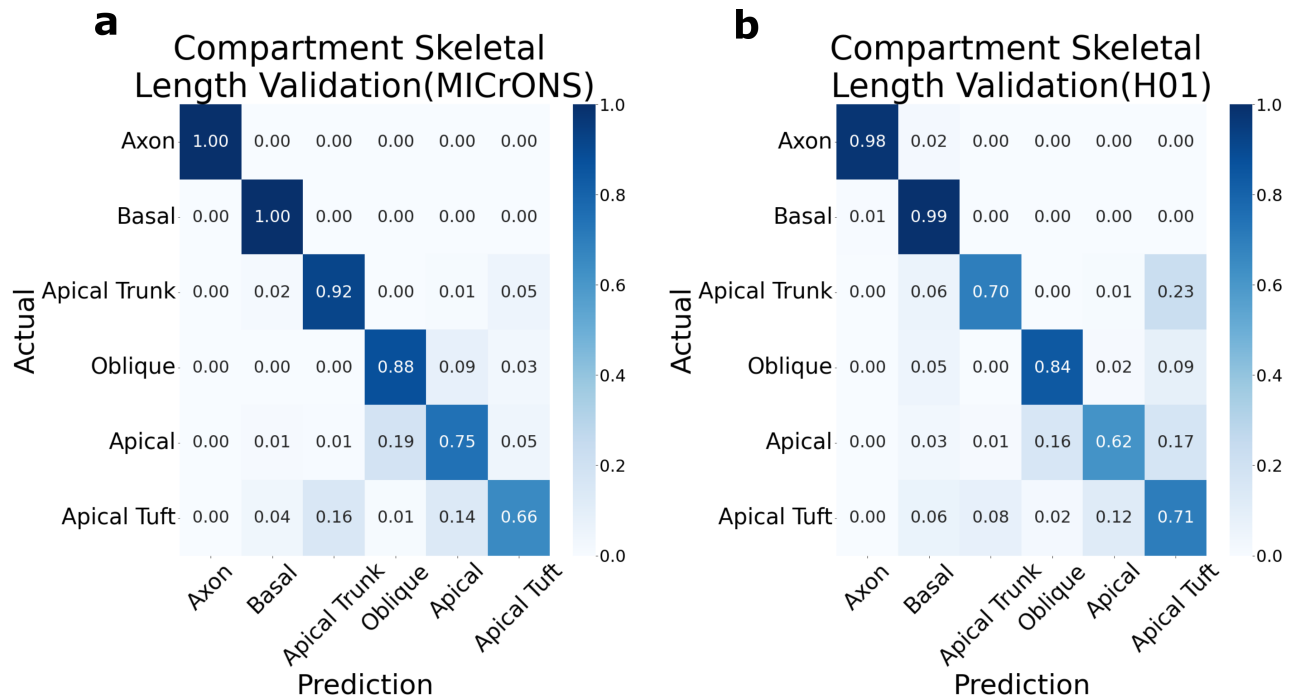


Fig. 9. Automatic Submesh (Compartment) Labeling Validation **a)** (MICrONS) Confusion matrix comparing automatic submesh labeling to human labels (random sampling of 158 processed cells). Here the TP,FP,TN,FN metrics are computed using skeletal length agreement for each compartment. Therefore, cells with longer stems or cells with more stems of a certain compartment type more heavily influenced the scores due to the skeletal length weighting. **b)** (H01) Confusion Matrix comparing automatic submesh labeling to human labeled submeshes (random sampling of 89 processed cells). Across both datasets, compartment labeling was nearly perfect for axon, basal, and apical compartments, but was less consistent for sub-compartments of the apical stem.

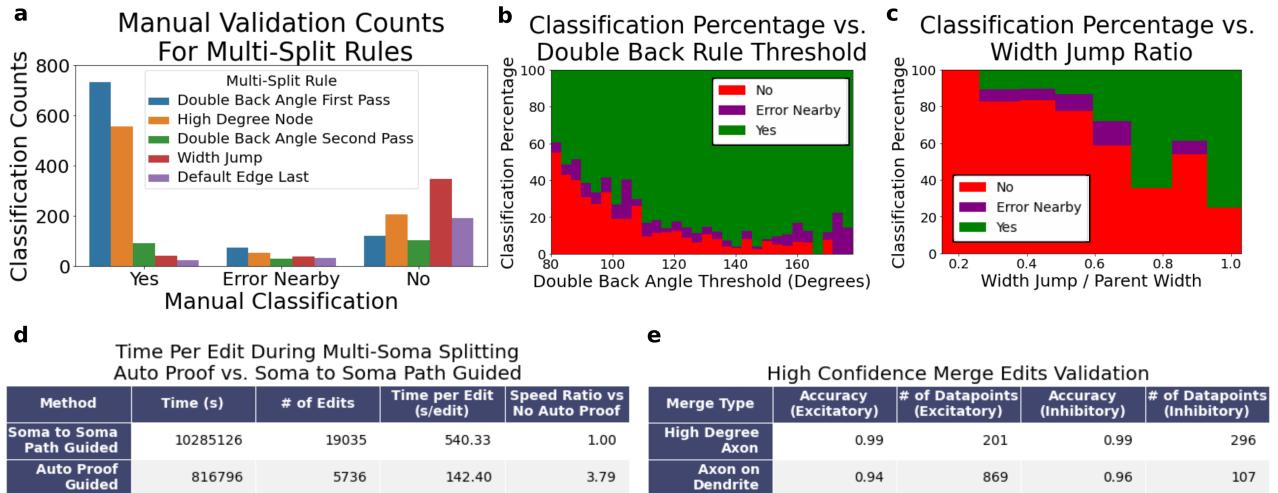


Fig. 10. Proofreading Validation All validation was performed by the proofreading team at Johns Hopkins University Applied Physics Laboratory (APL). In an initial round of validation, suggested error locations were evaluated in the context of splitting multi-soma cells in the MICrONS volume. As a result we were able to measure both the accuracy of these proofreading rules and the speed benefits of a semi-supervised approach compared to fully-manual proofreading. Additionally, the accuracy of a two automatic proofreading rules with high-confidence parameters (axon on dendrite, high degree axon on axon merges) were evaluated. **a**) Validation of split locations predicted by automatic multi-split algorithm. "Yes" (indicates that the proposed split can be executed immediately), "Error Nearby" indicates that the split location is correct within 20 μm , but that the human proofreader slightly modified the suggested split points, and "No" indicates that the true split location was far from the predicted location or no merge error was detected by the human proofreader.). The heuristic splitting rules are applied in the order indicated by the legend. The automated proofreading accuracy varied substantially over the different heuristic rules with an overall accuracy of 76.12 % when yes and Error Nearby are considered true positives. The best-performing rules can be selected for different datasets. **b**) Even for a single rule, thresholds can be tuned to optimize performance. Manual Classification of split locations predicted by the "Double Back" rule as a function of the angle measured at each predicted location illustrates that a higher accuracy could be achieved by setting a higher threshold for this algorithm. **c**) Manual Classification of split locations predicted by the width jump rule as a function of the width jump at each predicted location illustrating another example where interpretable thresholds can be adjusted for higher precision. **d**) Time statistics collected as humans performed manual tasks of splitting multi-soma neurons either using a tool that showed the path along the mesh between two somas or using the suggested split locations from our automatic multi-split algorithm. The speed at which humans could apply cuts in the correct locations more than tripled when using suggestions provided by the NEURD multi-split algorithm. Note: The validation is measured as the average amount of time for a single edit in the multi-soma splitting process; a single multi-soma split might require 20 or more edits to completely resolve the merge. **e**) Accuracy of two automatic proofreading rules with high-confidence parameters (high degree axon on axon, axon on dendrite).

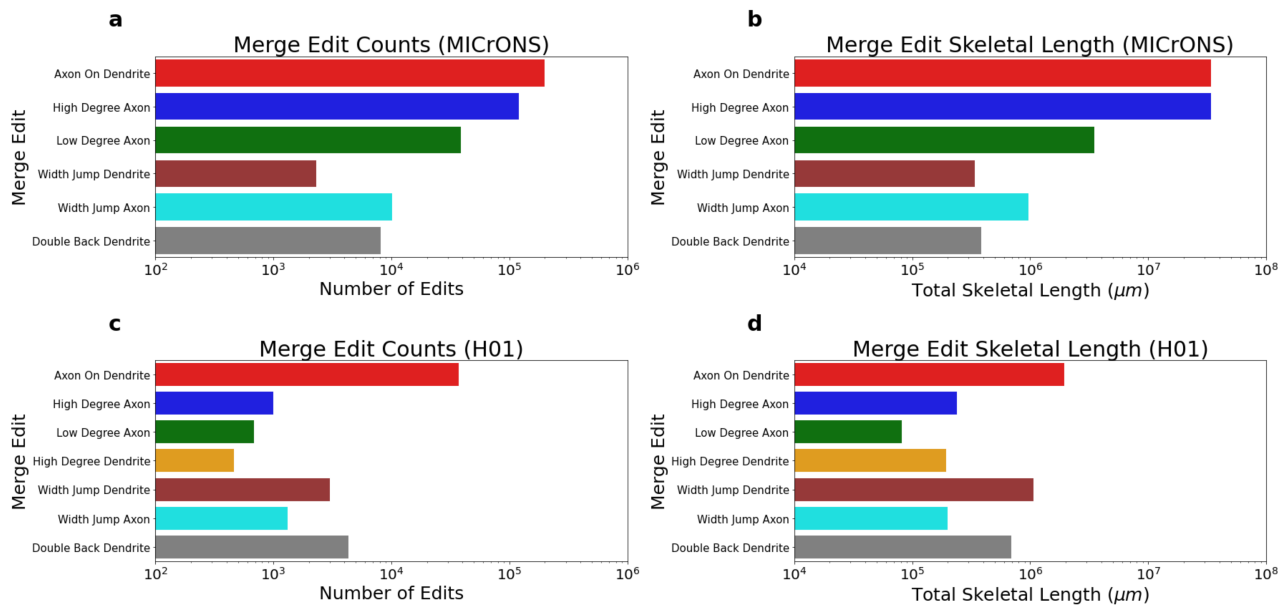


Fig. 11. Counts and total skeletal length of merge errors corrected during automated proofreading. The different heuristic rules are presented from top to bottom in the order that they are implemented in the automated proofreading workflow. Note that errors identified by rules later in the workflow are excluded from these statistics if they are found on already-errored segments identified earlier in the workflow. **a)** (MICrONS) Total number of separate locations where a specific heuristic rule corrected a merge error. **b)** (MICrONS) Total skeletal length eliminated by each heuristic rule. **c)** (H01) Total number of separate locations where a specific heuristic rule corrected a merge error. **d)** (H01) Total skeletal length eliminated by each heuristic rule.

