

High-throughput analysis of dendrite and axonal arbors reveals transcriptomic correlates of neuroanatomy

Received: 1 December 2023

Accepted: 16 July 2024

Published online: 27 July 2024

 Check for updatesOlga Gliko¹✉, Matt Mallory¹, Rachel Dalley¹, Rohan Gala¹, James Gornet², Hongkui Zeng¹, Staci A. Sorensen¹ & Uygur Sümbül¹✉

Neuronal anatomy is central to the organization and function of brain cell types. However, anatomical variability within apparently homogeneous populations of cells can obscure such insights. Here, we report large-scale automation of neuronal morphology reconstruction and analysis on a dataset of 813 inhibitory neurons characterized using the Patch-seq method, which enables measurement of multiple properties from individual neurons, including local morphology and transcriptional signature. We demonstrate that these automated reconstructions can be used in the same manner as manual reconstructions to understand the relationship between some, but not all, cellular properties used to define cell types. We uncover gene expression correlates of laminar innervation on multiple transcriptomically defined neuronal subclasses and types. In particular, our results reveal correlates of the variability in Layer 1 (LI) axonal innervation in a transcriptomically defined subpopulation of Martinotti cells in the adult mouse neocortex.

The shape of dendrites and axons, their distribution within the neuropil, and patterns of their long-range projections can reveal fundamental principles of nervous system organization and function. In the cortex, much of our understanding depends on the anatomical and functional descriptions of cortical layers. Yet, the origin and role of morphological and molecular diversity of individual neurons within cortical layers beyond broad subclass identities is poorly understood, in part due to low sample numbers. While molecular profiling techniques have recently improved by orders of magnitude, anatomical characterization remains time consuming due to continued reliance on (semi-)manual reconstruction.

Improvements in the throughput of the Patch-seq technique^{1–8} have enabled measurement of electrophysiological features, transcriptomic signatures, and local morphology in slice preparations for thousands of neurons in recent studies^{5,6}. In these repetitive experiments where maintaining a high throughput is a primary goal⁹, the brightfield microscope's speed, ease of use, and ubiquity make it an attractive choice to image local neuronal morphology. While this choice helps to streamline the experimental steps, morphological reconstruction remains a major bottleneck of overall throughput, in

part due to limited imaging resolution, even with state-of-the-art semi-manual tools⁵.

A rich literature exists on automated segmentation in sparse imaging scenarios. However, these methods typically focus on high-contrast, high-resolution images obtained by optical sectioning microscopy (i.e., confocal, two-photon, and light-sheet)^{10–18}, and are not immediately applicable to brightfield images because of the significantly worse depth resolution and the complicated point spread function of the brightfield microscope. Moreover, segmentation of full local morphology together with identification of the axon, dendrites, and soma has remained elusive for methods tested on image stacks obtained by the brightfield microscope^{19–22}. Therefore, we first introduce an end-to-end automated neuron reconstruction pipeline (Fig. 1a) to improve scalability of brightfield 3D image-based reconstructions in Patch-seq experiments by a few orders of magnitude. We note that our primary goal is not to report on a more accurate reconstruction method per se. Rather, we aim to demonstrate how automated tracing, with its potential mistakes, can be leveraged to rigorously address certain scientific questions by increasing the throughput. To this end, we select a set of brightfield images and use

¹Allen Institute, Seattle, WA, USA. ²California Institute of Technology, Pasadena, CA, USA. ✉ e-mail: olgag@alleninstitute.org; uygars@alleninstitute.org

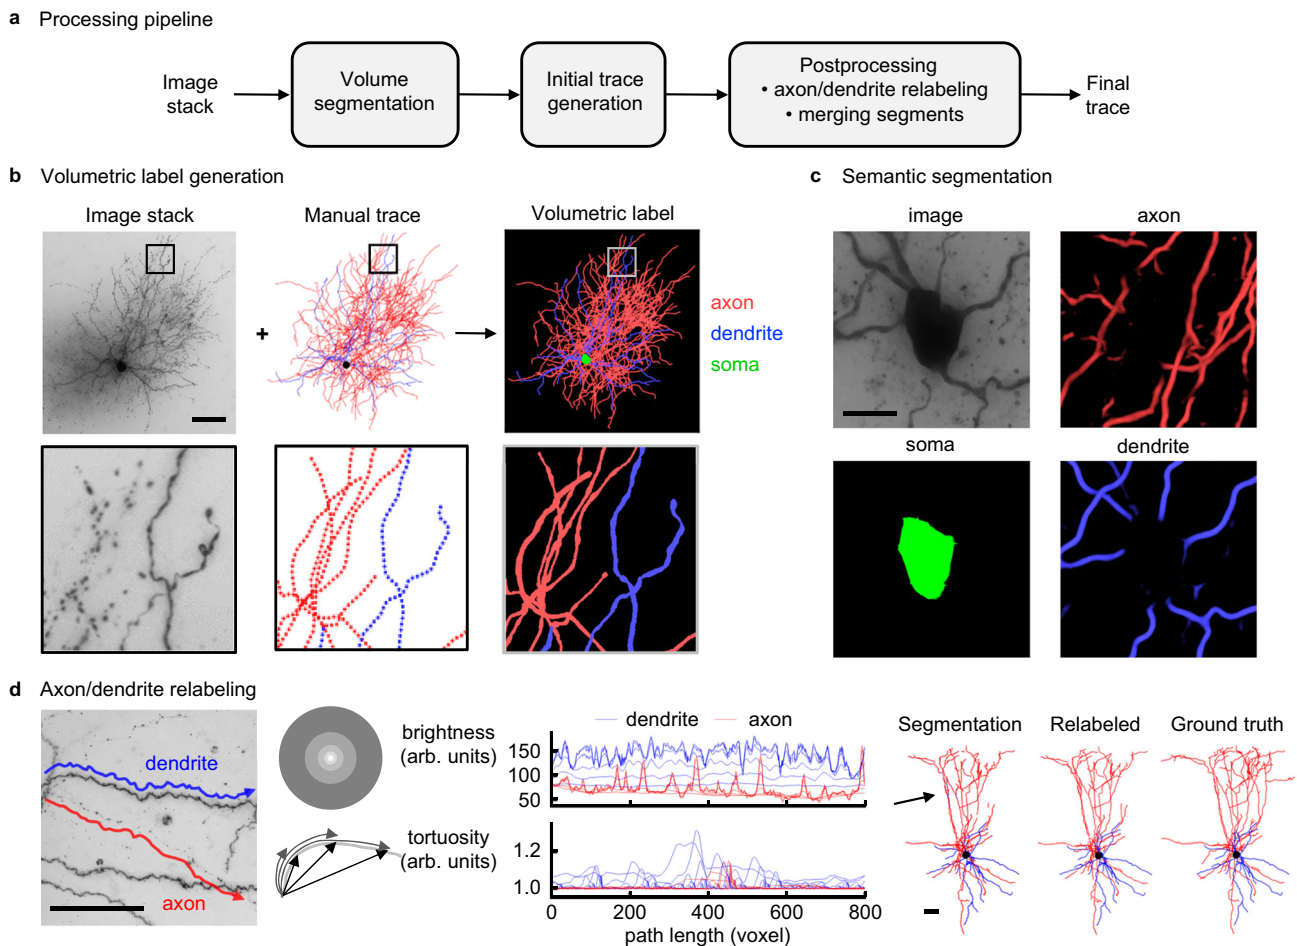


Fig. 1 | Neuron reconstruction pipeline for in-slice bright-field images of biocytin-filled neurons. **a** Processing pipeline. Convolutional neural network (CNN) segmentations of 3D image stacks are post-processed by custom machine learning tools to produce digital representations of neuronal morphologies. **b** Topology preserving fast marching algorithm generates the volumetric label from raw image stack and manual skeletonization. Image stack (minimum intensity projection), volumetric label (maximum intensity projection). Dendrites (blue), axons (red), soma (green) are separately labeled to train a supervised CNN model. Scale bar, 100 μm . **c** Semantic segmentation provides accurate soma location and

boundary. Image (maximum intensity projection), axon, soma, dendrite (maximum intensity projection). Scale bar, 10 μm . **d** Axon/dendrite relabeling. A neural network model predicts node labels from multiple image brightness and trace tortuosity features based on local contexts of different size along the initial trace. (left, example image of dendrite and axon segments (minimum intensity projection); middle, corresponding feature plots; right, automated traces of test neuron with/without relabeling vs. manual trace). Dendrites (blue), axon (red), soma (black). Arrow indicates nodes mislabeled by segmentation and corrected during post-processing. Scale bars, 50 μm . Source data are provided as a Source Data file.

the corresponding manually reconstructed neuron traces to assign voxel-wise ground truth labels (axon, dendrite, soma, background). Next, we design a custom deep learning model and train it on this curated ground truth dataset to perform 3D image segmentation using volumetric patches of the raw image as input. We implement fully automated post-processing steps, including basal vs. apical dendrite separation for pyramidal cells, for the resulting segmentations to obtain annotated traces for each neuron. We compare the accuracy of these automated traces with a held out set of manual reconstructions, based on geometrical precision, a suite of morphometric features, and arbor density representations derived from the traces.

We utilize this pipeline to reconstruct a large set of neurons from Patch-seq experiments, and use the transcriptomic profiles captured from the same cells to systematically search for gene subsets that can predict certain aspects of neuronal anatomy. The existence of a hierarchical transcriptomic taxonomy²³ enables studying subsets of neurons at different levels of the transcriptomic hierarchy. At the finest scale of the hierarchy (transcriptomic types or “t-types”²⁵), we study seven interneuron types and focus on a transcriptomically defined sub-population of LI-projecting, *Sst* gene expressing neurons (*Sst* cells) that correspond to Martinotti cells (See, for instance, ref. 24). While

previous studies have elucidated the role of Martinotti cells in gating top-down input to pyramidal neurons via their LI-innervating axons^{25,26}, the wide variability in the extent of LI innervation behind it is not well understood. Our results suggest transcriptomic correlates of the innervating axonal mass, which may control the amount of top-down input to canonical cortical circuits. Our approach represents a general program to systematically connect gene expression with neuronal anatomy in a high-throughput and data-driven manner.

