

1 Data-driven fine-grained region discovery in the 2 mouse brain with transformers

3
4 Alex J. Lee^{1,2}, Alma Dubuc^{1,2}, Michael Kunst³, Shenqin Yao³, Nicholas Lusk³, Lydia Ng³,
5 Hongkui Zeng³, Bosiljka Tasic³, Reza Abbasi-Asl^{1,2,*}

6
7 ¹ University of California, San Francisco
8

9 ² UCSF Weill Institute for Neurosciences

10 ³ Allen Institute for Brain Science

11 *Corresponding author: Reza Abbasi-Asl (Reza.AbbasiAsl@ucsf.edu)

12 Abstract

13 Spatial transcriptomics offers unique opportunities to define the spatial organization of tissues and
14 organs, such as the mouse brain. We address a key bottleneck in the analysis of organ-scale spatial
15 transcriptomic data by establishing a workflow for self-supervised spatial domain detection that is scalable
16 to multimillion cell datasets. This workflow uses a self-supervised framework for learning latent
17 representations of tissue spatial domains or niches. We use a novel encoder-decoder architecture, which
18 we named CellTransformer, to hierarchically learn higher-order tissue features from lower-level cellular
19 and molecular statistical patterns. Coupling our representation learning workflow with minibatched GPU-
20 accelerated clustering algorithms allows us to scale to multi-million cell MERFISH datasets where other
21 methods cannot. CellTransformer is effective at integrating cells across tissue sections, identifying
22 domains highly similar to ones in existing ontologies such as Allen Mouse Brain Common Coordinate
23 Framework (CCF) while allowing discovery of hundreds of uncataloged areas with minimal loss of domain
24 spatial coherence. CellTransformer domains recapitulate previous neuroanatomical studies of areas in
25 the subiculum and superior colliculus, and characterize putatively uncataloged subregions in subcortical
26 areas which currently lack subregion annotation. CellTransformer is also capable of domain discovery in
27 whole-brain Slide-seqV2 datasets. Our workflows enable complex multi-animal analyses, achieving nearly
28 perfect consistency of up to 100 spatial domains in a dataset of four individual mice with nine million cells
29 across more than 200 tissue sections. CellTransformer advances the state of the art for spatial
30 transcriptomics, by providing a performant solution for detection of fine-grained tissue domains from
31 spatial transcriptomics data.

32 Introduction

33 Hierarchical spatial organization is ubiquitous in tissue and organ biology. Systematic, high-
34 dimensional phenotypic measurements of this organization, generated through experimental tools such
35 as spatial transcriptomics, multiplex immunofluorescence, and electron microscopy, are also becoming
36 increasingly available as large, open datasets. However, transforming this abundance of data into a
37 useful representation can be difficult, even for fields with a wealth of prior knowledge, such as
38 neuroanatomy.

39 Datasets such as the Allen Brain Cell Mouse Whole Brain (ABC-MWB) Atlas¹⁻³, a multi-million
40 cell single-cell RNA sequencing (scRNA-seq) and spatial (MERFISH) atlas, provide unprecedented
41 opportunities to investigate whether computational tools can help biologists understand spatial cellular
42 and molecular organization. However, the size of these datasets presents computational challenges for
43 existing methods. Existing methods for spatial niche or spatial domain detection often operate on the
44 entire dataset at once, for example a tissue-section-wide cell by gene matrix. This precludes scale-up to
45 large multi-section datasets as most systems do not have the GPU memory required to load multiple
46 sections of data or store intermediary representations such as pairwise distance matrices⁴⁻⁶, particularly
47 as datasets scale into the millions or tens of millions. Some methods rely on Gaussian processes, which
48 feature a costly cubic computational scaling in the number of observations⁷. Other more scalable methods
49 are limited in capturing granular structure, integration across tissue sections, or require significant
50 neuroanatomical prior knowledge to manually audit, cluster, and hyperparameter tune for domain
51 discovery workflows^{8,9}.

52 Our method, CellTransformer, implements a robust representation learning and clustering
53 workflow to discover spatial niches at scale by representing not tissue sections but subgraphs that
54 represent individual cellular neighborhoods. We describe an innovative strategy to induce the encoder of
55 an encoder-decoder transformer to aggregate useful information into a neighborhood representation
56 token. This occurs by training the model to condition cell-type specific gene expression predictions using
57 this neighborhood context token. The model thus learns to predict expression of cell types in arbitrary cell
58 neighborhoods. This representation allows for recovery of important anatomically plausible spatial
59 domains while remaining computationally efficient.