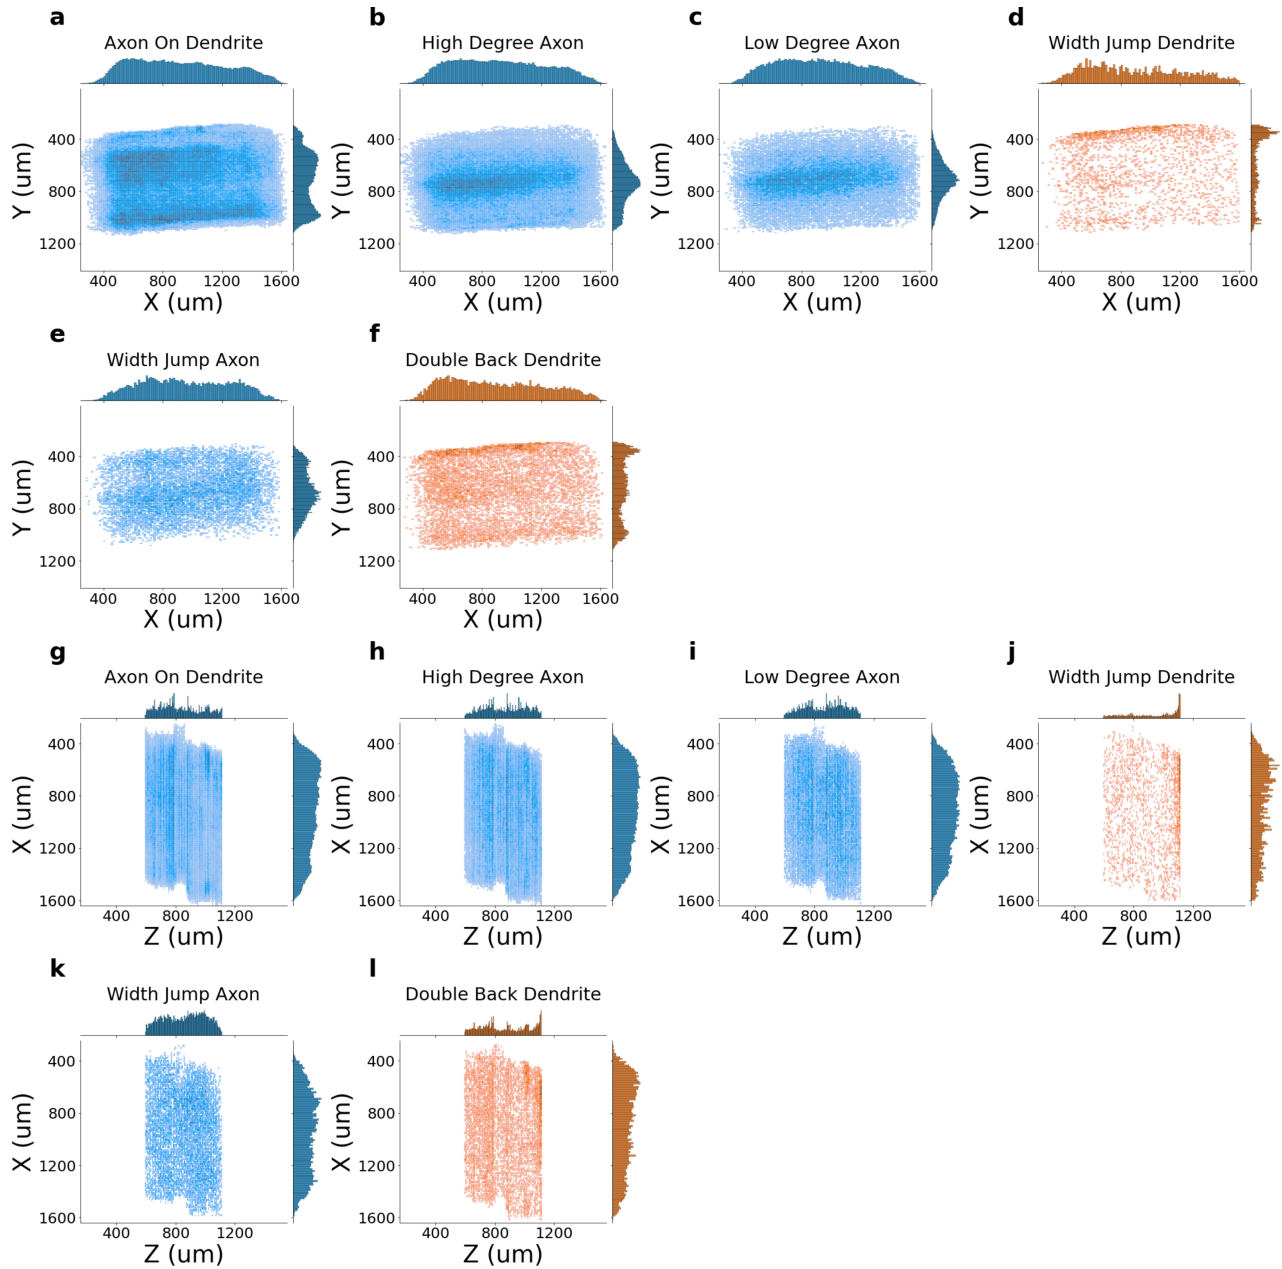


Fig. 12. Spatial Distribution of MICrONS Merge Edit Locations. The distribution of locations show biases in the volume for certain types of merge edits. These spatial biases may be due to segmentation or slicing defects, or differences in the concentration of different kinds of neuropil throughout the volume. **a - f)** X,Y merge edit locations for different heuristic rules **g - l)** X,Z merge edit locations for different heuristic rules

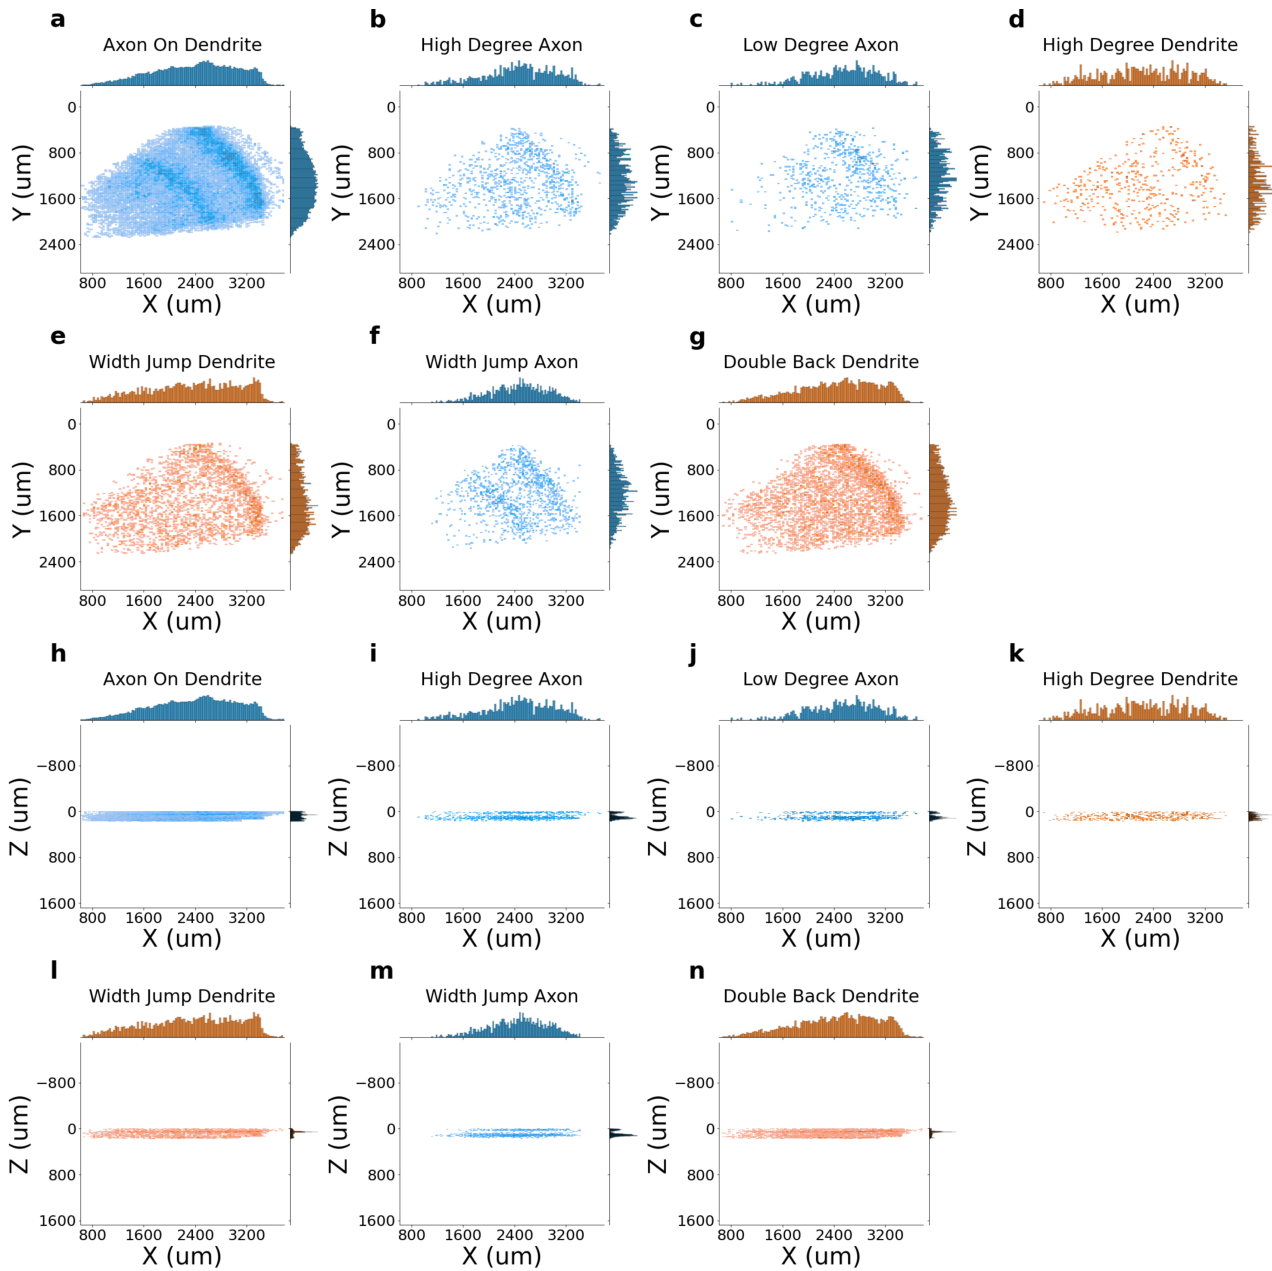


Fig. 13. Spatial Distribution of H01 Merge Edit Locations. The distribution of locations show biases in the volume for certain types of merge edits. These spatial biases may be due to segmentation or slicing defects, or differences in the concentration of different kinds of neuropil throughout the volume. **a - g)** X,Y merge edit locations for different heuristic rules **h - n)** X,Z merge edit locations for different heuristic rules