Results

An automated morphology reconstruction pipeline for bright-field microscopy

As the first step to automate the reconstruction of in-slice brightfield images of biocytin-filled neurons, we curate a set of manually traced neurons. While this set should ideally be representative of the underlying image space, it should also be as small as possible to facilitate downstream cross-validation studies via an abundance of cells not used during training. We thus choose 51 manually traced neurons as the training set to represent the underlying variability in the image quality and neuronal identity (via the available t-type labels). We develop a topology-preserving variant of the fast marching algorithm¹³

to generate volumetric labels from manual traces (Fig. 1b). We train a convolutional neural network (U-Net)^{16,27–29} using image stacks and labels as the training set and employing standard data augmentation strategies to produce initial segmentations of neuronal morphologies (Fig. 1b,c). While knowledge of axonal vs. dendritic branches informs most existing insight, their automated identification poses a challenge due to the limited field-of-view of artificial neural networks. We find that image and trace context that is in the vicinity of the initial segmentation is sufficient to correct many axon vs. dendrite labeling mistakes in an efficient way because this effectively reduces the problem to a single dimension, i.e., features calculated along the 1D initial trace (Fig. 1d). We further algorithmically post-process the segmentations to correct connectivity mistakes introduced by the neural network and obtain the final reconstruction of axons and dendrites (Fig. 1d, Fig. 2a, Figs. S1–S11, Methods). We observe that this approach offers marked improvements in tracing quality compared to a previous large-scale effort focusing on fluorescent, optical-sectioning microscopy^{30,31} (Fig. S12). Moreover, Fig. S13 qualitatively shows that its segmentation quality remains robust when tested, without any tuning, with images from different species and brain structures^{32–34}. (Fine-tuning the existing model with a small training set representing the tissue of interest should further improve performance.) The overall pipeline produces neuron reconstructions in the commonly-used swc format³⁵ from raw image stacks at a rate of ~6 cells/day with a single GPU card (Methods). Our setup uses 16 cards to achieve two orders of magnitude improvement in speed over semi-manual segmentation⁵ with one anatomist. We have so far processed the cells reported in ref. 5, which mapped neurons to an existing taxonomy of transcriptomic cell types²³ and introduced a transcription-centric, multi-modal analysis of inhibitory neurons. We have also processed a set of ~700 excitatory neurons which are analyzed in ref. 36. (While we typically display only the apical and basal dendrite segmentations for excitatory cells, the method can also trace and label the local axon when it is captured in the slice (Fig. S14).) After quality control steps (Methods), we focus on a set of 813 interneurons for further study in this paper.

The proposed pipeline produces end-to-end automated reconstructions in single-cell imaging scenarios. However, in practice, neurons are patched near each other to increase the throughput of physiological and transcriptomic characterization. The resulting image, which typically centers on the neuron of interest, can therefore contain neurites from other neurons. Neurites from off-target neurons within the image stack cannot be properly characterized because they rarely remain in the field of view. As part of algorithmic post-processing and quality control, disconnected segments are removed automatically when they remain relatively far from the cell of interest (Methods). When multiple neurons are patched in close proximity, quality control by a human is needed to check for and remove nearby extraneous branches. To ensure the integrity of presented results with minimal manual effort, the cell is not used if quality control suggests the existence of nearby branches and a manual trace is not available. If the manual trace already exists, we simulate manual branch removal based on a mask obtained from the manual trace (Methods). We report quantification results separately for reconstructions obtained with/without nearby branch removal.

Evaluation of reconstruction accuracy

We evaluate the quality of automated traces by comparing them to the manual traces which we regard as the ground truth. To compare a pair of automated and manual traces, we perform a bi-directional nearest-neighbor search to find correspondence nodes in both traces within a certain distance¹³. A node in the automated trace that has (does not have) a corresponding node in the manual trace is referred to as a true (false) positive node, and a node in the manual trace that does not have a corresponding node in the automated trace is referred to as a false

negative node. We calculate this metric separately for axonal and dendritic nodes, as well as for all nodes regardless of the type, and compute corresponding precision, recall, and f1-score (harmonic mean of the precision and recall) values. These metrics indicate how well the automated trace captures the layout of the axon/dendrites/neurites in a reconstructed neuron. Figure 2b and Table S1 display that, at a search radius of 10 μm , the mean f1-score is above 0.8 for both axonal and dendritic morphologies. Therefore, we expect these automatically generated traces to perform comparably to their manually generated counterparts in analyses that do not require a resolution better than 10 μm , as we demonstrate below. (Fig. S15 shows a basic qualitative estimate of cross-human tracing discrepancy based on one test cell and Table S2 quantifies this study.)

To further assess the similarity in the arbor layout and other aspects of morphology that are not captured by the node correspondence study, we use standard morphometric features. We find that while many features summarizing the overall morphology can be accurately predicted, features related to the topology of arbor shapes, such as maximum branch order, are prone to mistakes (Fig. 2d and Fig. S16).

While the quantitative analyses described above both suggest that automated reconstruction succeeds in broadly capturing neuronal morphology, including separation of axonal vs. dendritic branches, they also demonstrate that important differences nevertheless remain between automated and manual traces. Therefore, to robustly analyze anatomical innervation patterns against potential topological mistakes introduced by the automated routine, we develop a 2D arbor density representation (ADR)^{12,37–40} of axons and dendrites registered to a common laminar axis defined by cortical pia and white matter boundaries. Here, the vertical axis represents the distance from pia and the horizontal axis represents the radial distance from the soma node (Fig. 2c). Note that this 2D representation still requires 3D imaging because many branches become undetectable in 2D projected images due to noise (e.g., Fig. 1b). Moreover, standardizing the orientations of the brain and the tissue slice is challenging in high-throughput experiments so that the rotation around the laminar axis would be hard to control in 3D representations.

At the level of transcriptomic types, the ADRs calculated from automated reconstructions appear similar to those calculated from manual segmentations based on the root-mean-squared difference between them (Fig. 2d).

To better quantify this similarity, we compare the performance of the ADR against that of morphometric features⁴¹ by training classifiers to predict t-types and subclasses (Sst, Pvalb, Vip, Sncg, Lamp5)⁵. We find that the ADR is not statistically significantly worse than the morphometric features in terms of classification accuracy (Boschloo's exact test, asymptotically exact harmonic mean of p-values over multiple runs⁴²: $p = 0.81$ for t-types, $p = 0.26$ for subclasses, Methods), consistent with ref. 43 (Fig. 3a, b).

We also test robustness against imperfections due to fully automated tracing by comparing the classification accuracy obtained from automated tracing versus manual tracing on the same set of cells. End-to-end automation appears to perform similarly as manual tracing in cell type prediction based on ADRs (Fig. 3c, d) and morphometric features (Fig. S19). We finally compare cell type identification based on automatically generated ADRs vs. manually generated morphometric features. We find that they are not significantly different in t-type classification (Boschloo's test, harmonic mean $p = 0.72$), but the ~5% advantage of manual morphometric features in subclass classification is statistically significant (Boschloo's test, harmonic mean $p = 0.01$).

Beyond the comparative aspect, these results demonstrate a correspondence between gene expression and the anatomy of local arbors as represented by the proposed registered 2D ADRs, which agrees with previous findings with morphometric features for these cells⁵. (subclass accuracy of ~79% vs. random at 20%, most abundant