60 We evaluate CellTransformer on using the ABC-MWB dataset (3.9 million cells collected with a
61 500 gene MERFISH panel)¹ demonstrating its effectiveness in producing completely data-driven spatial
62 domains of the mouse brain by comparing the results to the Allen Mouse Brain Common Coordinate
63 Framework version 3 (CCFv3)¹⁰. CCF is a consensus hand-drawn 3D reference space compiled from a
64 large multimodal data corpus. Annotations feature labels at three levels of coarseness (from 25 regions at
65 coarse-grain to 670 at fine-grain), which we use to show that CellTransformer excels at identifying spatial
66 domains which are spatially coherent and biologically relevant. CellTransformer domains reproduce
67 known regional architecture observed in targeted studies of the subiculum and in the superior colliculus
68 superficial layers. Beyond the 670 regions currently annotated in ABC-MWB, we show our workflow
69 produces meaningful data-driven domains in regions which currently lack subregion annotation. As
70 examples, we focus on data-driven subdomains we define in superior colliculus and midbrain reticular
71 nucleus.

72 We also demonstrate CellTransformer's strength in integrating domains across animals,
73 leveraging a separate whole-brain dataset within ABC-MWB¹¹ comprising 6.5 million cells distributed
74 across four animals and 239 sections and with a separate gene panel with 1129 genes. We find that
75 CellTransformer produces consistent subregions across all 5 animals (1 coronal and 4 sagittal),
76 suggesting a successful integration across animals with heterogeneous measurements. Notably we also
77 find that identified domains are highly consistent across animals. To our knowledge, this work provides
78 the first demonstration that large scale data-driven discovery of domains at CCF-like resolution can be
79 based on spatial transcriptomics data. Finally, we show that our framework can perform domain detection
80 in a different spatial transcriptomics modality, Slide-seqV2, using the whole-brain dataset of cellularly
81 deconvoluted results¹².