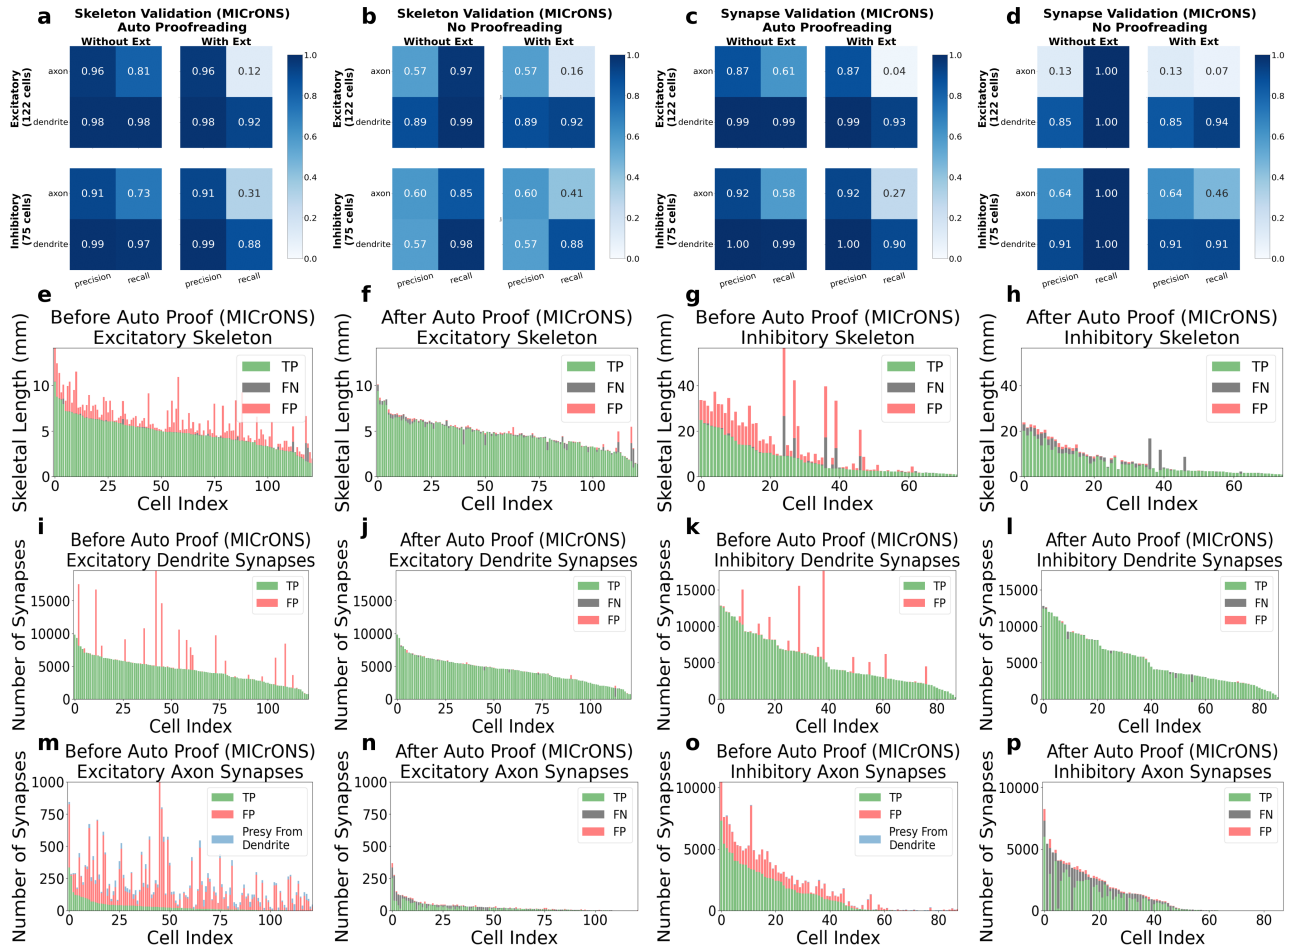


Fig. 14. Supplemental MICrONS Auto Proofreading Validation. Validation metrics and visualizations of the automatic proofreading step when comparing the edits made by the automatic proofreading algorithm to edits made by human proofreaders. Metrics reported for “With Ext” refer to the skeleton or synapse validation compared to a human proofreader who both cleaned the existing cell of merge errors and added back missing axon and dendrite segments, whereas “Without Ext” refers to a comparison of only cleaning away merge errors. The number of cells in the test set were 122 excitatory and 75 inhibitory. Metrics reported for “No Proofreading” refer to how the raw cells without any proofreading would compare to after human proofreading, thus giving a sense of error rates of the raw segmentation. Histograms (panels e - p) give a visual representation of the metrics reported in the precision/recall tables (panels a - d). FN classifications can exist before automatic proofreading because of dropping axon/dendritic segments in the mesh and graph processing pipeline prior to the automatic proofreading step. Note: neurons with multi-soma merges are included in these visualizations and metrics. **a)** The precision/recall metrics comparing the skeleton length of cells after automatic proofreading for the Axon/Dendrite compartments and for different exc/inh cell types when compared to human proofreading with extension and without. **b)** The precision/recall metrics comparing the skeleton length of cells with no automatic proofreading. **c)** The precision/recall metrics comparing the synapse counts of cells with automatic proofreading. **d)** The precision/recall metrics comparing the synapse counts of cells with no proofreading. **e - h)** TP/FN/FP classification of each test cell’s skeletons before and after automatic proofreading for both excitatory and inhibitory cells, demonstrating a large percentage of the FP skeleton segments are removed after the process. **i - l)** TP/FN/FP classification of each test cell’s dendrite synapses (postsyns) before and after automatic proofreading for both excitatory and inhibitory cells, demonstrating a large percentage of the FP postsyns are removed after the application of dendrite proofreading heuristics. **m - p)** TP/FN/FP classification of each test cell’s axon synapses (presyn) before and after automatic proofreading for both excitatory and inhibitory cells, demonstrating a large percentage of the FP presyns are removed after the application of axon proofreading heuristics. Those axon presyns located not on the main axon but on dendritic segments are filtered away and designated as “Presyn From Dendrite”, which does not include the heuristic rule of “Axon on Dendrite” but instead just filters away any presyns located on dendritic segments that were not filtered away using the heuristic rules.

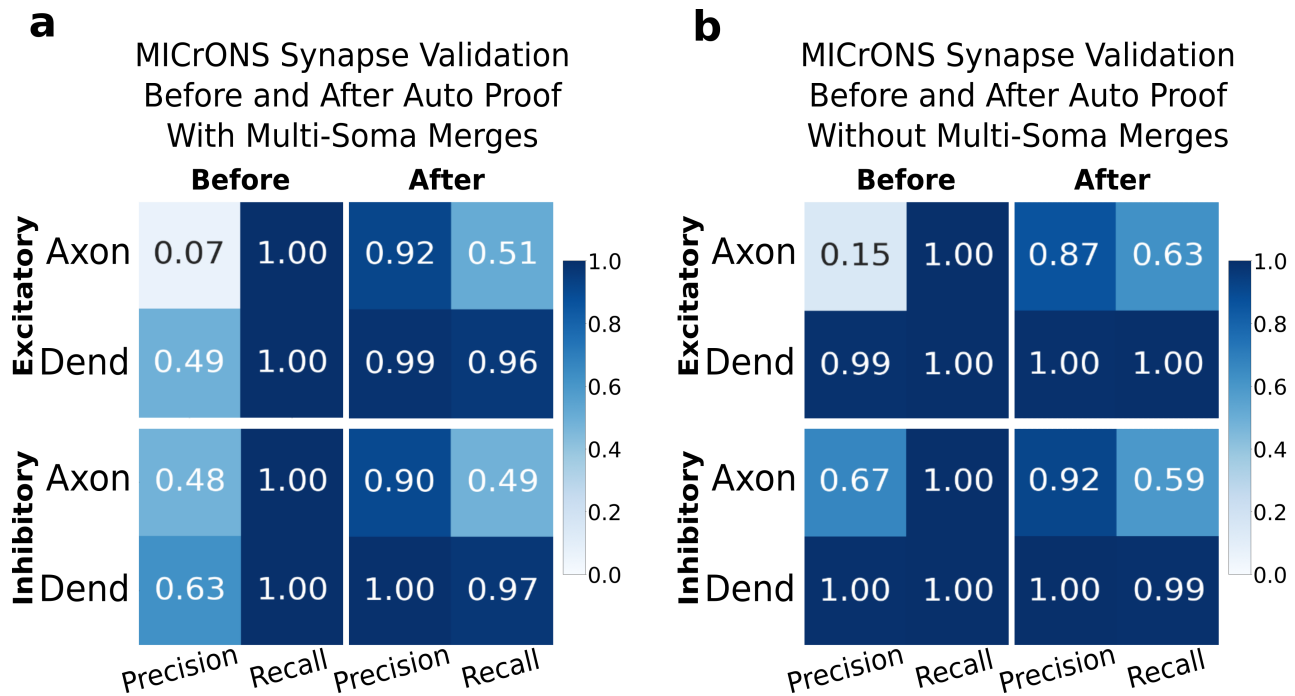


Fig. 15. Supplemental MICrONS Multi-Soma Auto Proofreading Validation. The precision/recall metrics comparing the synapse counts of cells before and after automatic proofreading for the Axon/Dendrite compartments and for different excitatory/inhibitory cell types when compared to human proofreading. **a)** Validation when only considering neurons with at least one soma merge to the main segment (19 excitatory, 12 inhibitory). **b)** Validation when only considering neurons with no soma merged to the main segment (103 excitatory, 76 inhibitory).