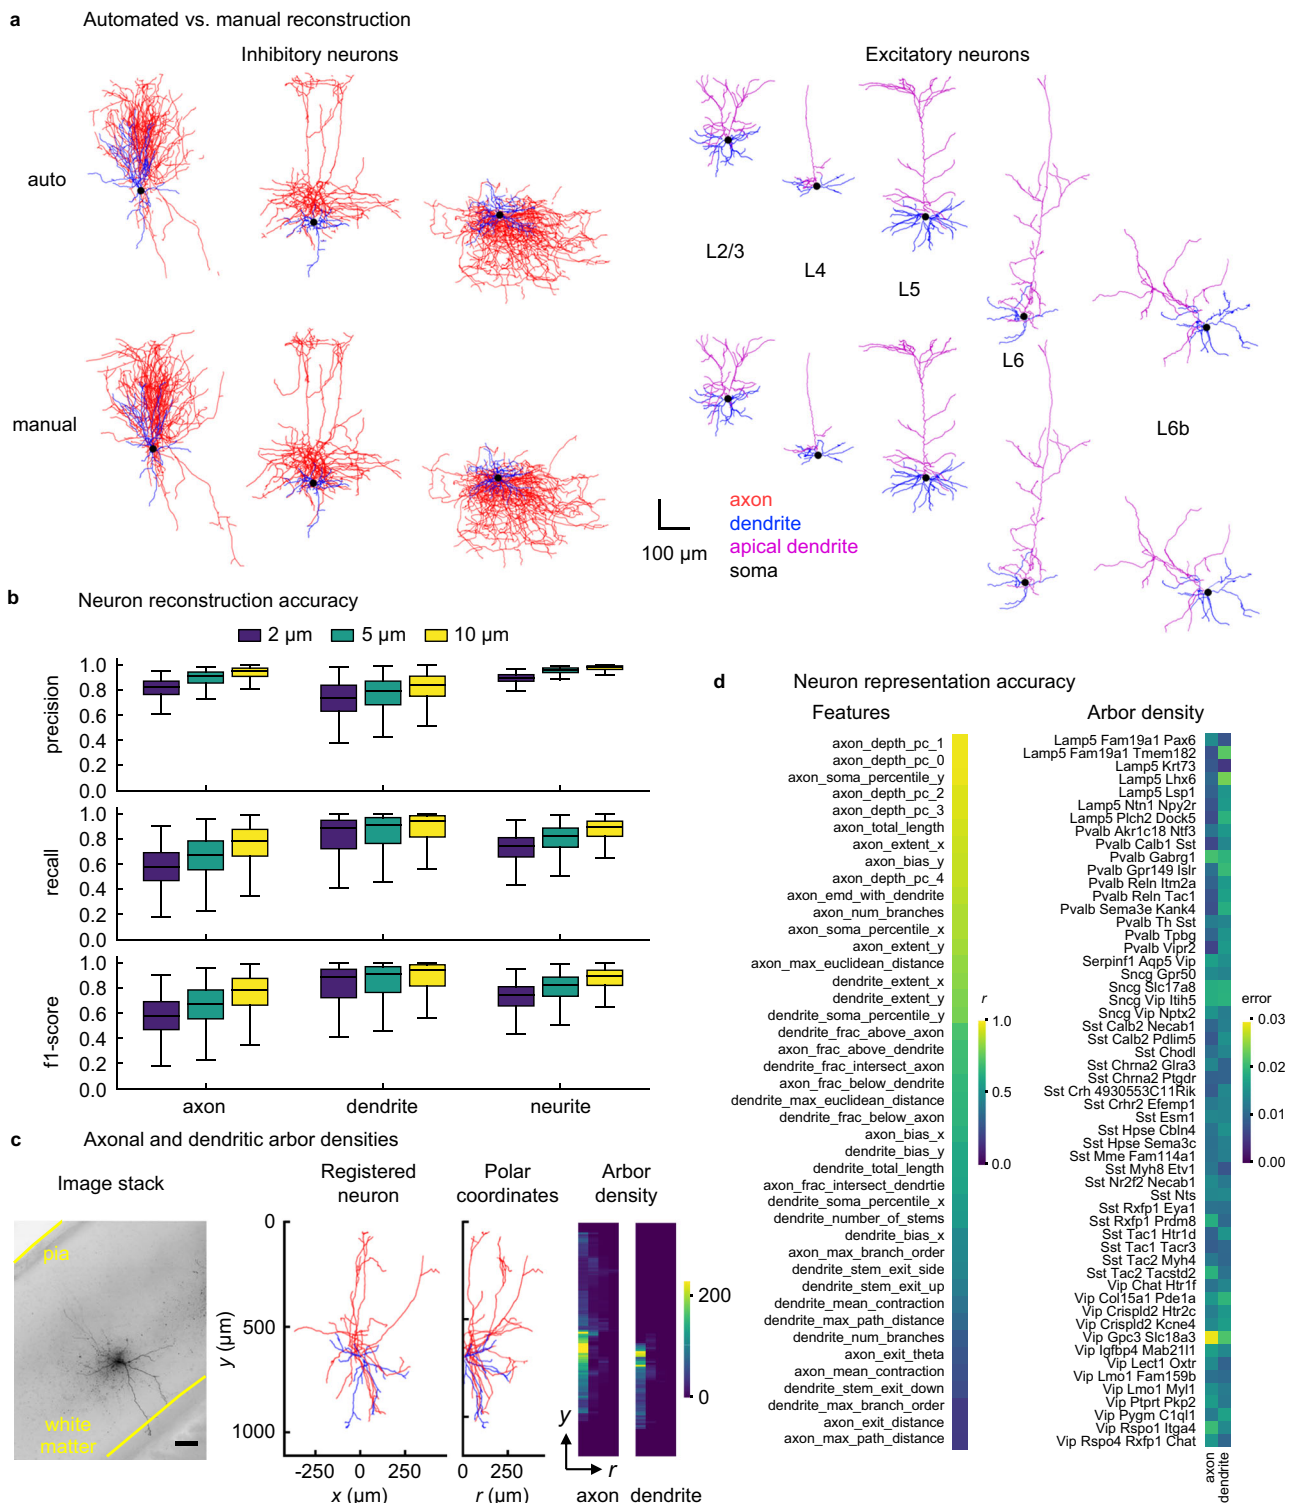


Fig. 2 | Assessing quality of automated reconstructions. **a** Automated and manual traces of example test neurons (left, inhibitory neurons; right, excitatory neurons—apical and basal dendrites are assigned for excitatory cells.) **b** Neuron reconstruction accuracy. Precision, recall, and f1-score values are calculated by comparing automated and manual trace nodes within a given distance (2, 5, and 10 μ m). Box plots ($n = 340$ cells) are shown for axonal, dendritic and neurite (combined axonal and dendritic) nodes. The box extends from the first quartile (Q1) to the third quartile (Q3) of the data, with a line at the median. The whiskers extend

from the box to the farthest data point lying within 1.5x the inter-quartile range (IQR) from the box. Mean \pm s.d. of the values over 340 cells are provided in Table S3. Scatter plots are shown in Fig. S17. **c** Generation of 2D axonal and dendritic ADRs. Scale bar, 100 μ m. **d** Pearson's r values (left) and average root-mean-squared error (right) between automatically vs. manually generated features (left and Fig. S16) and ADRs for each t-type (right and Fig. S18). ADRs are normalized to have unit norm. Source data are provided as a Source Data file.

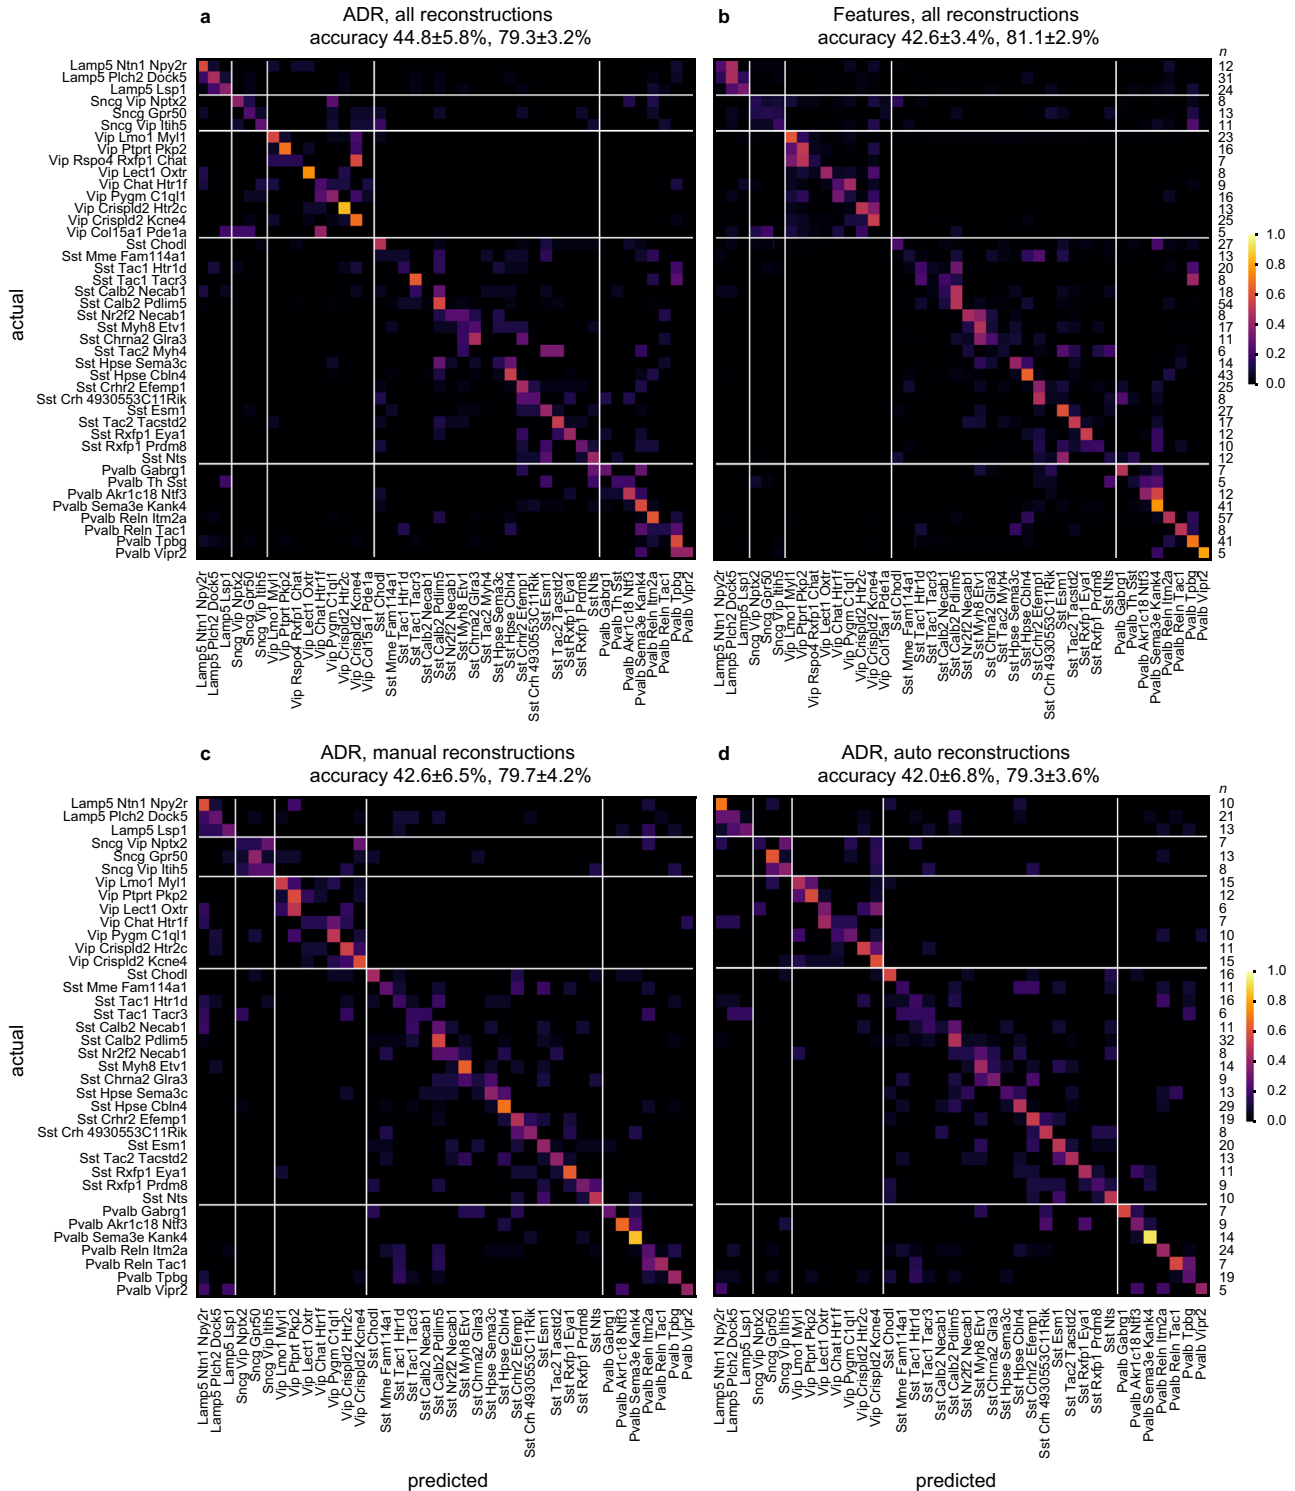


Fig. 3 | Comparison of cell type classification accuracy based on the ADR vs. a set of classical morphometric features. Confusion matrix for the classification of 42 t-types based on axonal and dendritic ADRs (a) and morphometric features (b), using a combination of 246 automatically and 501 manually reconstructed cells. Confusion matrix for the classification of 38 t-types based on ADRs, using