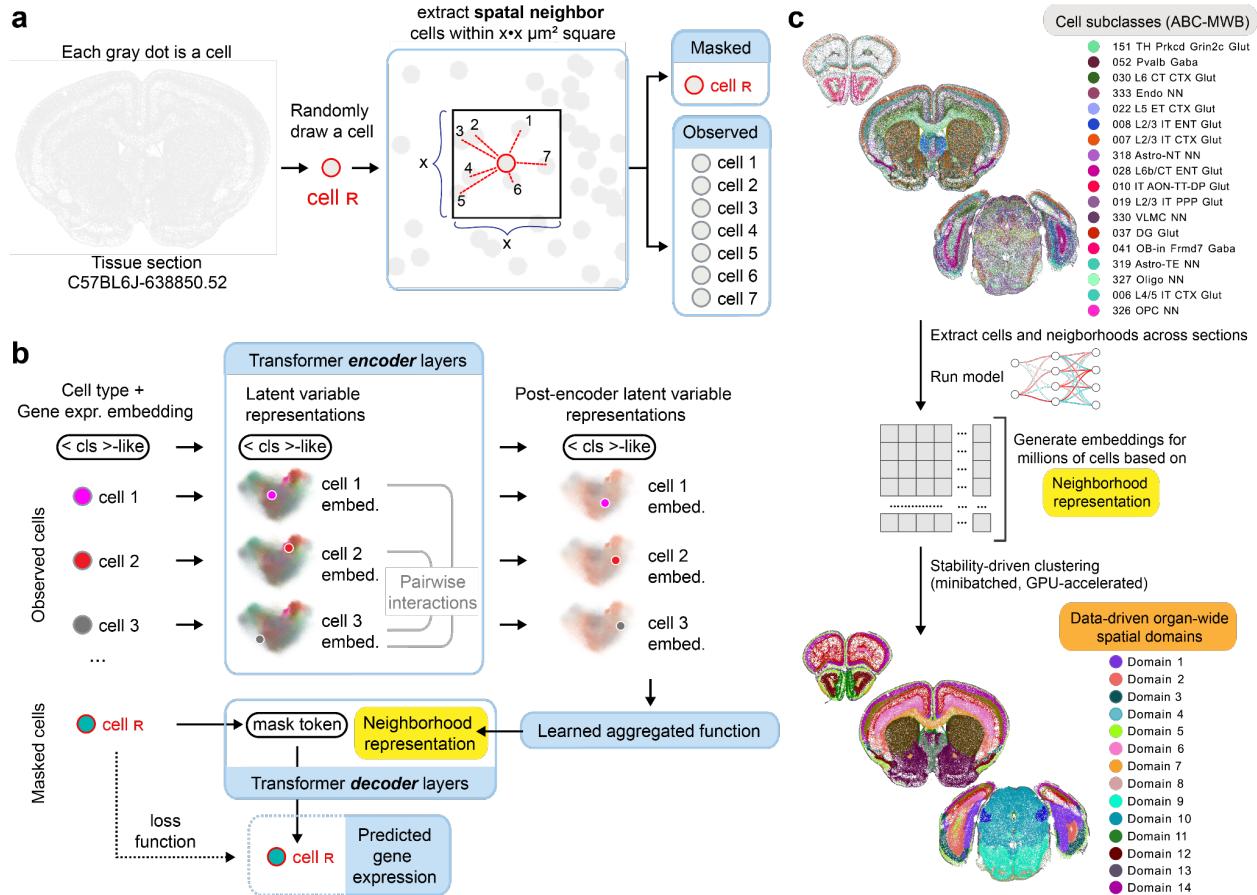
82 Results

83 The CellTransformer architecture and domain detection workflow

84 CellTransformer is a graph transformer¹³ neural network that is trained to learn latent
85 representations of cell neighborhoods by conditioning single-cell gene expression predictions on
86 neighborhood spatial context. We define a cellular neighborhood as any cells within a user-specified
87 distance cutoff in microns away from a reference or center cell. As input, our model requires the gene
88 expression profiles and cell type classifications for cells in a neighborhood and outputs a latent variable
89 representation for that neighborhood. One of the principal operations in a transformer is the self-attention
90 operation, which computes a feature update based on pairwise interactions between elements in a
91 sequence, which are referred to as tokens (here, cells). Accordingly, one interpretation of our model is of
92 learning an arbitrary and dynamic pairwise interaction graph among cells.

93 Restricting this graph to a small neighborhood subgraph of the whole-tissue-section graph has
94 benefits for both computational resource usage and biological interpretability. We interpret the size of the
95 neighborhood as a constraint on the physical distance at which statistical correlations between the
96 observed cells and their gene expression profiles can be directly captured. Truncating neighborhoods
97 using a fixed spatial threshold instead of choosing a fixed number of neighbors also allows the network to
98 account for the varying density of cells in space. Accordingly, our framework incorporates a notion of both
99 cytoarchitecture (relative density and proximity) and molecular variation (cell type and RNA-level
100 variation) in the data.

101 To induce our model to learn biologically relevant latent features from cell neighborhoods, we
102 designed a self-supervised training scheme requiring only cell-type labels, which many large-scale
103 studies make available via scRNA-seq atlas reference mapping^{1,11}. Specifically, we train the model to
104 extract features from cellular neighborhoods, modeled as sets of cell tokens that are within a box of fixed
105 size centered around a center, or reference cell, and use them to predict the observed gene expression of
106 the cell at the center of the neighborhood. We refer to this cell as the reference cell (indicated by “cell R”
107 in **Figure 1a**). Cell tokens are generated by composing cell-type and gene expression information
108 (**Methods**). After encoding with a series of transformer layers (where cells are only allowed to attend to
109 each other if they are in the same neighborhood), these tokens are then aggregated using a learned
110 pooling operation to produce a single token representation of the entire tissue context. The model
111 receives a new mask token representing the reference cell’s type which is used to predict its gene
112 expression following the operation of several transformer decoder layers (**Figure 1b**). Importantly, during
113 this process, only the mask token and the neighborhood representation can attend to each other. This
114 operation captures a hierarchical encoding and decoding process where low level information (gene and
115 cell type) is produced at the cell token level and aggregated into a high-level representation. This high-
116 level representation is then used to conduct the reverse decoding process (prediction of gene expression
117 from cell type and tissue context information). Unlike closely related method NCEM¹⁴, which predicts
118 expression of a reference masked node, we aggregate information across tokens (nodes) in a cellular
119 neighborhood using a learned pooling which strongly bottlenecks the information distributed across the
120 tokens prior to masked cell prediction.



121
122 **Figure 1.** Overall training and architectural scheme for CellTransformer. (a.) During training, a single cell is drawn
123 (we denote this the reference cell, highlighted in red). We extract the reference cell's spatial neighbors and partition
124 the group into a masked reference cell and its observed spatial neighbors. (b.) Initially, the model encoder receives
125 information about each cell and projects those features to d -dimensional latent variable space. Features interact
126 across cells (tokens) through the self-attention mechanism. These per-cell representations and an extra token acting
127 as a register token are then aggregated into a single vector representation, which we refer to as the neighborhood
128 representation. This representation is concatenated to a mask token which is cell type-specific and chosen to
129 represent the type of the reference cell. A shallow transformer decoder (dotted lines) further refines these
130 representations and then a linear projection is used to output parameters of a negative binomial distribution modeling
131 of the MERFISH probe counts for the reference cell. (c.) Once the model is trained, we compute embeddings (one for
132 each neighborhood/reference-cell pairing) and concatenate these embeddings within the tissue section datasets and
133 across tissue sections. Concatenating embeddings across tissue sections produces region discovery at organ level.
134 We then cluster these embeddings using k -means to discover tissue domains across sections.

135
136
137

138 At test time, we extract this neighborhood representation for each cell and use k -means clustering
139 to identify discrete spatial domains (**Figure 1c**). We will use the term spatial domain to refer to the output
140 of clustering on embeddings and cluster to refer to single-cell clusters transferred from the ABC-WMB
141 single cell taxonomy. We emphasize that the input embedding matrix for k -means is conducted by
142 concatenating all cells across the dataset across tissue sections. Since minibatching is used during
143 training (unlike methods such as STAligner and GraphST), for generating embeddings, and during k -
144 means (using cuml for GPU-acceleration), overall computational costs of our algorithm are limited in

145 principle only by the memory required for storage of cellular neighborhoods rather than entire sections or
146 datasets.