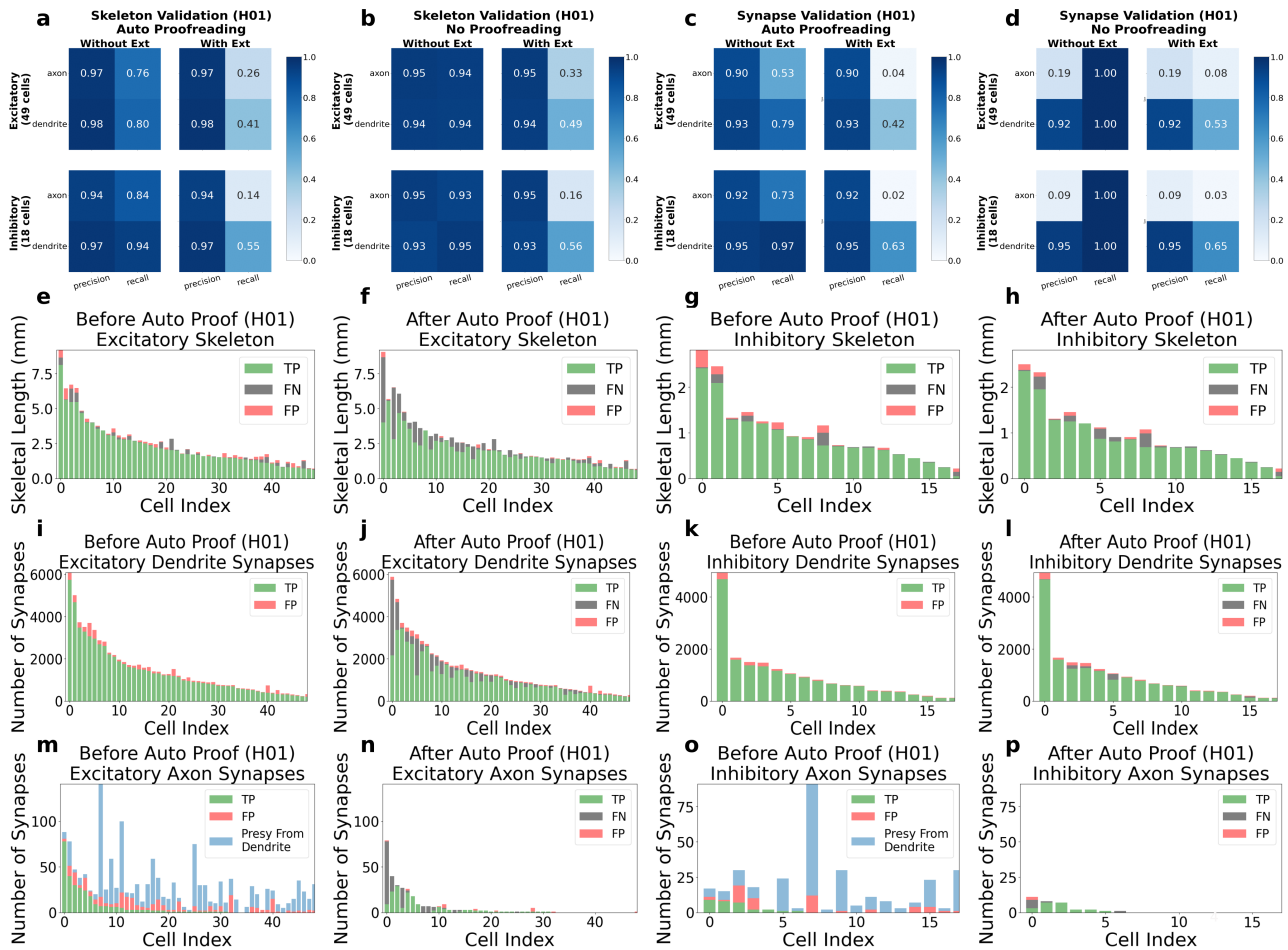


Fig. 16. Supplemental H01 Auto Proofreading Validation. Validation metrics and visualizations of the automatic proofreading step when comparing the edits made by the automatic proofreading algorithm to edits made by human proofreaders. Metrics reported for “With Ext” refer to the skeleton or synapse validation compared to a human proofreader who both cleaned the existing cell of merge errors and added back missing axon and dendrite segments, whereas “Without Ext” refers to a comparison of only cleaning away merge errors. The number of cells in the test set were 49 excitatory and 18 inhibitory. Metrics reported for “No Proofreading” refer to how the raw cells without any proofreading would compare to after human proofreading, thus giving a sense of error rates of the raw segmentation. Histograms (panels e - p) give a visual representation of the metrics reported in the precision/recall tables (panels a - d). FN classifications can exist before automatic proofreading because of dropping axon/dendritic segments in the mesh and graph processing pipeline prior to the automatic proofreading step. Note: While perfectly extending all axonal and dendritic processes is not yet possible, the extent to which neurons were extended in the manually proofread set from the H01 dataset are much less extensively extended in comparison to those of the MICrONS dataset; therefore, the recall numbers for the “With Ext” categories in the H01 validation are much more likely an over-estimate in comparison with those of the MICrONS dataset. **a)** The precision/recall metrics comparing the skeleton length of cells after automatic proofreading for the Axon/Dendrite compartments and for different exc/inh cell types when compared to human proofreading with extension and without. **b)** The precision/recall metrics comparing the skeleton length of cells with no automatic proofreading. **c)** The precision/recall metrics comparing the synapse counts of cells with automatic proofreading. **d)** The precision/recall metrics comparing the synapse counts of cells with no proofreading. **e - h)** TP/FN/FP classification of each test cell’s skeletons before and after automatic proofreading for both excitatory and inhibitory cells, demonstrating a large percentage of the FP skeleton segments are removed after the process. **i - l)** TP/FN/FP classification of each test cell’s dendrite synapses (postsyns) before and after automatic proofreading for both excitatory and inhibitory cells, demonstrating a large percentage of the FP postsyns are removed after the application of dendrite proofreading heuristics. **m - p)** TP/FN/FP classification of each test cell’s axon synapses (presyn) before and after automatic proofreading for both excitatory and inhibitory cells, demonstrating a large percentage of the FP presyns are removed after the application of axon proofreading heuristics. Those axon presyns located not on the main axon but on dendritic segments are filtered away and designated as “Presyn From Dendrite”, which does not include the heuristic rule of “Axon on Dendrite” but instead just filters away any presyns located on dendritic segments that were not filtered away using the heuristic rules.

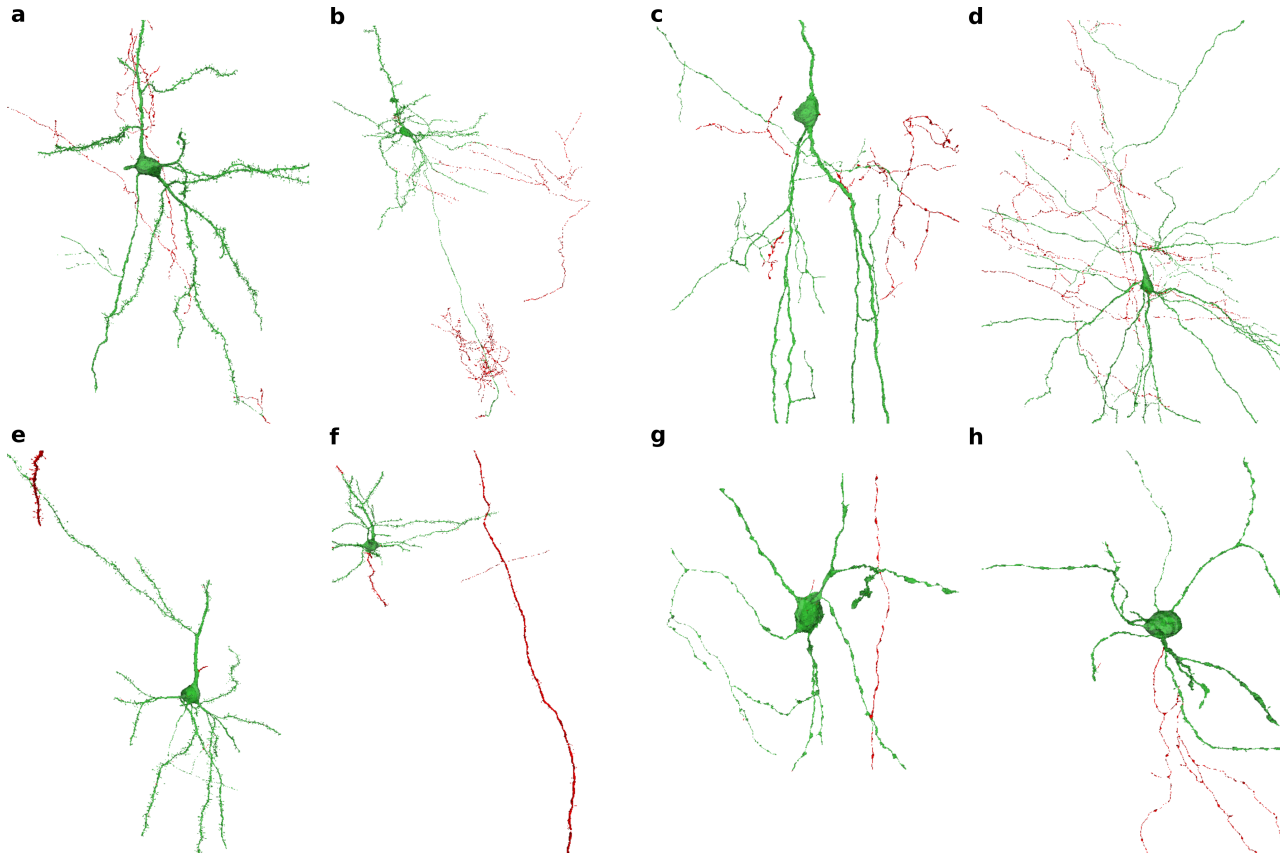


Fig. 17. Proofread Neuron Examples with Merge Errors Labeled Examples of excitatory and inhibitory neurons from both the MICrONS and H01 dataset after automatic proofreading (green) with the removed merge errors shown (red). **a,b**) (MICrONS) Example excitatory cells in the 50th and 90th percentile of merge error skeletal length removed. **c,d**) (MICrONS) Example inhibitory cells in the 50th and 90th percentile of merge error skeletal length removed. **e,f**) (H01) Example excitatory cells in the 50th and 90th percentile of merge error skeletal length removed. **g,h**) (H01) Example inhibitory cells in the 50th and 90th percentile of merge error skeletal length removed.