488 manually (c) and automatically (d) reconstructed cells. Accuracy values reported in the headers refer to mean ± s.d. of the overall t-type and t-subclass classifiers, respectively, across cross-validation folds. Rightmost columns list the number of cells in each t-type (*n*). Source data are provided as a Source Data file.

label at 47%; t-type accuracy of ~45% vs. random at ~2%, most abundant label at ~8%). When the transcriptomic type assignments are incorrect, cells are rarely assigned to transcriptomically far-away clusters based on the ADR or morphometric features, as demonstrated by the dominance of the entries around the main diagonal in Fig. 3. Note that the rows and columns of these confusion matrices are

organized based on the reference taxonomy to reflect transcriptomic proximity (Figs. S20 and S21). Therefore, the relative inaccuracy at the t-type level could be attributed to aspects of morphology not captured by the ADR or morphometric features (e.g., synapse locations), or other observation modalities (e.g., physiological, epigenetic) being key separators between closely related t-types.

Table 1 | Statistical significance and effect size values for predicting anatomical features from gene expression via sparse linear regression for seven different cell types

	Sst Calb2 Pdlm5	Sst Hpse Cbln4	Sst Chodl	Pvalb Reln Itm2a
L1 axon	0 / 0.40	1.000	0.637	0.023 / 0.13
L2/3 axon	0.013 / 0.30	1.000	1.000	0 / 0.35
L4 axon	0.689	0.332	0.607	0 / 0.35
L5 axon	0.033 / 0.15	0.013 / 0.25	1.000	0.102
L1 dendrite	0.042 / 0.17	1.000	0.689	0.088
L2/3 dendrite	0.033 / 0.15	1.000	0.332	0.013 / 0.27
L4 dendrite	0.697	0 / 0.40	0.013 / 0.40	0.058
L5 dendrite	0 / 0.19	0.102	0.393	0 / 0.25
soma depth	0 / 0.26	0.023 / 0.21	0.246	0 / 0.34
axon centroid	0.023 / 0.16	0.210	0.058	0 / 0.34
dendrite centroid	0.150	0.023 / 0.26	0.096	0 / 0.39

	Pvalb Tpbpg	Lamp5 Lsp1	Lamp5 Plch2 Dock5
L1 axon	0.033 / 0.21	0.013 / 0.56	0.042 / 0.40
L2/3 axon	0.023 / 0.27	0.042 / 0.52	0.323
L4 axon	0.216	1.000	0.351
L5 axon	0.058	0.135	0.081
L1 dendrite	0 / 0.42	0.283	0 / 0.40
L2/3 dendrite	0.042 / 0.21	1.000	1.000
L4 dendrite	0.074	0.393	1.000
L5 dendrite	0.013 / 0.26	1.000	1.000
soma depth	0.013 / 0.29	0.023 / 0.48	0.067
axon centroid	0.058	0.013 / 0.44	0.074
dendrite centroid	0.013 / 0.38	0.023 / 0.46	0.074

For each entry, the FDR-corrected p -value as calculated by a non-parametric one-sided shuffle test is listed. If the value is considered statistically significant at $p \leq 0.05$, the R^2 value is also displayed (p / R^2). p -values less than or equal to 0.05 and R^2 values larger than or equal to 0.25 are shown in bold. A p value of 0 indicates that the calculated p value is less than 0.001, the sensitivity of the shuffle test, and less than 0.013 after FDR correction. Source data are provided as a Source Data file.

Correlates of gene expression and laminar innervation

Having established that registered 2D ADRs are as successful as a standardized, rich set of morphometric features in predicting transcriptomic identity and that ADRs can be generated in a fully automated manner from raw images with only mild loss in performance, we aim to uncover more explicit connections between gene expression and anatomy as captured by the ADR. Since layer-specific axon and dendrite innervations are prominently and reliably captured by the ADR, we study their transcriptomic correlates. We treat the search for genes that are predictive of laminar innervation strength (neurite length innervating a given layer) as a sparse regression problem^{44,45} (Methods), and focus on 7 t-types whose morphologies are well sampled in our dataset with the help of automated reconstruction (3 Sst, 2 Pvalb, 2 Lamp5 types). That is, we aim to uncover minimal gene sets whose expression can predict the amount of axonal and dendritic innervation of individual laminae as well as the locations of the soma and centroids of the axonal and dendritic trees along the laminar axis. Throughout, we control the false discovery rate (FDR) by applying multiple testing correction (Methods). Tables 1 and S3-S5 summarize these results. We observe that no single anatomical feature is significantly predictable from gene expression for all inhibitory t-types and every studied t-type has at least one significantly predictable anatomical feature. (Only the L4 dendritic innervation strength is significantly correlated with gene expression for cells of type Sst Chodl.)

Perhaps more interestingly, we find that the sets of laminae innervation-predicting genes within transcriptomically defined subclasses and t-types are highly reproducible (Table S5) and almost mutually exclusive (Fig. 4g). These observations support a connectivity-related organization of cortical cells⁴⁶. (Multiple discrete and continuous factors of variability may shape neuronal phenotypes^{4,47} and their dissection may not be possible by studying a subset of non-adjacent t-types.) They also put forth a related question: can gene expression further predict innervation strength of a single layer in a continuum?

Tuning laminar innervation within a cell type: a comparative study

To elucidate this question, we choose a transcriptomically defined subpopulation that is well-sampled in our dataset with the help of automated reconstruction, produces a large effect size in the gene regression study (Table 1), and has been a source of confusion due to its anatomical variability: Sst Calb2 Pdlm5 neurons²³ represent a transcriptomically homogeneous subset of Martinotti cells, which are inhibitory neurons with L1-innervating axons that gate top-down input to pyramidal neurons^{25,26}. However, the amount of axon reaching L1 varies widely across cells⁵. Tables 1 and S5 show that a small set of genes, including genes implicated in synapse formation, cell-cell recognition and adhesion, and neurite outgrowth and arborization⁴⁸⁻⁵⁰, can nevertheless predict the L1-innervating skeletal mass of neurons belonging to this homogeneous population ($R^2 = 0.40$, $p < 0.001$, non-parametric shuffle test). Since the somata of this population are distributed across L2/3, L4, and L5 (Fig. 4b), one potential explanation for this result is that gene expression is correlated with the overall depth location of the cells rather than L1 innervation strength in particular (Fig. 4c). Therefore, we repeat the sparse regression study after removing the piecewise linear contribution of soma depth to L1 innervation (linear fit and subtraction for only the L1 innervating subpopulation because the relationship is trivially nonexistent for the non-innervating subpopulation, Methods). We find that the expression levels of a small set of genes are still statistically significantly predictive of L1 innervation: $R^2 = 0.31$, $p < 0.001$, non-parametric shuffle test. (Repeating with a linear fit and subtraction for the whole population does not change the qualitative result: $R^2 = 0.30$, $p < 0.001$.)

Next, we obtain a comparative perspective on the L1 innervation result for the Sst Calb2 Pdlm5 subset of Martinotti cells by juxtaposing this result with that for the cells of the Lamp5 Lsp1 type. Somata of these cells are also distributed across multiple cortical layers and their axons have highly variable levels of L1 innervation. Sparse regression again succeeds in finding a small set of genes whose expression level can predict L1 innervation ($R^2 = 0.56$, $p = 0.01$, non-parametric shuffle test, Table 1). However, it fails to uncover a statistically significant gene set after removing the piecewise linear contribution of soma depth: $R^2 = 0.06$, $p = 0.26$, non-parametric shuffle test. (Repeating with a linear fit and subtraction for the whole population does not change the qualitative result: $R^2 = 0.08$, $p = 0.39$.) That is, in contrast to the Sst Calb2 Pdlm5 cells, soma depth almost completely explains the variability in L1 innervation for cells in the Lamp5 Lsp1 population (Fig. 4e, f).

Lastly, we consider the possibility that the cells whose axons do not reach L1 are simply irrelevant to this study and bias the statistics. (Axons of 3 out of 52 cells in the Sst Calb2 Pdlm5 population, and 7 out of 22 cells in the Lamp5 Lsp1 population do not reach L1.) We repeat the above comparison after removing the cells whose axons don't reach L1 altogether from this study. Sparse regression still uncovers a statistically significant relationship between L1 innervation strength and a set of genes for the Sst Calb2 Pdlm5 population after removing the linear contribution of soma depth ($R^2 = 0.28$, $p < 0.001$). In contrast, it again fails to find a statistically significant relationship for the Lamp5 Lsp1 population ($R^2 = 0.05$, $p = 0.39$).