147 **Data-driven discovery of fine-grained spatial domains in the**
148 **mouse brain using ABC-WMB**

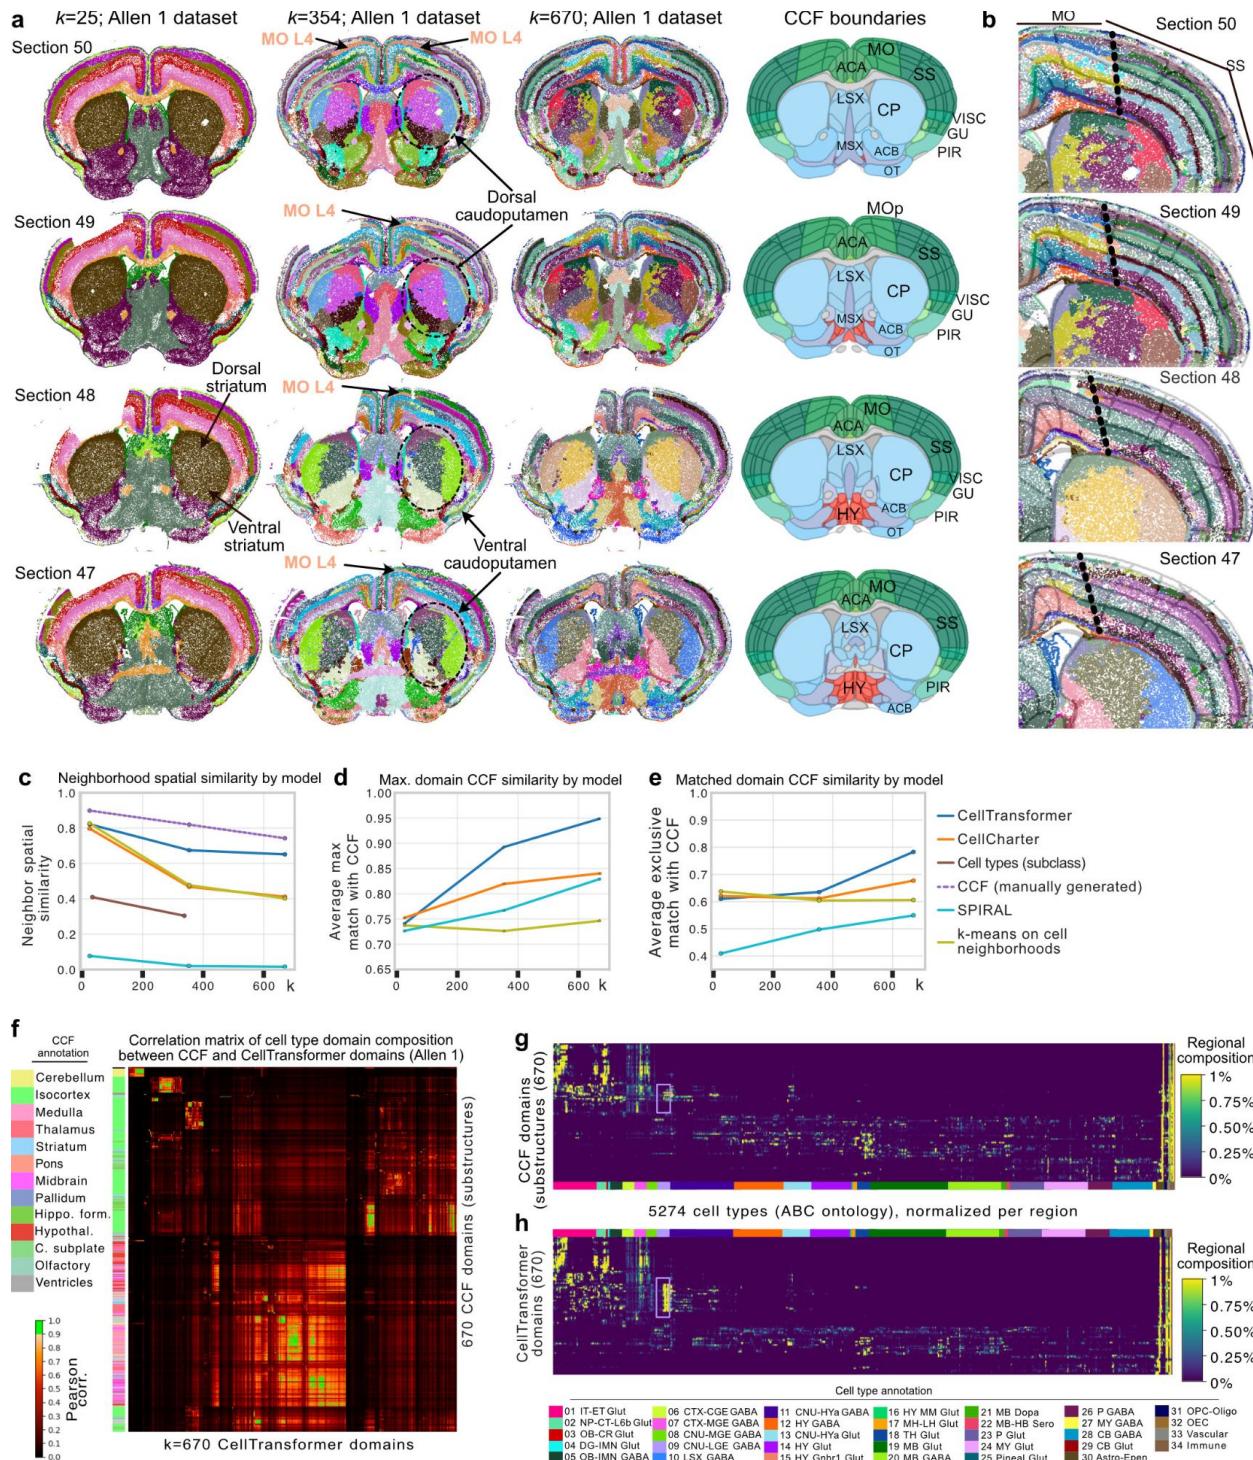
149 The ABC-WMB spatial transcriptomics dataset contains data from five mouse brains^{1,11}. One
150 animal was processed by the Allen Institute for Brain Science with a 500 gene MERFISH panel and 53
151 coronal sections (Yao et al, 2023)¹ The remaining four other animals, generated in Zhang et al. (2023)[11]
152 were collected with a 1129 gene panel. Sections from two of these animals (“Zhuang 1”, 147 sections;
153 and “Zhuang 2”, 66 sections) were sampled coronally. The other two animals in the dataset (“Zhuang 3”,
154 23 sections; and “Zhuang 4”, 3 sections) were sampled sagittally.