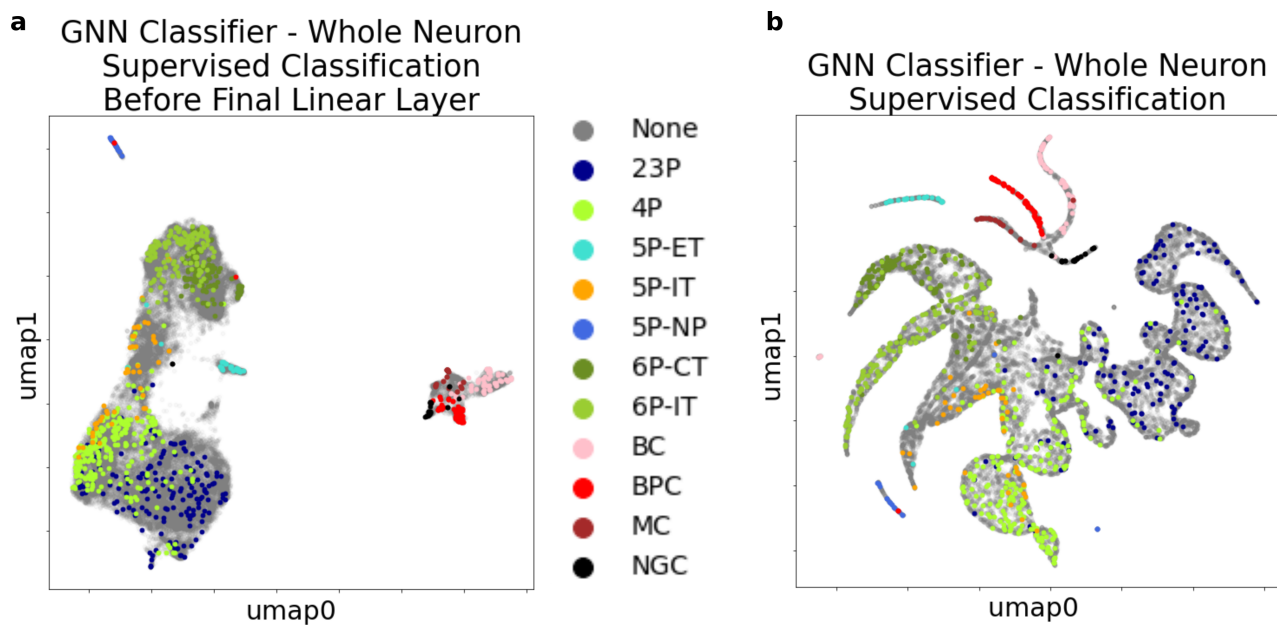


Fig. 18. GNN Classifier Whole Neuron UMAP Embeddings. **a)** Embeddings before the final linear layer and softmax function with hand labeled cells from (Schneider-Mizell et al., 2023) overlaid (these labels were used for the training and validation process of GNN). Cell-type separation is evident at this stage indicating that the classifier has learned useful features prior to the readout. **b)** Embeddings after the final linear layer with cell type labels from (Schneider-Mizell et al., 2023).

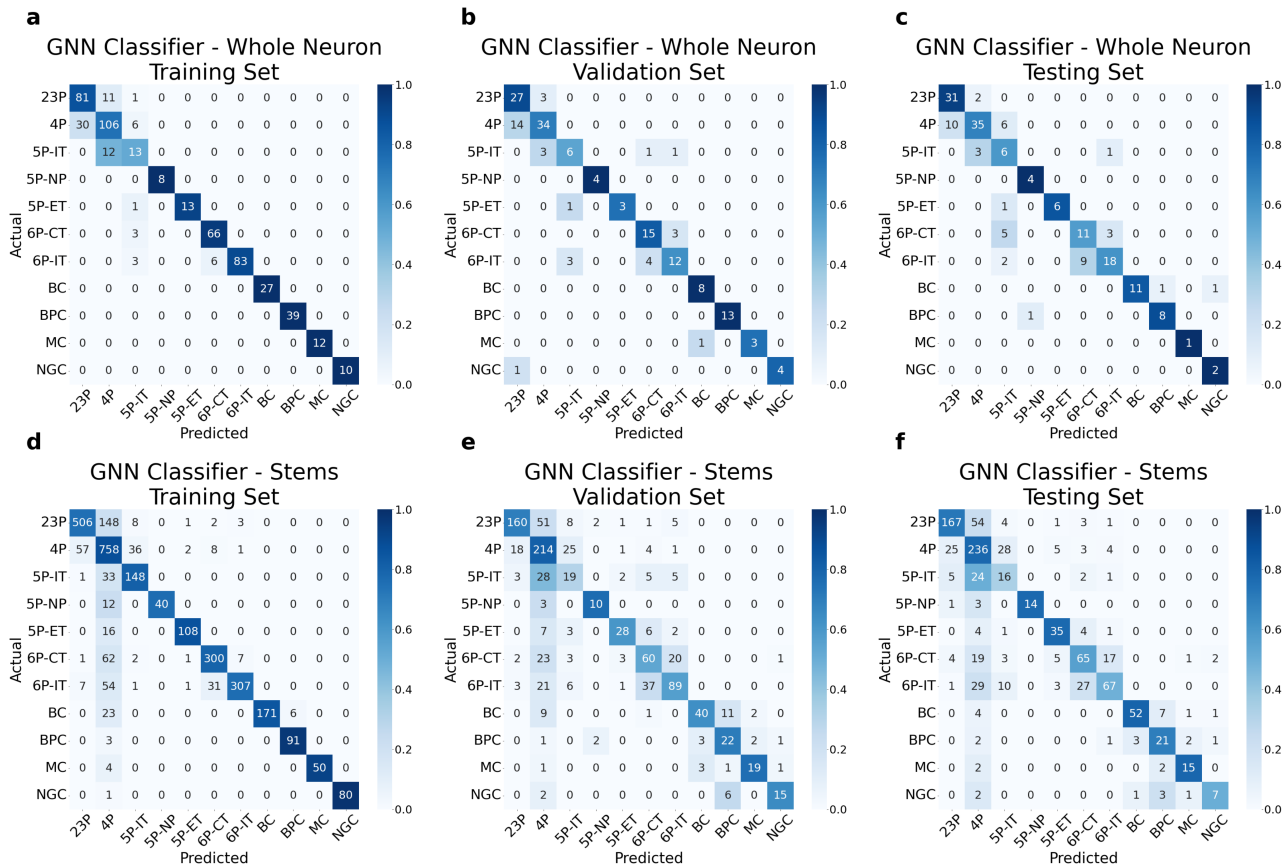


Fig. 19. GNN Classifier Train/Validation/Test Confusion. Confusion matrix of predicted and actual neuron counts from each of the 60:20:20 Train/Validation/Test splits used in the supervised training process of the Graph Neural Network cell type classifier. All ground truth labels were from hand-annotated classes described in (Schneider-Mizell et al., 2023). Color intensities are from normalized values in reference to the summation of a given row. **a-c)** Confusion matrix (neuron counts) for the GNN Classifier using the entire neuron graph including the somatic root node. **d-f)** Confusion matrix (stem counts) for the GNN Classifier applied to a single stem (a single projection from the soma), without any information about the soma itself.