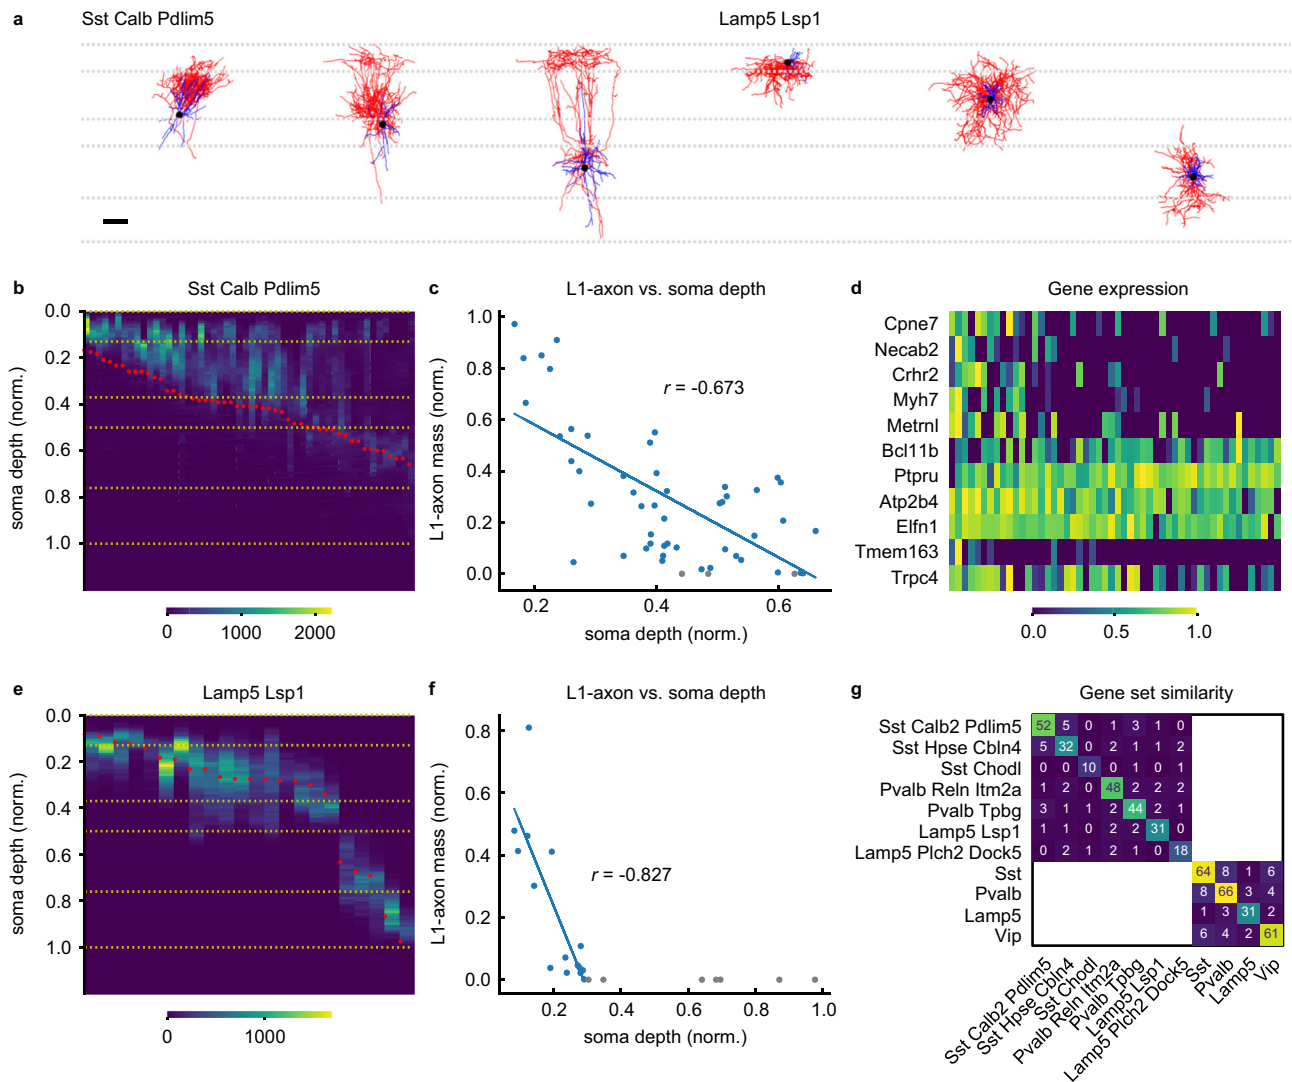


Fig. 4 | L1 axonal innervation correlates with expression of subset of genes in Martinotti cells. Example neurons of Sst Calb2 Pdlim5 and Lamp5 Lsp1 t-types (a). (Horizontal dashed lines indicate cortical layer boundaries. Scale bar, 100 μm .) 1D axonal arbor density for the 52 cells in the Sst Calb2 Pdlim5 t-type (b) and the 22 cells in the Lamp5 Lsp1 t-type (e). (Yellow horizontal dashed lines and red dots indicate cortical layer boundaries and soma depth, respectively). Normalized L1-axon skeletal mass vs. normalized soma depth (0:pia, 1:white matter boundary) for the Sst Calb2 Pdlim5 cells (c) and the Lamp5 Lsp1 cells (f). Lines fitted to cells with

nonzero L1 innervation. Cells whose axons don't reach L1 are shown in gray. Pearson's r values are shown. **d** Gene expression vs. L1-axon skeletal mass for the genes selected by the sparse regression analysis. (L1-axon mass decreases from left to right.) **g** Similarity matrix for the sets of laminae-predicting genes within transcriptomic types and subclasses. (See Table S5.) Each entry denotes the number of genes in the intersection between the corresponding row and column. Source data are provided as a Source Data file.

To summarize, while axons of Lamp5 Lsp1 cells appear to shift along the laminar axis according to their soma location within the cortical depth, soma location does not seem to dictate the axonal L1 innervation of Sst Calb2 Pdlim5 neurons, whose strength can nevertheless be predicted by gene expression. For both of these t-types, the automated reconstruction pipeline increased the sample size by more than 60% (Sst Calb2 Pdlim5: 63%, Lamp5 Lsp1: 69%), empowering the statistical analysis pursued here. Similarly, the sample counts for the t-types studied in Table 1 increased between 48% and 138%. (The increase over the whole dataset is 50%, from 543 to 813 cells.) Since t-types correspond to leaf nodes of the cell type hierarchy, their sample sizes are much smaller than the subclass-level counts. Therefore, automated reconstruction can be beneficial both by capturing more of the biological variability in single-cell morphologies of populations at the finest level of transcriptomically defined taxonomies and by enabling cross-validation schemes similar to the ones pursued here.

Discussion

While the classification of neuronal cell types is increasingly based on single cell and nucleus genomic technologies, characterization of neuron morphology—a classical approach—captures an aspect of neuronal identity that is stable over long time scales, is intimately related to connectivity and function, and can now be connected with genomic attributes through the use of simultaneous profiling techniques such as Patch-seq. Nevertheless, light microscopy-based methods of neuronal reconstruction often inadequately reproduce the determinant attributes of morphological signature, especially in high-throughput settings. Here, we have presented an end-to-end automated neuronal morphology reconstruction pipeline for brightfield microscopy, whose simple setup supports flexible, single or multimodal, characterization protocols. We have also proposed an arbor density representation as a descriptor of cortical neuronal anatomy that is robust against noise in high-throughput imaging scenarios as well as mistakes of automated reconstruction. Its success suggests that

detailed morphological reconstructions may ultimately not be necessary if the only aim is inferring the cell type label.

Through the use of sparsity arguments and statistical testing, we demonstrated that this pipeline can help reveal relationships between gene expression and neuronal anatomy, where a large number of anatomical reconstructions enables accurate inference in the presence of a large gene set. As an application, we studied the correlation between gene expression and laminar innervation on a Patch-seq dataset of cortical neurons⁵ and showed that the gene correlates of different innervation patterns have little overlap across transcriptomically defined subpopulations. While the same program can potentially also address the relationship between morphological and electrophysiological properties of neurons, the accuracy of automated reconstructions should be further improved for use in detailed compartmental models⁵¹.

Finally, we focused on axonal innervation of L1 by a transcriptomically defined subpopulation of Somatostatin-expressing Martinotti cells. We found that the innervation strength is relatively weakly correlated with soma depth for this cell type, but not all types. Moreover, a subset of genes can predict the remaining variability in the innervation strength after the effect of soma depth is removed, suggesting a control mechanism beyond simple shifting of the morphology within the cortical depth for this cell type. Considering that neurons in this population are thought to gate top-down input to cortical pyramidal neurons^{25,26}, this result suggests tuning of innervation strength in a continuum within the discrete laminar organization of the mouse cortex^{47,52–54}, potentially to improve task performance of the underlying neuronal network.

From a segmentation perspective, we believe our work represents a significant step forward by producing hundreds of automatically reconstructed morphologies obtained from the brightfield microscope (Figs. S1–S11). As demonstrated in the main text, these cortical neuron morphologies are statistically indistinguishable from their manually generated counterparts in certain aspects (e.g., cell type identification), but not in many others (e.g., arbor topology). Indeed, much further improvement is needed to achieve complete and accurate tracing of neurons. Nevertheless, advances in computer vision algorithms and computing infrastructure that can support complicated models and large datasets suggest that qualitative improvements may be within reach in the next few years. While our training set occupies a nontrivial amount of disk space (~0.5 Teravoxels for the inhibitory model and ~0.8 Teravoxels for the excitatory model), larger sets will improve generalization, enable the use of larger image contexts and the effective tuning of more parameters (e.g., the use of the popular transformer architecture^{55,56}.) The presented method can increase the speed of manually verified trace generation. In addition, existing manually traced neurons without transcriptomic characterization (e.g., ref. 41) can still be useful in training the segmentation model. We note, however, that admitting a neuron to the training set is currently labor-intensive: it requires ensuring all branches in the image, not only those of the neuron of interest, are traced, and branches and somata are properly inflated by the volumetric label generation routine. Finally, while voxel-based loss functions, such as the one used in our model, are easier and faster to train, single voxel mistakes can change the connectivity due to the filamentous appearance of the arbor under the light microscope. Therefore, topology-aware objective functions^{57,58} can improve the topological accuracy of the segmentations, a relative weakness of the proposed model. If perfect segmentation is required, we expect a human expert to remain in the loop in the near future, primarily to verify the accuracy of the branching points.

Methods

Dataset

The dataset profiling local morphology and transcriptome of GABAergic mouse visual cortical neurons was generated as part of the

Patch-seq recordings described in ref. 5. This dataset includes 2,341 cells with transcriptomic profiling and high-resolution image stacks, where the brain sections were imaged on an upright AxioImager Z2 microscope (Zeiss) equipped with an Axiocam506 monochrome camera. Tiled image stacks of individual cells were acquired with a 63× objective lens (Zeiss Plan-Apochromat 63×/1.4 NA or Zeiss LD LCI Plan-Apochromat 63×/1.2 NA) at an interval of 0.28 μm or 0.44 μm along the Z axis. Individual cells were manually placed in the appropriate cortical region and layer within the Allen Mouse Common Coordinate Framework (CCF)^{59,60} by matching the 20× image of the slice with a “virtual” slice at an appropriate location and orientation within the CCF. 1,259 cells were removed from the dataset either because they were mapped to nearby regions (instead of visual cortex) or because their images had incomplete axons. 543 of the remaining 1,082 bright-field image stacks of biocytin-filled neurons were reconstructed both manually and automatically, and this set is used for training and testing of the t-type classification algorithms and for error/ R^2 -value quantification. The remaining 539 cells were reconstructed only automatically. To ensure the quality of scientific results presented in Table 1 and Fig. 4, we excluded 118 images with multiple neurons in the field of view and 151 images that failed at different stages of the pipeline (missing pia/white matter annotations, annotation-related failed upright transformation, reconstruction failing visual inspection) from those analyses. Finally, 16 cells in the manually and automatically reconstructed population and 20 cells in the automatically-only reconstructed population were not used for analyses involving t-types because these cells were deemed to not have “highly consistent” t-type mappings in ref. 5.