155 We first trained CellTransformer on the Allen 1 dataset, subsequently extracting embeddings for
156 each cell’s neighborhood, which we defined as a set of cells within a fixed size square around that cell.
157 We then clustered these embeddings using k -means. We emphasize that to generate spatial domains
158 across the brain, all k -means clustering in this paper was performed by concatenating cells in the dataset
159 across tissue sections. All further references to visualizations of domains, including those only visualized
160 for a subset of domains, were fit at a given number of domains across the entire dataset. We also
161 optionally introduced a smoothing step prior to k -means, which we applied to spatially smooth the
162 embeddings. See **Supplementary Note 1** for a discussion on the effects of smoothing on detected
163 domains.

164 We generated domains at $k=25$, 354, and 670, to match the division, structure, and substructure
165 annotations in CCFv3, displaying domains for four consecutive tissue sections (**Figure 2a**). We also
166 provide representative images of spatial clusters across the brain (28/53 sections) at different k in
167 **Supplementary Figures 1-3**. Low domain numbers such as $k=25$ broadly divide the brain into
168 neuroanatomically plausible patterns, with subregions of striatum (dorsal and ventral marked in **Figure**
169 **2a**) and cortical layers clearly visible. A comparison of cortical layers across these sections shows that
170 CellTransformer domains at $k=25$ are well matched to CCF (**Supplementary Figure 4b**) and correctly
171 identify major classes of layers (1, 2/3 4, 5, and 6) across somatosensory and somatomotor cortex. In
172 particular, we point out the excellent correspondence of domains across tissue sections at $k=25$ across
173 the entire dataset (**Supplementary Figure 1**), with nearly perfect consistency across regions. This
174 suggested that our neighborhood representation method was robust enough to enable integration without
175 modeling of batch or tissue-level covariates.

176 At $k=354$, anterior-posterior subdivisions emerge such as the presence of layer 4 in the motor
177 cortex¹³ (**Figure 2a**, see **Supplementary Figure 4d, e**). Historically, the mouse motor cortex was thought
178 to lack a granular layer 4, however recently, MERFISH, transcriptomic and epigenomic studies have
179 confirmed its existence^{1,15,16}. At $k=100$ and $k=354$, we find a domain corresponding to Layer 4 in the
180 somatosensory cortex which clearly extends to layer 4 in the motor cortex.

181 At $k=670$, the cortical layers identified at lower resolution are further partitioned into superficial,
182 intermediate, and deep strata within several layers. We visualize cortical layers across sections in depth
183 (**Figure 2b**), showing CellTransformer not only identifies fine superficial-deep structure within cortical
184 layers but also preserves the boundary between somatosensory and motor cortex (marked in thick black
185 dotted lines in **Figure 2b**). Taken together these results showed that CellTransformer robustly describes
186 previously known anatomical structures.