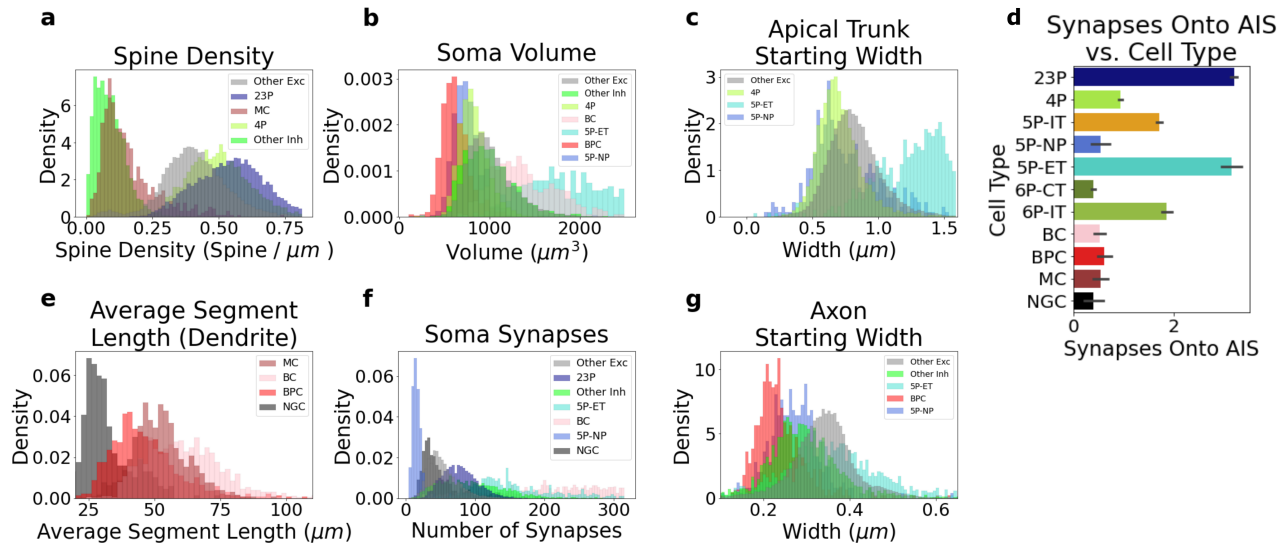


Fig. 20. Various morphological features computed by NEURD. Histograms and bar graphs of a variety of salient features computed by NEURD, colored by the labels generated from the GNN classifier. Some plots are replicating previous work from (Elabbady et al., 2022). **a**) Spine density (number of spines per μm of skeletal length) distributions from automatic spine detection. As has been previously reported (Azouz et al., 1997), layer 2/3 pyramidal cells are more densely spiny than layer 4, and MC spine density is higher than other inhibitory cells. **b**) Soma volume computed during the soma detection step in the NEURD mesh processing pipeline. As expected, 5P-NP, 4P and BPC generally have smaller somas than other cells from their same excitatory or inhibitory class while 5P-ET and BC are larger than other cells in the same class (Elabbady et al., 2022). **c**) Width measurements generated from the average distance of the inner skeleton to the mesh surface (radius approximation) at the beginning of the apical trunk protrusion. Compared to other cell types, 4P and 5P-NP cells have smaller trunks, while 5P-ET are larger. **d**) Average number of synapses onto the axon initial segment (AIS) for different cell types. As expected, 23P, 5P-ET, and 5P-IT cell types are more densely innervated on their AIS (Schneider-Mizell et al., 2021) AIS is defined as within $10 \mu\text{m}$ - $40 \mu\text{m}$ skeletal distance of the soma, and error bars are standard deviation. **e**) Average skeletal length of non-branching dendritic segments for stems of different cell type subclasses, illustrating that NGC have significantly shorter distances between branch points in their dendrites than other inhibitory cells. **f**) Distributions of synapses onto the soma illustrating the expected larger average number of soma synapses for 5P-ET and BC and smaller numbers for 5P-NP and NGC (Elabbady et al., 2022). **g**) Distributions of radius approximation for the start of the axon protrusion from either a dendrite or the soma, showing smaller typical widths for 5P-NP and BPC and larger starting widths for 5P-ET.

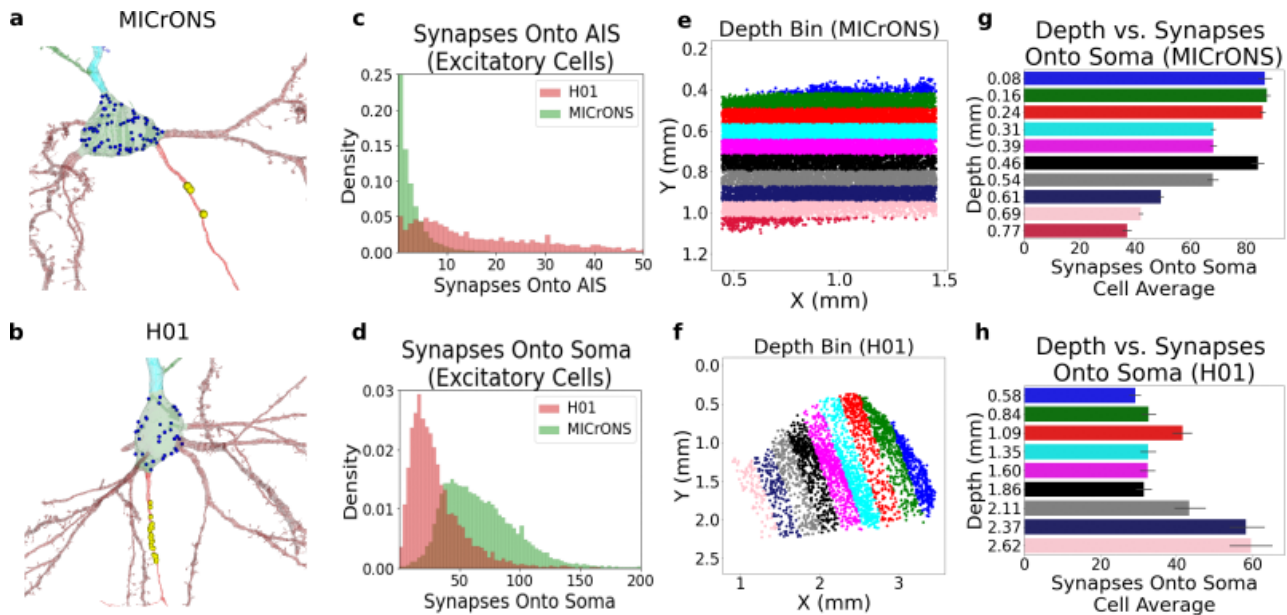


Fig. 21. Synapses onto AIS and Soma Comparison a-b) Example neuron with synapses onto the axon (AIS synapses) in yellow and synapses onto the soma in blue. Example neurons are both in the 75th percentile of the number of AIS synapses in their respective volumes. **c)** Distribution of the number of AIS synapses per cell compared across datasets, emphasizing the increased innervation of the AIS for neurons in the H01 dataset in comparison to MICrONS. **d)** Distribution of the number of soma synapses per cell. As expected, neurons in the MICrONS volume have more identified synapses onto their soma, despite the smaller surface area of mouse somas compared to human (Wildenberg et al., 2021). **e-f)** Depth bins used for analysis of both synapses onto AIS vs depth (Main Fig. 5f,g) and synapses onto soma vs. depth (this figure, panels g,h). **g-h)** Average number of synapses onto the soma of cells varies across depth (mean +/- std), decreasing in deeper layers of the MICrONS volume, but increasing in deeper layers of the H01 dataset.




Fig. 22. Neuron Dendritic Branching Characteristics. Measurements related to leaf nodes (terminating ends of the dendritic stem) excluding apical dendrites. **a)** Distributions of the number of total leaf nodes for the non-apical dendrites of each neuron separated by inhibitory cell type. As expected, NGC cells have the most leaf nodes of any inhibitory cell type, while BPC have fewer leaf nodes compared to other interneurons. **b-c)** Histogram for all the non-apical dendritic stems of every neuron in the volume comparing the initial width of the stem to the number of leaf nodes. For both the MICrONS (b) and H01 (c), there is a significant positive correlation.