Volumetric training data generation from skeletonized morphologies

Segmentation of neuronal morphologies from 3D image stack requires voxel-wise labels while manual reconstructions specified by traces only provide a set of vertices and edges corresponding to the layout of the underlying morphology. We developed a topology-preserving fast-marching algorithm to generate volumetric ground truth using raw image stacks and manual traces by adapting a fast-marching-based segmentation algorithm^{13,61} initialized with trace vertices to segment image voxels. This segmentation should be consistent with the layout of morphology traces, without introducing topological errors (e.g., splits, merges, holes). We ensured this by incorporating simple point methods in digital topology⁶² into the segmentation algorithm. (i.e., the proposal generated by fast marching is disallowed if the proposed voxel value flip changes the underlying topology.) We noticed that the soma can be incompletely filled by the fast marching algorithm. Therefore, we treated the soma region separately and used a sequence of erosion and dilation operations followed by manual thresholding to achieve complete labeling. Each voxel was labeled as axon, dendrite, soma, background.

Neural network architecture and training

We used a 3D U-Net convolutional neural network²⁸ to perform multi-class segmentation (i.e. each voxel is assigned a probability of belonging to classes specified in the label set). We trained two separate models using raw images and volumetric ground truth; one with 51 inhibitory neurons, and another with 75 excitatory neurons from mouse cortex. The weights of the excitatory model were initialized with those of the trained inhibitory model, except for the classification layer. The U-Net architecture consists of a contracting path to capture context and a symmetric expanding path that enables precise localization and has been shown to achieve state-of-the-art performance in many segmentation tasks^{27,28}. Building on previous work¹⁶, we developed an efficient Pytorch implementation that runs on graphical processing units (GPUs). To address GPU memory constraints, during each training epoch the training set was divided randomly into subsets

of 3 stacks. Training proceeded sequentially using all subsets in an epoch, and data loading time did not exceed 10% of the total training time. We trained the model on batches of 3D patches ($128 \times 128 \times 32$ px³, XYZ), which were randomly sampled from the subset of 3 stacks loaded into memory. All the models were trained using the Adam optimizer⁶³ with a learning rate of 0.1. Training with a GeForce GTX 1080 GPU for ~50 epochs took ~3 weeks. Since the neuron occupies a small fraction of the image volume, we chose patches that contained at least one voxel belonging to the neuron. To add salient, negative examples to the training set, we also included a number of patches with bright backgrounds produced by staining artifacts and pial surface. To enable the model to generalize from relatively small number of training examples and improve the segmentation accuracy, we augmented training data by 90° rotations and vertical/horizontal flips in the image (XY) plane.

End-to-end neuron reconstruction pipeline

An end-to-end automated pipeline combined segmentation of raw image stacks into soma, axon, and dendrite structures with post-processing routines to produce a swc file. Segmentation with trained models using a single high-end GPU takes ~5.8 min per Gvoxel, or ~186 min for average 32 Gvoxel image stack. Our pipeline had access to 16 GPU (NVIDIA Titan X) cards. Even though the model trained using inhibitory neurons generalized well on all types of neurons, we found that the model trained using excitatory neurons improved the axon/dendrite assignment accuracy on excitatory neurons. Segmentation was post-processed by thresholding the background, followed by connected component analysis to remove short segments, skeletonization, and converting to a digital reconstruction in the swc format.

Axon/dendrite relabeling

Since the initial segmentation by the UNet-based neural network assigns every foreground voxel one of {soma, axon, dendrite} based on local context defined by the patch size, it is prone to occasional errors, particularly in distinguishing between axon and dendrite. These errors propagate to the labeling of nodes in the tree representation. In order to improve this initial node labeling, we developed an error correction approach which utilizes the initial segmentation to make a decision based on a larger spatial context.

We trained a secondary neural network model to predict the labels based on features calculated using the raw image stack and the initial trace. First, for each connected component in the skeleton of the initial segmentation, we identified the longest path from the node closest to soma and calculated features on the node set defining that path: (i) The first 6 features are 1D arrays of image brightness values calculated for every node in the set at different spatial scales using spherical kernels of varying radii (1, 2, 4, 8, 16, 32). (ii) A second set of 6 features are 1D arrays of neurite tortuosity values calculated for every node in the set as ratios of path and Euclidean distances between each node and its n -th neighbor away from the soma ($n = 4, 8, 16, 32, 64, 128$). (iii) Two additional features are 1D arrays of node type of every node in the set and a single number representing distance of the closest node to the soma. For each 1D array, we used the first 2048 nodes and zero-padded shorter arrays to have a uniform array size.

The neural network architecture consists of three arms, two arms have two convolutional layers with 4 and $8 \times 7 \times 3$ filters followed by 4×1 max pooling, a fully connected layer and a dropout layer each. These arms process feature sets (i) and (ii) above by stacking the sets of 6 1D arrays along a second dimension. The third arm processes feature set (iii) and has two convolutional layers with 4 and $8 \times 7 \times 1$ filters followed by 4×1 max pooling, a fully connected layer and a dropout layer. The outputs of these three arms are concatenated with the 'distance to soma feature'. Finally, the concatenated hidden feature map is processed by two fully connected layers to produce a single scalar softmax

output indicating the inferred label type. The network model was trained using examples from the training dataset of the semantic segmentation model, and is used to relabel the neuron traces during postprocessing. Averaging predicted label type over the 9 longest paths improved relabeling accuracy for connected components longer than 2048 nodes.

Connecting disconnected segments

Due to a combination of staining artifacts and limitations of bright-field imaging, the initial reconstruction is often characterized by multiple disconnected subtrees. We introduced artificial breaks to manual traces to train a random forest classifier to predict whether nearby pairs of connected components should be merged. First, we find all pairs of segment (subtree) end points that are located within a certain distance. Next, for each pair of end points, we calculate the Euclidean distance and four collinearity values which measure the segments' orientation relative to each other. Specifically, for each end point, we calculate two vectors representing segment end orientation at two different spatial scales (i.e., the orientation of the branch terminating at that end point). The collinearity values are the dot products of each of these vectors with the vector between end points of the pair. Finally, for each end point in the pair we also calculate the above features for the closest four other end points. As a result, for every pair of end nodes we have a total of 45 features. Only segments of the same type, axon or dendrite, are considered for merging.

Additional postprocessing

We passed the reconstructions through a series of additional post-processing for extraneous cell/artifact removal, down sampling, node sorting and pruning. We used quality control by a human to check for the presence of disconnected branches of extraneous cells that were not removed during postprocessing. We excluded samples that did not pass this quality control if they did not have a manual trace. For samples with a manual trace, we used a neighborhood of the manual trace to simulate manual removal of extraneous cells by masking with that neighborhood. We excluded these samples from reconstruction accuracy quantification and used them only for cell type classification. For excitatory neurons (Fig. 2a, and ref. 36), we trained a random forest classifier to identify apical dendrite segments in excitatory auto-trace reconstructions. The classifier was trained on geometric features that distinguish apical dendrite segments from the basal dendrite (e.g. upright distance from soma). This classifier achieved a mean accuracy of 85% percent across 10-fold cross validation.

Morphometric feature calculation

Reconstructions were transformed to an upright position that is perpendicular with respect to pia and white matter surfaces. Morphometric features were calculated using the skeleton_keys python package⁴¹. Following ref. 41, z-dimension features were not included.

Arbor density generation of axons and dendrites

We represented axonal and dendritic morphologies as 2D density maps registered to a common local coordinate axis using the pia/white matter boundaries. First, we applied upright transform to the reconstructed neuron followed by the correction of z-shrinkage and the variable tilt angle of the acute slice⁴¹. Next, adapting our previous work^{12,40} we conformally mapped pia and white matter boundaries to flat surfaces, calculated a nonlinear transformation on the whole tissue by a least-square fit to pia/white matter mappings, and applied this transformation to the morphology trace. Finally, we used the registered trace to generate a 2D density representation. The polar axes representing the cortical depth and the lateral distance from the soma (Fig. 2c) make this representation invariant to rotations around

the laminar axis. We calculated these maps separately for axons and dendrites. We downsampled these maps to $120 \text{ px} \times 4 \text{ px}$ images with a pixel size of $8 \mu\text{m} \times 125 \mu\text{m}$ to be robust to minor changes in morphology. We normalized the intensity by the lateral area corresponding to each pixel so that each pixel value represents local arbor density.

Assessing neuron reconstruction accuracy

To assess the quality of automated neuron reconstructions, we used manual reconstructions as the ground truth. We quantified the correspondence of trace nodes, as described in the main text, to evaluate the accuracy of the trace layout. We calculated precision, recall and f1-score metrics at three distances (2, 5, and $10 \mu\text{m}$), and reported mean and standard deviation values for the test set of 340 neurons (all samples that have manual traces excluding the ones used for training models or required masking). In addition, we evaluated the accuracy of neuron morphology representations, morphometric features and ADRs, for the same set. After realizing that the image coordinates used for manual vs. automated tracing were nonlinearly warped with respect to each other for one cell (penultimate cell in Fig. S10), we removed that cell from the single cell-level node correspondence, ADR, and feature-based comparison studies. (It was used in all cell typing studies.) We reported the Pearson's correlation coefficient r for each morphometric feature. We calculated the average root-mean-squared error per t-type between normalized axon/dendrite ADRs derived from automated and manual reconstructions.