187
188
189
190
191
192
193
194

Figure 2. Representative images of spatial domains discovered using CellTransformer on the Allen 1 dataset (53 coronal sections and 500 gene MERFISH panel¹) and comparison to CCF. (a.) Four sequential tissue sections (the inter-section distance is 200 μ m) from anterior (first row, corresponding to section 50) to posterior (bottom row, section 47). In the first three columns, each dot is a cell, colored by spatial domain identified by CellTransformer when clustering was conducted with $k = 25, 354$, and 670 domains (the CCF division, structure, and substructure domain resolutions). Spatial domain labels are depicted with the same colors across sections within the same column. Fourth column shows CCF region registration to the same tissue section. Select regions are annotated

195 with CCF labels. MO: motor cortex, SS: somatosensory cortex; ACA: anterior cingulate, CP: caudoputamen; LSX:
196 lateral septum; MSX: medial septum; VISC: visceral cortex; GU: gustatory cortex; PIR: piriform cortex; OT: olfactory
197 tubule; ACB: nucleus accumbens; HY: hypothalamus **(b.)** Single hemisphere images of same tissue sections in **a.**
198 domains fit at $k=670$, zoomed in on cortical layers of motor cortex (MO) and somatosensory cortex (SS). CCF
199 boundaries are shown in semi-transparent lines, with the boundary between SS and MO outlined in larger black
200 dotted lines. **(c.)** Spatial homogeneity (see **Methods**) of domains from different methods including recently published
201 methods CellCharter and SPIRAL. **(d.)** Average Pearson correlation (averaging over number of domains and method)
202 of the maximum Pearson correlation between the cell type composition (at subclass level, 338 types) vectors of data-
203 driven regions with CCF ones. **(e.)** Average Pearson correlation (averaging over number of domains and method) of
204 optimal matched pairs between data-driven and CCF regions, where CCF regions are only allowed to pair with one
205 data-driven region per comparison. Matches fit using linear programming. **(f.)** Region-by-region Pearson correlation
206 matrix comparing cell type composition vectors from 670 CCF regions (at substructure level) with 670 spatial domains
207 from CellTransformer. The CCF regions are shown on the left with their structure annotations from CCF at division
208 level on the side of the plot. Correlations above 0.9 are shown in bright green to assist in visualization. **(g.)** Cell type
209 (cluster level) by region matrix for 670 CCF regions at substructure level. **(h.)** Cell type (cluster level) by region matrix
210 for 670 CellTransformer regions. Rows are normalized to sum to 1 in both **g.** and **h.** Colors along x-axis in both **g.**
211 and **h.** show cell class annotations from ABC-MWB cell type taxonomy at class level to allow for visualization of
212 composition in terms of known types. Cell types in the “09 CNU-LGE GABA” class are boxed in purple in **g.** and **h.**,
213 matching their color in the legend. Rows of both **g.** and **h.** are grouped using clustering to produce approximately
214 similar structure.

215
216
217
218 We also examined the caudoputamen at various choices of k . At $k=25$, the caudoputamen is one
219 domain, which separates into broad spatially contiguous domains at $k=100$. Interestingly, at $k=354$ and
220 $k=670$, we observe domains that intermingle in a grid-like pattern (**Figure 2a, Supplementary Figure 5**)
221 that strongly resembles the Voronoi parcellation established in Hintiryan et al. (2016)¹⁷ through systematic
222 projection mapping to caudoputamen. Notably, CellTransformer also captures the transition between the
223 quadrant pattern in intermediate caudoputamen (sections 52, 50 and 49 in **Supplementary Figure 5**) to
224 the sequential strip organization (sections 44, 43) which Hintiryan et al. (2016) attributed to the
225 differences in subnetwork reorganization. The correspondence of our transcriptomic domains to the
226 Hintiryan et al. (2016) results, which are exclusively based on projection mapping (non-transcriptomic
227 data), suggests the biological relevance of our representation learning workflow.

228 We compared CellTransformer to several other workflows to capture spatial coherency and
229 multiresolution neuroanatomical annotations in CCF at the division, structure, and substructure levels. For
230 comparison, we used two recent methods, CellCharter¹⁸ and SPIRAL¹⁹ that are scalable to millions of
231 cells as benchmarks. CellCharter builds spatially informed embeddings for domain detection by
232 concatenating the embeddings across scales followed by dimensionality reduction and batch correction,
233 while SPIRAL uses graph-neural networks for batch effect correction and integration across scales.
234 Additionally, we implemented a machine learning baseline that employs k -means clustering on cellular
235 neighborhoods (represented as cell type count vectors). Many of the other GPU-accelerated methods
236 such as scENVI⁴, STACI²⁰, spaGCN⁵, STAligner⁶, STAGATE²¹ or GraphST²² cannot be run on datasets
237 that contain millions of cells due to computational constraints (see **Methods**). Many of these methods
238 require instantiation of a dataset-wide pairwise distance matrix between all cells either on GPU or in RAM,
239 which is a prohibitively large matrix (~60TB for ~4M cells) even for enterprise-level hardware. In contrast,
240 our workflow does not require very large system RAM or extensive preprocessing steps due to our
241 training and inference, maxing out at less than 100GB but requiring significantly less in practice.