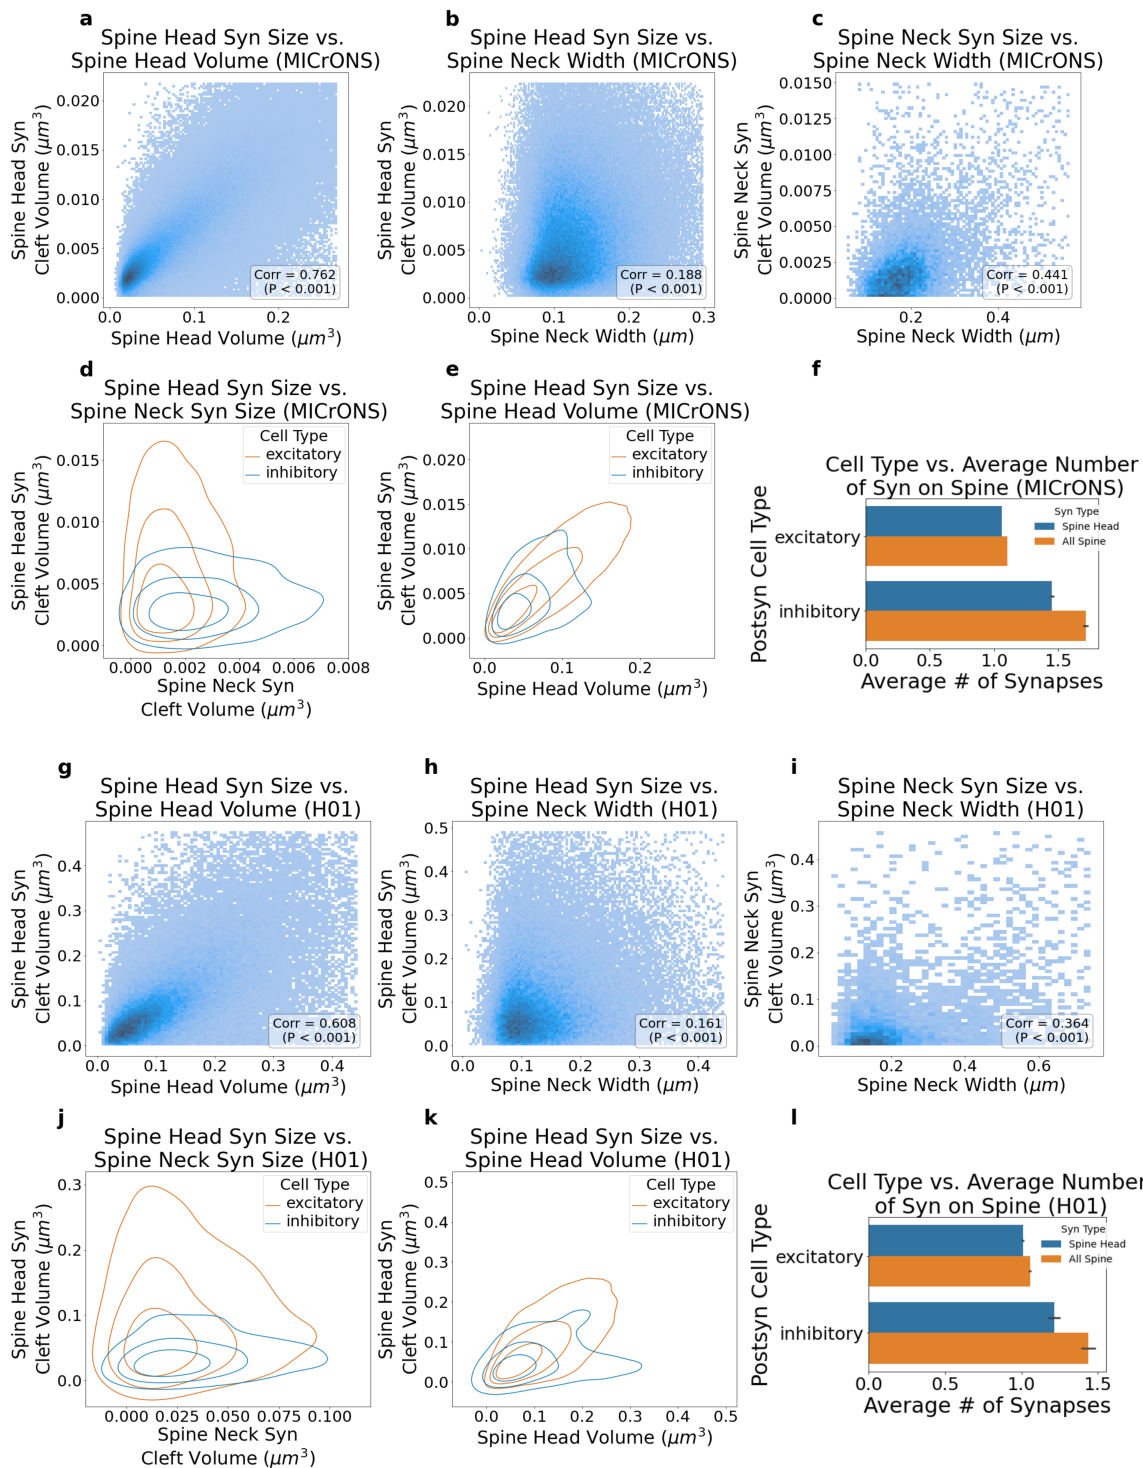


Fig. 23. Postsynaptic Spine Feature Analysis Here we compare the distributions and correlations of certain spine features; replicating and expanding on previous work. The MICrONS dataset is analyzed in **a-f**) and the same analysis is repeated for H01 in **g-l**). **a-b,g-h**) As expected, for synapses onto the spine **head**, the size of the synaptic cleft and the volume of the spine head mesh are strongly positively correlated, while cleft size and neck width are positively but more weakly correlated ([Harris and Stevens, 1989](#); [Arellano et al., 2007](#)). **c,i**) For synapses onto the spine **neck**, the width of the spine neck and the synaptic cleft volume of synapses are positively correlated. **d,j**) KDE of the joint distribution of the spine neck synaptic cleft volume with spine head synaptic cleft volume for different postsynaptic cell types (exc/inh), illustrating the different joint distributions for each cell type. For synapses onto spine **heads**, synaptic size has a wider range for excitatory cell spines than inhibitory cell spines in both volumes. For synapses onto spine **necks**, the range of synaptic size is larger for inhibitory cells in the MICrONS volume, but similar in the H01 dataset. **e,k**) Spine head volume is positively correlated with spine head synaptic cleft volume for both excitatory and inhibitory neurons in both datasets. **f,l**) Average number of synapses on all spines and spine heads for different cell types, indicating that in both datasets inhibitory spines receive more synapses per spine than excitatory spines.

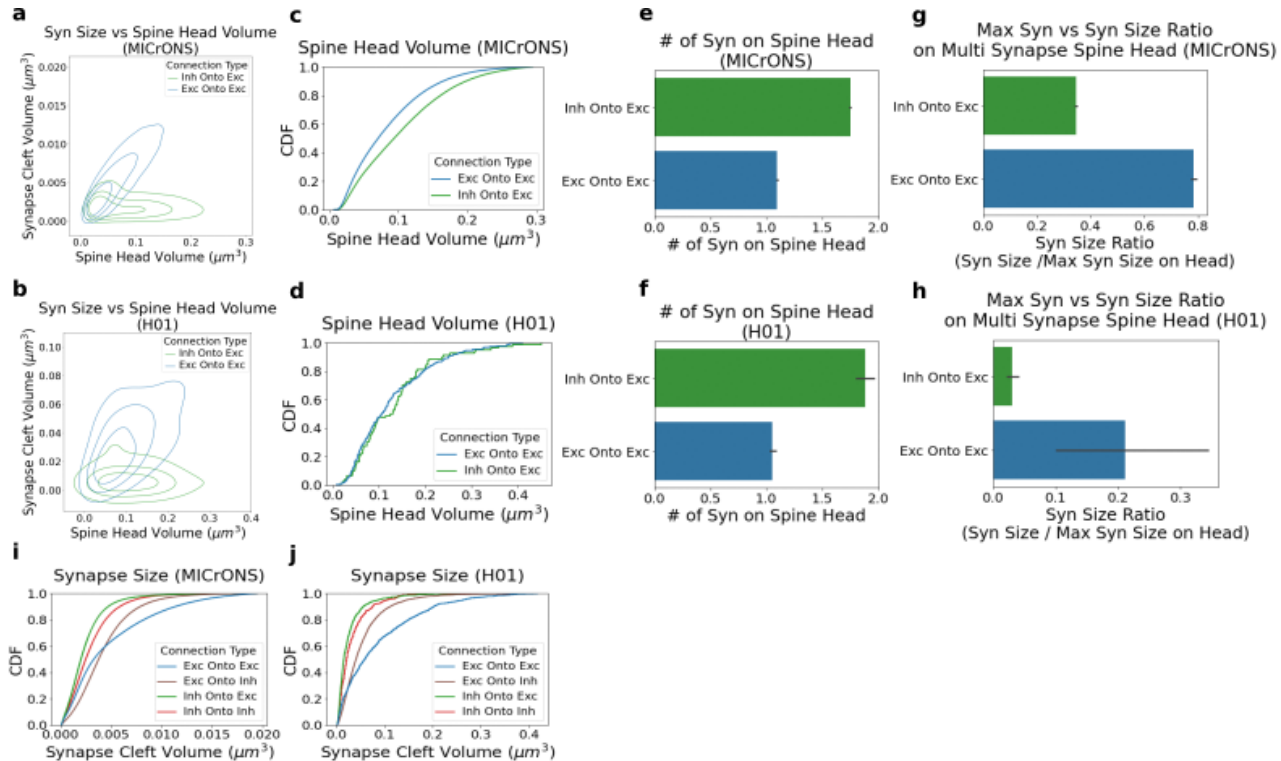


Fig. 24. Spine and Synapse Connectivity Analysis. We revisited the spine analysis in Fig. 23 taking into account information about the identity of the presynaptic neuron for each synapse. **a-b)** KDE distribution relating the postsynaptic spine head volume and the synapse cleft volume for synapses onto excitatory cells given different presynaptic cell types. In both datasets a significant positive correlation between spine head volume and synapse size is observed only when the source cell is excitatory but not when the source is inhibitory. **c-d)** CDF of the spine head volume for postsynaptic excitatory cells given different presynaptic cell types. For the MICrONS dataset, inhibitory presynaptic cells typically target larger spine heads but this trend is not significant in H01. **e-f)** As a possible explanation of why inhibitory cells target larger spine heads, a plot of the average number of synapses on a spine head conditioned on the presynaptic cell type shows that spine heads targeted by inhibitory neurons generally have two synapses as opposed to a mean closer to one synapse per spine for excitatory synapses. **g,h)** Expanding on the observation that inhibitory cells typically synapse onto spines with more than one synapse, for spines with multiple synapses, we plot the relative size of a spine head synapse to the size of the largest synapse on that same spine head (mean \pm std) given different presynaptic cell types. We observe that the synapse from an inhibitory source is typically much smaller than the largest synapse on the spine head in both the MICrONS and H01 dataset. **i,j)** CDF of the distribution of synapse cleft volumes for different connections types show a similar trend between the MICrONS and H01 dataset where the synapses with excitatory presynaptic cells are typically larger than inhibitory cells and synapses onto inhibitory cells are typically smaller than those onto excitatory.

a

Graph	MICrONS (Raw)	MICrONS (Auto)	H01 (Raw)	H01 (Auto)
# nodes (Exc/Inh)	44683 / 5104	44683 / 5104	2085 / 1236	2085 / 1236
# out edges (Exc/Inh)	2608436 / 1462930	263643 / 692756	12567 / 3924	5646 / 1236
mean in degree (Exc/Inh)	76.37 / 129.14	17.76 / 31.89	3.45 / 7.53	0.88 / 4.09
mean out degree (Exc/Inh)	58.38 / 286.62	5.90 / 135.70	6.03 / 3.17	2.71 / 1.00
mean axon sk length (μm)	1044 / 3256	421 / 1397	593 / 329	510 / 251
mean dendr sk length (μm)	3220 / 3150	2618 / 2330	3248 / 1369	2876 / 1229
% nodes in giant component	99.95	99.58	97.68	83.14
mean s.p. (shortest path)	3.61	4.78	9.20	4.37
mean s.p. undirected	3.06	3.87	6.47	9.73
95th percentile s.p.	5.00	7.00	16.00	9.25
95th percentile s.p. undirected	4.00	5.00	11.00	18.00
mean s.p. along Exc edges	3.90	8.33	13.15	1.00
mean s.p. along Exc edges undirected	3.24	5.31	7.39	9.43
mean s.p Exc-Exc/Inh-Inh	3.60 / 3.50	4.80 / 4.43	9.30 / 9.38	5.10 / 1.83
mean s.p Exc-Exc/Inh-Inh undirected	3.06 / 3.00	3.88 / 3.80	6.44 / 6.51	9.68 / 9.75

Fig. 25. Connectome Network Statistics Table a) Network statistics of the MICrONS and H01 connectomes where "edges" refer to synapses between neurons, "raw" refers to synaptic data before any processing with NEURD, and "auto" refers to the connectome produced after the decomposition pipeline and automated proofreading.