Supervised classification

Supervised classification using morphometric features was performed by training a random forest classifier implemented in the scikit-learn Python package⁶⁴ using 5-fold cross-validation. This was repeated 20 times. Supervised classification using ADRs was done by training a feed-forward neural network classifier using our Pytorch implementation. The network architecture consists of a convolutional layer with $7 \times 3 \times 1$ filters, a $4 \times 1 \times 1$ max pooling layer, a convolutional layer with $7 \times 3 \times 1$ filters, a $3 \times 1 \times 1$ max pooling layer, a layer that concatenates the hidden features with the soma depth value, and a fully connected layer with the number of units corresponding to the number of classes. Each convolutional layer uses the rectified linear function as the non-linear transformation.

Since the depth locations of neurons vary within the cortex, we introduced a data augmentation scheme based on simulation of cell type-dependent neuronal shift along the laminar axis. Namely, for each t-type we calculated the range of soma depth variations, and applied a random shift within that range to the input soma depth value, as well as the corresponding shift to the ADR intensity in the laminar direction. This cell type-dependent random shift of the ADR and the soma depth together with a modulation of the intensity values of the ADR improved the accuracy of classification.

The networks were trained using the cross-entropy loss function and the Adam optimizer with a learning rate of 0.001. Training using 10-fold cross-validation with GeForce GTX 1080 GPU for 50,000 epochs took ~24 h. A set of 246 automatically and 501 manually reconstructed cells was used for training classifiers shown in Fig. 3a and b, and a set of the same 488 automatically and manually reconstructed cells were used for Fig. 3c, d and Fig. S19. Both sets included only cells from t-types with at least 5 cells. Confusion matrices, mean and standard deviation of accuracy across cross-validation folds were reported.

We performed Boschloo's exact test on 2×2 contingency tables where columns/rows store total numbers of correct and incorrect predictions for two given classifiers. When comparing ADR-based to morphometric feature-based classifiers, we calculated the contingency table for each of the 20 repetitions used in the feature-

based classifier study. We calculated the p -value of one-sided Boschloo's test to evaluate the null hypothesis of ADR-based accuracy being less than feature-based accuracy. To aggregate the 20 p -values, we used ref. 42 and the Python implementation at <https://github.com/benjaminpatrickevans/harmonicmean> to report the p -value of the asymptotically exact harmonic mean p -value test for t-type and subclass predictions.

Cophenetic agreement. To take the hierarchical organization of transcriptomically defined mouse cortical cell types²³ into account when evaluating the accuracies of the different classification tasks, we defined resolution index per cell as the scaled height of the closest common ancestor of assigned and predicted labels in the t-type hierarchical tree⁶⁵ (Figs. S20, S21). Accordingly, the resolution index for a correctly classified t-type ("leaf node" label) is 1. In the worst case, the closest ancestor for an assigned and predicted label can be the root node of the taxonomy, which corresponds to a resolution index of 0. We report the mean and s.e.m. values for this measure of cophenetic agreement between true and predicted assignments for each cell type in Figs. S20, S21.

Sparse feature selection analysis

Following ref. 66, a set of 1,252 genes were used for this analysis. This set was obtained by excluding genes if they satisfy any of the following criteria: (1) they are highly expressed in non-neuronal cells, (2) they have previously reported sex or mitochondrial associations, and (3) they are much more highly expressed in Patch-seq data vs. Fluorescence Activated Cell Sorting (FACS) data (or vice versa) and therefore may be associated with the experimental platform⁵. Further, we removed gene models and some other families of unannotated genes that may be difficult to interpret. We also used the β score, a published measure to evaluate the degree to which gene expression is exclusive to t-types⁶⁷, to exclude genes expressed broadly across t-types. Gene expression values were CPM normalized, and then $\log_e(\bullet + 1)$ transformed for all the downstream analyses.

A set of 777 neurons was used for the feature selection analysis where automated reconstructions comprise ~32% of this set (~44% of the subset of 7 t-types studied in Fig. 4). Every neuron in the dataset was characterized by the expression levels of the set of 1252 genes, and their axonal and dendritic 1D arbor density representations were organized into two 120×1 vectors. For each neuron, axon/dendrite layer-specific skeletal masses normalized by total skeletal mass were calculated to quantify layer-specific innervation for axonal and dendritic morphologies, and axon/dendrite centroids were calculated to characterize laminar position of the morphology. To select a small subset of genes that are responsible for the variability in individual anatomical features within each transcriptomic type or subclass, we solved the Lasso regression problem⁶⁸ (LassoCV command in the scikit-learn library⁶⁴) for each anatomical feature using the cells in that transcriptomic set. We analyzed only the sets corresponding to the types and subclasses with at least 20 cells. Briefly, let K_t denote the number of cells of type or subclass t for which we have both anatomical features y_t (a $K_t \times 1$ vector) and gene expression values X_t (a $K_t \times N$ matrix). We solve

$$\text{minimize } \frac{1}{2K_t} \|y_t - X_t w_t\|_2^2 + \alpha \|w_t\|_1 \quad (1)$$

by performing nested 5-fold cross-validation. For each cross-validation fold, we passed the training set into LassoCV which performed another splitting of the data to determine the hyperparameter α and the set of selected genes. We selected the 10 genes with maximum absolute weight values and calculated the coefficient of determination, R^2 , for the test set. Finally, we selected the 10 most frequent genes across the 5 folds and calculated the mean test R^2 value. To evaluate statistical

significance, we shuffled the rows of the gene expression matrix X_t 1000 times and used the same procedure to calculate mean test R^2 value for each shuffled run. We calculated the one-sided p -value as the fraction of shuffled runs with R^2 values greater than or equal to the true R^2 . Finally, we performed multiple testing correction of p -values using Benjamini-Yekutieli method⁶⁹ to control the false discovery rate (multipletests command in the statsmodels library⁷⁰). We report resulting p -values and test R^2 values for t-types in Table 1 and for t-types and subclasses in Tables S3 and S4.

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

Transcriptomic and morphological data supporting the findings of this study is available online at <https://portal.brain-map.org/explore/classes/multimodal-characterization> ("Neurons in Mouse Primary Visual Cortex"). Additional dataset of automated morphological reconstructions is available at <https://github.com/ogliko/patchseq-autorecon>. Source data are provided with this paper as a Source Data file. Source data are provided with this paper.

Code availability

Code pertaining to this study as well as the trained neural network model for automated segmentation are available at <https://github.com/ogliko/patchseq-autorecon>⁷¹ and <https://github.com/rhngla/topo-preserve-fastmarching>. Morphometric features are calculated using the skeleton_keys package at <https://skeleton-keys.readthedocs.io>.

References

- Lipovsek, M. et al. Patch-seq: Past, present, and future. *J. Neurosci.* **41**, 937–946 (2021).
- Fuzik, J. ános et al. Integration of electrophysiological recordings with single-cell rna-seq data identifies neuronal subtypes. *Nat. Biotechnol.* **34**, 175 (2016).
- Cadwell, C. R. et al. Electrophysiological, transcriptomic and morphologic profiling of single neurons using patch-seq. *Nat. Biotechnol.* **34**, 199–203 (2016).
- Földy, C. et al. Single-cell rnaseq reveals cell adhesion molecule profiles in electrophysiologically defined neurons. *Proc. Natl Acad. Sci.* **113**, E5222–E5231 (2016).
- Gouwens, N. W. et al. Integrated morphoelectric and transcriptomic classification of cortical gabaergic cells. *Cell* **183**, 935–953 (2020).
- Scala, F. et al. Phenotypic variation of transcriptomic cell types in mouse motor cortex. *Nature* **598**, 144–150 (2021).
- Lee, B. R. et al. Scaled, high fidelity electrophysiological, morphological, and transcriptomic cell characterization. *Elife* **10**, e65482 (2021).
- Que, L., Lukacsovich, D., Luo, W. & Földy, C. Transcriptional and morphological profiling of parvalbumin interneuron subpopulations in the mouse hippocampus. *Nat. Commun.* **12**, 1–15 (2021).
- Marx, V. Patch-seq takes neuroscience to a multimodal place. *Nat. Methods* **19**, 1340–1344 (2022).
- Donohue, D. E. & Ascoli, G. A. Automated reconstruction of neuronal morphology: an overview. *Brain Res. Rev.* **67**, 94–102 (2011).
- Türetken, E., Benmansour, F. & Fua, P. Automated reconstruction of tree structures using path classifiers and mixed integer programming. In: *2012 IEEE Conference on Computer Vision and Pattern Recognition* (eds Belongie, S., Blake, A., Luo, J. & yuille, A.) 566–573 (IEEE, 2012).
- Sümbül, U. et al. A genetic and computational approach to structurally classify neuronal types. *Nat. Commun.* **5**, 1–12 (2014).
- Gala, R., Chapeton, J., Jitesh, J., Bhavsar, C. & Stepanyants, A. Active learning of neuron morphology for accurate automated tracing of neurites. *Front. Neuroanatomy* **8**, 37 (2014).
- Feng, L., Zhao, T. & Kim, J. neutube 1.0: a new design for efficient neuron reconstruction software based on the swc format. *eneuro* **2** (2015).
- Quan, T. et al. Neurogps-tree: automatic reconstruction of large-scale neuronal populations with dense neurites. *Nat. Methods* **13**, 51–54 (2016).
- Gornet, J. et al. Reconstructing neuronal anatomy from whole-brain images. In: *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)* (eds Liebling, M. & Greenspan, H.) 218–222 (IEEE, 2019).
- Winnubst, J. et al. Reconstruction of 1,000 projection neurons reveals new cell types and organization of long-range connectivity in the mouse brain. *Cell* **179**, 268–281 (2019).
- Peng, H. et al. Morphological diversity of single neurons in molecularly defined cell types. *Nature* **598**, 174–181 (2021).
- He, W. et al. Automated three-dimensional tracing of neurons in confocal and brightfield images. *Microsc. Microanal.* **9**, 296–310 (2003).
- Peng, H. et al. Automatic tracing of ultra-volumes of neuronal images. *Nat. Methods* **14**, 332–333 (2017).
- Zhou, Z., Kuo, Hsien-Chi, Peng, H. & Long, F. Deepneuron: an open deep learning toolbox for neuron tracing. *Brain Inform.* **5**, 1–9 (2018).
- Jin, D. Z. et al. Shutu: open-source software for efficient and accurate reconstruction of dendritic morphology. *Front. Neuroinform.* **13**, 68 (2019).
- Tasic, B. et al. Shared and distinct transcriptomic cell types across neocortical areas. *Nature* **563**, 72–78 (2018).
- Markram, H. et al. Reconstruction and simulation of neocortical microcircuitry. *Cell* **163**, 456–492 (2015).
- Murayama, M. et al. Dendritic encoding of sensory stimuli controlled by deep cortical interneurons. *Nature* **457**, 1137 (2009).
- Schuman, B., Dellal, S., Prönneke, A., Machold, R. & Rudy, B. Neocortical layer 1: an elegant solution to top-down and bottom-up integration. *Annu. Rev. Neurosci.* **44**, 221–252 (2021).
- Ronneberger, O., Fischer, P. & Brox, T. U-net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical Image Computing And Computer-assisted Intervention* (eds Navab, N., Hornegger, J., Wells, W. M. & Frangi, A. F.) 234–241 (Springer, 2015).
- Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T. & Ronneberger, O. 3d u-net: learning dense volumetric segmentation from sparse annotation. In: *International Conference on Medical Image Computing and Computer-assisted Intervention* (eds Ourselin, S., Joskowicz, L., Sabuncu, M. R., Uanl, G. & Wells, W.) 424–432 (Springer, 2016).
- Funke, J. et al. Large scale image segmentation with structured loss based deep learning for connectome reconstruction. *IEEE Trans. Pattern Anal. Mach. Intel.* **41**, 1669–1680 (2018).
- Xiao, H. & Peng, H. App2: automatic tracing of 3d neuron morphology based on hierarchical pruning of a gray-weighted image distance-tree. *Bioinformatics* **29**, 1448–1454 (2013).
- Manubens-Gil, L. et al. Bigneuron: a resource to benchmark and predict performance of algorithms for automated tracing of neurons in light microscopy datasets. *Nat. Methods* **20**, 824–835 (2023).
- Chartrand, T. et al. Morphoelectric and transcriptomic divergence of the layer 1 interneuron repertoire in human versus mouse neocortex. *Science* **382**, eadf0805 (2023).
- Brown, K. M. et al. The diadem data sets: representative light microscopy images of neuronal morphology to advance automation of digital reconstructions. *Neuroinformatics* **9**, 143–157 (2011).