242 To quantify the spatial coherence of domains, for each cell we identified its nearest 100 spatial
243 neighbor cells. We then quantified the proportion of neighbor cells within the same spatial domain label as
244 the starting cell (**Figure 2c**). Ideally, we would expect a high proportion of neighbor cells to be in the

245 same spatial domain as the starting cell. In this comparison of neighborhood spatial smoothness,
246 CellTransformer outperforms CellCharter (58.2% better spatial coherence at 670 domains) and SPIRAL
247 (4091.2%). CellTransformer also outperforms the machine learning baseline based on k -means clustering
248 on cellular neighborhoods (61.9% better spatial coherence). For reference, we include the CCF
249 parcellation (dashed purple line) in this comparison to provide an upper bound, as well as spatial
250 coherence using single cell type calls at subclass level (338 types, see **Methods**).

251 To quantify the similarity of detected domains with CCF annotations, we compared the cell type
252 composition of domains using cell type calls from the ABC-WMB taxonomy. We again chose the subclass
253 cell type level, extracting for each domain and for each method a 338-long cell-type vector. We calculated
254 the Pearson correlation of cell type composition vectors computed using the CCF regional annotations at
255 division (25), structure (354) and substructure (670) levels against those of the various methods at the
256 corresponding number of spatial domains. First, for each data-driven domain, we computed the maximum
257 correlation to any CCF domain at the same CCF annotation resolution averaging these numbers across
258 domains. CellTransformer outperforms other methods at mid-granularity and fine-granularity (**Figure 2d**).
259 In this comparison, several data-driven regions can match the same CCF region, which in the worst case
260 could provide an overly optimistic picture of the correspondence between data-driven domains and CCF.
261 To address this, we conducted a second analysis where only one CCF region could be matched to a
262 given data-driven one. We used linear programming to optimize 1:1 pairing of data-driven regions to CCF
263 ones based on their Pearson correlation and averaged these across regions and methods (**Figure 2e**).
264 CellTransformer is highly performant, showing that increase in correlation is not due to redundant
265 matches to a single area in CCF. Visualization of spatial clusters from CellCharter (**Supplementary**
266 **Figure 6-7**) at $k=670$ domains across the brain and in midbrain shows lack of spatial coherence in cortical
267 layers and midbrain, with detected domains distributed in a what appear to be non-biological patterns. In
268 contrast, CellTransformer identified spatially coherent domains and uncovered plausible neuroanatomical
269 structures.

270 To further characterize the similarity of CellTransformer domains with CCF, we plotted the
271 Pearson correlation matrix (**Figure 2f**) between cell type composition vectors generated at 670 domains
272 (substructure level in CCF). Block structures with very high correlations (>0.9 , shown in bright green) in
273 the matrix clearly show that CellTransformer is able to identify regions that are highly similar with known
274 ones without any labels. We also investigated correspondence of cell type composition with more
275 granular single cell annotations, employing the “cluster” (5274 cell types) level annotations from ABC-
276 MWB. We observed high similarity between the “substructure” CCF domain set (**Figure 2g**) and 670
277 CellTransformer domains (**Figure 2h**) with average Pearson correlation of CellTransformer to CCF
278 domains of 0.853. This shows the high correspondence of CCF and CellTransformer (**Figure 2g** and
279 **Figure 2h**) is robust to cell type resolution at which comparison occurs. CellTransformer identified an
280 increase in number of domains containing the 09 CNU-LGE GABA class (striatal/pallidal GABAergic
281 neurons from lateral ganglionic eminence compared with the 670 CCF substructures, shown in light
282 purple box in **Figure 2g** and **Figure 2h**), potentially suggesting the presence of uncharacterized
283 developmental populations.

284 The observation of hierarchical grouping of domains at different choices of k (for example
285 delineation of cortical layers and sublayers with increasing number of domains) prompted us to develop a
286 strategy to evaluate an optimal number of spatial domains based on two metrics. We implemented a
287 previously published strategy²³ to determine the optimal number of domains using a stability criterion. We
288 reasoned that the optimal choice of spatial domain number would feature minimal variability across
289 clustering runs. In brief, we computed 20 clustering instances with different random initializations for a
290 large range k values (100-2000) and quantified their variability over these initializations (see **Methods**).
291 Interestingly, stability increased with increasing k (**Supplementary Figure 8a, 8b**). To facilitate the choice
292 of a particular resolution for analysis, we also computed the inertia (sum of squared errors) for each