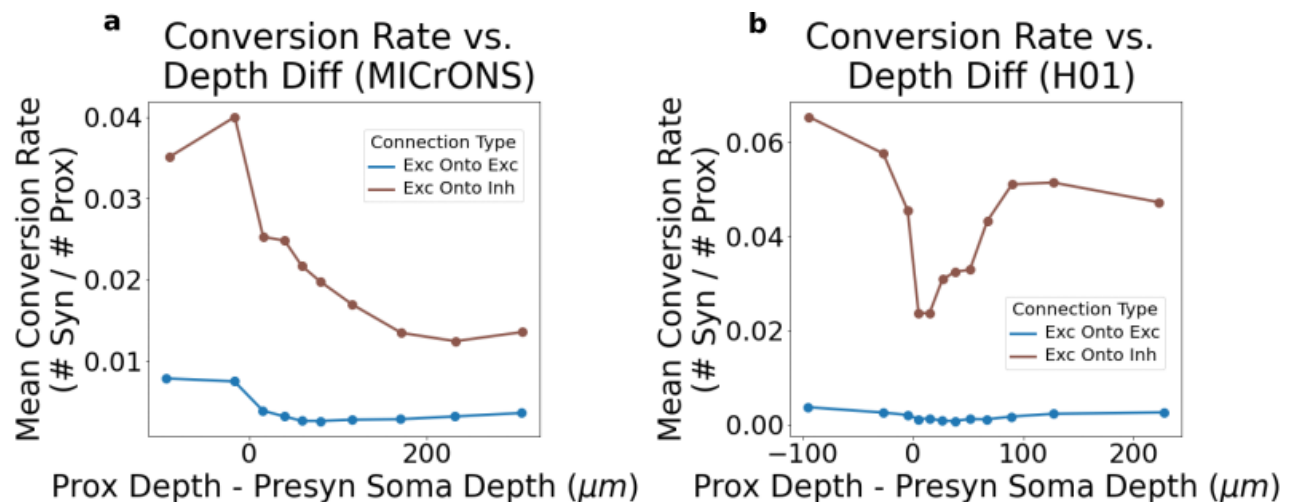


Fig. 26. Conversion Rate vs Proximity Relative Depth. Proximities are binned (approximate equal depth bins) in terms of their relative depth to the presynaptic soma center (proximity depth - presynaptic soma depth) and then the mean conversion rate (number of synapses/ number of proximities) for that bin is computed for different connection types. **a)** Conversion rate as a function of relative proximity depth in the MICrONS volume, demonstrating that the conversion rate for excitatory connections onto both excitatory and inhibitory postsynaptic cells is greater when the proximity is above the soma (for both connection types) **b)** Same plot for H01; the conversion rate as a function of relative proximity depth, demonstrating a greater conversion rate above the soma than below, but with an additional reduction in conversion rates close to the soma that is not seen in MICrONS.

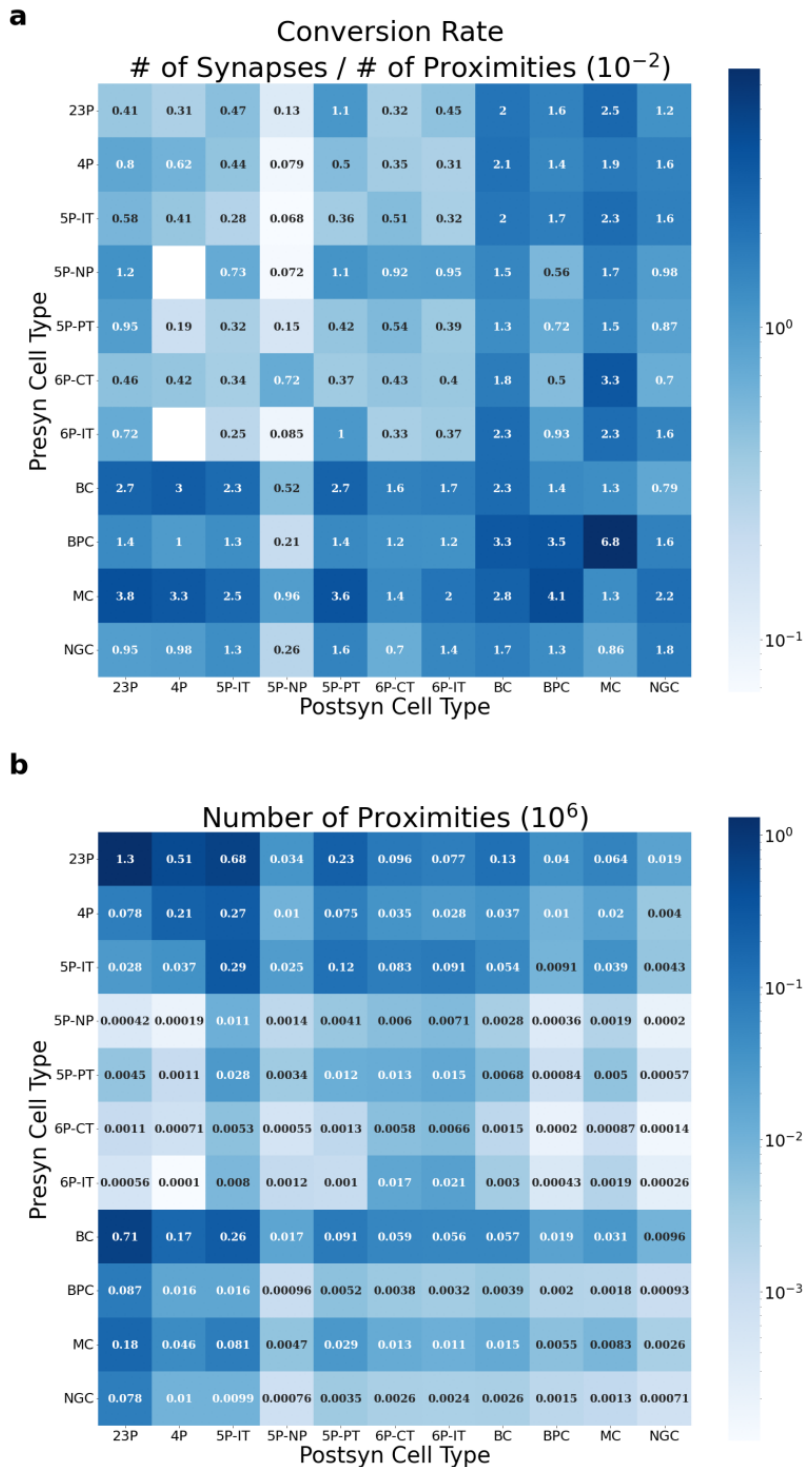


Fig. 27. Conversion Rate Cell Type Matrix The conversion rates (number of synapses / number of proximities) for different presynaptic and postsynaptic cell type pairs. The cell type labels are determined by the GNN whole neuron classifier. Proximities are filtered to only include those with the following features: less than $3 \mu\text{m}$ proximity distance, dendrite only postsynaptic compartment, presynaptic proximity width less than 130 nm (to exclude myelinated axon), presynaptic and postsynaptic cell type labels with at least a 70% confidence for each from the GNN classifier. **a)** Conversion rate for different cell type presynaptic and postsynaptic combinations **b)** Number of proximities in dataset used to calculate conversion rate

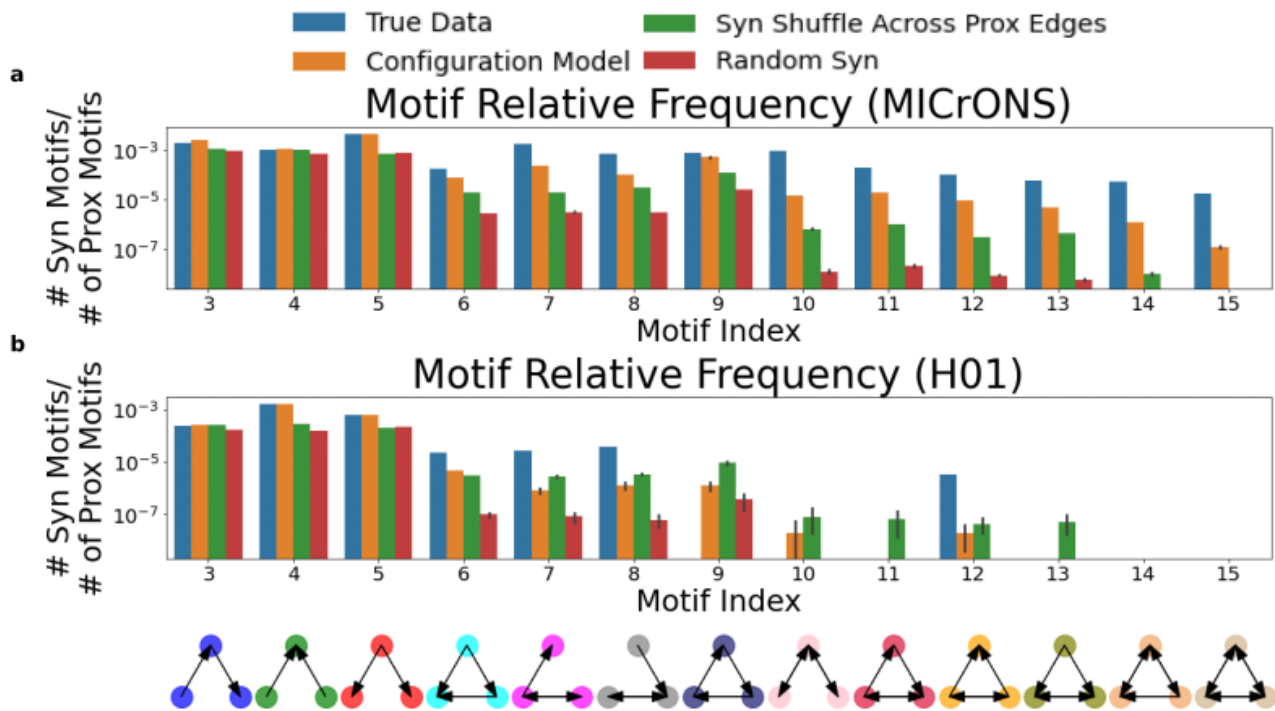


Fig. 28. Higher Order Triangle Motif Analysis. We count the number of directed triangle motifs in the synaptic and proximity connectome and compare the observed ratios to null ratios from three different models: first, a model where synaptic degree distribution is held the same but edges are shuffled (configuration model), second, a model where the synaptic edges are shuffled only between neurons with an existing proximity edge, or third, a model where synapses are random shuffled between neurons regardless of proximity. **a)** MICrONS dataset relative frequencies (duplicated from Fig. 6j) showing that the relative frequency of higher motifs in the synaptic connectome decreases as the number of edges in the motif increases (more higher order), but are consistently higher than the null model controls (250 random graph samples for each null distribution comparison). **b)** H01 dataset relative frequencies showing that the relative frequency of higher motifs in the synaptic connectome decreases as the number of edges in the motif increases. The observed motif frequencies are again higher than the null model, but many of the motifs with more than three directed edges are not observed due to the more incomplete reconstruction of neurons in H01 (400 random graph samples for each null distribution comparison, more samples were computed than MICrONS because computation was faster with a smaller connectome).