34. Martinez, L. M. et al. Receptive field structure varies with layer in the primary visual cortex. *Nat. Neurosci.* **8**, 372–379 (2005).
35. Ascoli, G. A. Mobilizing the base of neuroscience data: the case of neuronal morphologies. *Nat. Rev. Neurosci.* **7**, 318–324 (2006).
36. Sorensen, S. A. et al. Connecting single neuron transcriptomes to the projectome in mouse visual cortex. *bioRxiv* <https://doi.org/10.1101/2023.11.25.568393> (2023).
37. Stepanyants, A., Hof, P. R. & Chklovskii, D. B. Geometry and structural plasticity of synaptic connectivity. *Neuron* **34**, 275–288 (2002).
38. Cuntz, H., Forstner, F., Borst, A. & Häusser, M. One rule to grow them all: a general theory of neuronal branching and its practical application. *PLoS Comput. Biol.* **6**, e1000877 (2010).
39. Oberlaender, M. et al. Cell type-specific three-dimensional structure of thalamocortical circuits in a column of rat vibrissa cortex. *Cerebral Cortex* **22**, 2375–2391 (2012).
40. Sümbül, U., Zlateski, A., Vishwanathan, A., Masland, R. H. & Seung, H. S. Automated computation of arbor densities: a step toward identifying neuronal cell types. *Front. Neuroanatomy* **8**, 139 (2014).
41. Gouwens, N. W. et al. Classification of electrophysiological and morphological neuron types in the mouse visual cortex. *Nat. Neurosci.* **22**, 1182–1195 (2019).
42. Wilson, D. J. The harmonic mean p-value for combining dependent tests. *Proc. Natl Acad. Sci.* **116**, 1195–1200 (2019).
43. Laturnus, S., Kobak, D. & Berens, P. A systematic evaluation of interneuron morphology representations for cell type discrimination. *Neuroinformatics* **18**, 591–609 (2020).
44. Yuan, M. & Lin, Y. Model selection and estimation in regression with grouped variables. *J. Roy. Stat. Soc.: Ser. B (Stat. Methodol.)* **68**, 49–67 (2006).
45. Kobak, D. et al. Sparse reduced-rank regression for exploratory visualisation of paired multivariate data. *J. Roy. Stat. Soc.: Ser. C (Appl. Stat.)* **70**, 980–1000 (2021).
46. Favuzzi, E. et al. Distinct molecular programs regulate synapse specificity in cortical inhibitory circuits. *Science* **363**, 413–417 (2019).
47. Marghi, Y., Gala, R., Baftizadeh, F. & Sumbul, U. Joint inference of discrete cell types and continuous type-specific variability in single-cell datasets with mmidas. *bioRxiv* <https://doi.org/10.1101/2023.10.02.560574> (2023).
48. Inoue, A. & Sanes, J. R. Lamina-specific connectivity in the brain: regulation by n-cadherin, neurotrophins, and glycoconjugates. *Science* **276**, 1428–1431 (1997).
49. Kolodkin, A. L. & Tessier-Lavigne, M. Mechanisms and molecules of neuronal wiring: a primer. *Cold Spring Harbor Perspectives Biol.* **3**, a001727 (2011).
50. Sun, Y.-C. et al. Integrating barcoded neuroanatomy with spatial transcriptional profiling enables identification of gene correlates of projections. *Nat. Neurosci.* **24**, 873–885 (2021).
51. Gerstner, W., Kistler, W. M., Naud, R. & Paninski, L. *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition* (Cambridge University Press, 2014).
52. Harris, K. D. et al. Classes and continua of hippocampal ca1 inhibitory neurons revealed by single-cell transcriptomics. *PLoS Biol.* **16**, e2006387 (2018).
53. Stanley, G., Gokce, O., Malenka, R. C., Südhof, T. C. & Quake, S. R. Continuous and discrete neuron types of the adult murine striatum. *Neuron* **105**, 688–699 (2020).
54. O’Leary, T. P. et al. Extensive and spatially variable within-cell-type heterogeneity across the basolateral amygdala. *Elife* **9**, e59003 (2020).
55. Dosovitskiy, A. et al. An image is worth 16 × 16 words: transformers for image recognition at scale. In: *International Conference on Learning Representations* (eds Mohamed, S., White, M., Cho, K. & Song, D.) (OpenReview.net, 2020).
56. Liu, Z. et al. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision 10012–10022* (eds Damen, D., Hassner, T., Pal, C. & Sato, Y.) (IEEE, 2021).
57. Hu, X., Samaras, D. & Chen, C. Learning probabilistic topological representations using discrete morse theory. In: *The Eleventh International Conference on Learning Representations* (Liu, Y., Finn, C., Choi, Y. & Deisenroth, M.) (OpenReview.net, 2022).
58. Grim, A., Chandrashekar, J., Svoboda, K. & Sümbül, U. Instance segmentation with supervoxel based topological loss function Preprint at <https://openreview.net/forum?id=NhLBhx5BVY> (2014).
59. Kuan, L. et al. Neuroinformatics of the allen mouse brain connectivity atlas. *Methods* **73**, 4–17 (2015).
60. Wang, Q. et al. The allen mouse brain common coordinate framework: a 3d reference atlas. *Cell* **181**, 936–953 (2020).
61. Sethian, J. A. Fast marching methods. *SIAM Rev.* **41**, 199–235 (1999).
62. Bertrand, G. & Malandain, Grégoire A new characterization of three-dimensional simple points. *Pattern Recog. Lett.* **15**, 169–175 (1994).
63. Kingma, D. P., & Ba, J. Adam: A method for stochastic optimization. In: *International Conference on Learning Representations* (eds Kingsbury, B., Bengio, S., Freitas, de N. & Larochelle, H.) (2015).
64. Pedregosa, F. et al. Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).
65. Smith, S. J. et al. Single-cell transcriptomic evidence for dense intracortical neuropeptide networks. *Elife* **8**, e47889 (2019).
66. Gala, R. et al. Consistent cross-modal identification of cortical neurons with coupled autoencoders. *Nat. Comput. Sci.* **1**, 120–127 (2021).
67. Bakken, T. E. et al. Single-nucleus and single-cell transcriptomes compared in matched cortical cell types. *PLoS ONE* **13**, e0209648 (2018).
68. Tibshirani, R. Regression shrinkage and selection via the lasso. *J. Roy. Stat. Soc.: Ser. B (Methodological)* **58**, 267–288 (1996).
69. Benjamini, Y. & Yekutieli, D. The control of the false discovery rate in multiple testing under dependency. *Ann. Stat.* 1165–1188 (2001).
70. Seabold, S. & Perktold, J. statsmodels: Econometric and statistical modeling with python. In: *9th Python in Science Conference* (Walt, van der S.) (SciPy.org, 2010).
71. Gliko, O. et al. High-throughput analysis of dendritic and axonal arbors reveals transcriptomic correlates of neuroanatomy. <https://doi.org/10.5281/zenodo.12215416> (2024).

Acknowledgements

The authors wish to thank the Allen Institute for Brain Science founder, Paul G Allen, for his vision, encouragement and support. We thank Michael Hawrylycz and Claire Gamlin for helpful feedback on the manuscript. This work was supported by the National Institute Of Mental Health of the National Institutes of Health under award numbers 1U01MH114824-01 and 1RF1MH128778-01. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

Author contributions

Conceptualization: O.G., H.Z., S.A.S., U.S.; Research design: O.G., M.M., R.D., S.A.S., U.S.; Perform research: O.G., M.M., R.D., R.G., J.G., S.A.S., U.S.; Data analysis: O.G., M.M.; Writing original draft: O.G., M.M., R.D., R.G., S.A.S., U.S.; Editing manuscript: O.G., M.M., R.D., R.G., S.A.S., U.S.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41467-024-50728-9>.

Correspondence and requests for materials should be addressed to Olga Gliko or Uygur Sümbül.

Peer review information *Nature Communications* thanks Shreejoy Tripathy and the other anonymous reviewer(s) for their contribution to the peer review of this work. A peer review file is available.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2024