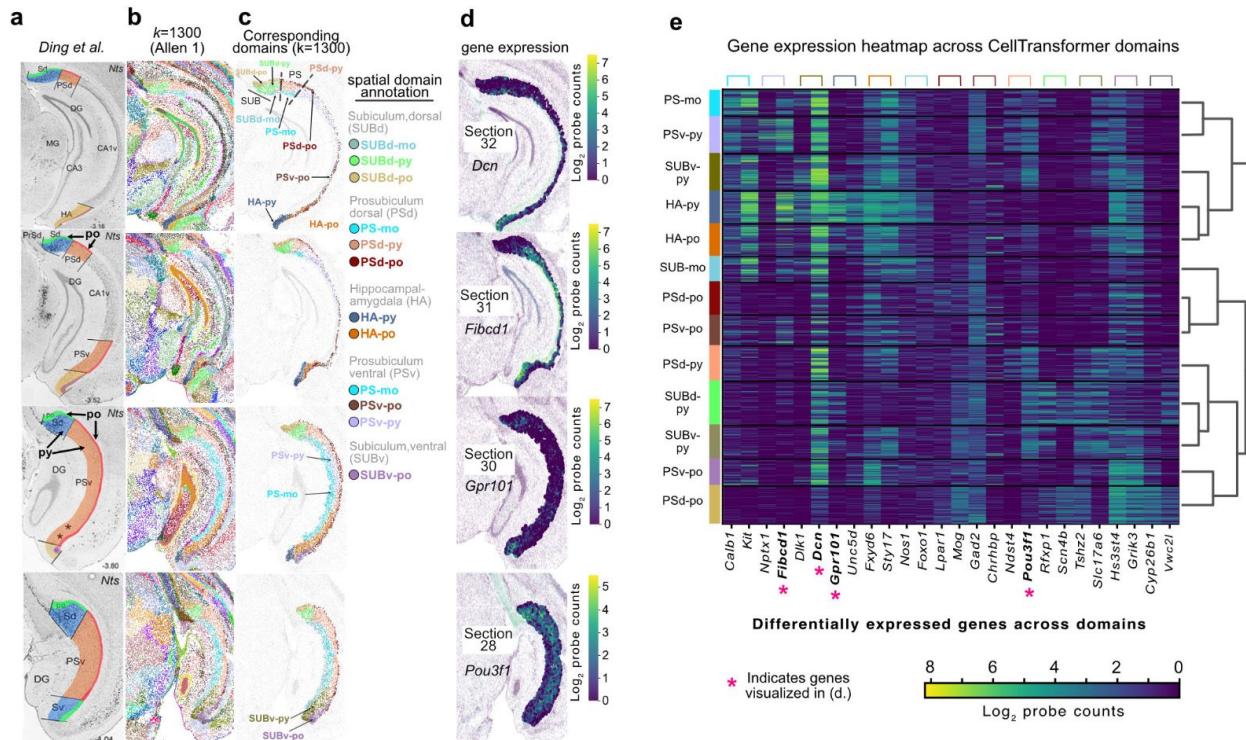
293 clustering solution. Low stability at small numbers of domains may partially explain subpar results for
294 CellTransformer in the $k=25$ CCF evaluations. We averaged the inertia curve and instability and
295 computed the point of second derivative crossing to identify $k=1300$ as our resolution for analysis
296 (crossing point shown with red dot in **Supplementary Figure 8c**).

297 CellTransformer is the only method out of the three we implemented (including six other pipelines
298 which were unable to cope with the size of ABC-MWB dataset) to allow discovery of spatially coherent
299 divisions at greater than CCF resolution. To our knowledge, this study establishes the first instance of a
300 data-driven method using spatial transcriptomics data to identify brain regions at resolutions finer than
301 previously defined in the CCF. We next sought to establish correspondence of particular domains at
302 $k=1300$ to known neuroanatomy.

303 Mapping of spatial domains in the hippocampal formation

304 We characterized CellTransformer's ability to capture known anatomical structure in the
305 hippocampal formation, notably the subiculum (SUB) and prosubiculum (PS), in the Allen 1 dataset. We
306 focused on this area because it is well characterized with respect to both connectivity²⁴ and transcriptomic
307 composition^{25,26}. These structures were investigated in Ding et al. (2020)²⁷, where the authors performed
308 consensus clustering of glutamatergic neurons and subsequent ISH experiments were used to
309 comprehensively map domains in dorsal subiculum (SUBd) and dorsal and ventral prosubiculum (PSd
310 and PSv). Specifically, this and other recent works have noted the extensive laminar organization
311 (superficial layers to deeper layers), and the dorsal-ventral organization of the subiculum²⁸⁻³⁰. This
312 organization has been attributed to distinct and correlated patterns of gene expression and connectivity.

313 We qualitatively compared spatial domains discovered by CellTransformer with $k=1300$ to the
314 anatomical borders identified in Ding et al. (**Figure 3a**). The subiculum features a three-layer organization
315 referred to as molecular (mo) layer, a pyramidal cell (py) layer, and polymorphic (po) cell layer. **Figure 3a**
316 shows a diagram of SUB and PS regions based on Ding et al. (2020) with the pyramidal and polymorphic
317 layers of SUB and PS annotated in bold black text. **Figure 3b** shows discovered spatial domains at
318 $k=1300$ across four sequential sections corresponding to those in Ding et al. (2020). A subset of domains
319 corresponding to SUB and PS are shown in **Figure 3c** along with putative regional annotations.
320 CellTransformer identifies a three-layer organization in the dorsal subiculum corresponding to that in Ding
321 et al. (2020) labeled SUBd-py (light green), SUBd-po (gold), and SUBd-mo (gray-blue). CellTransformer
322 also correctly splits the SUBd and PSd shown with black dotted lines on the image of section 32. Three-
323 layer strata are also observed in PSd, although notably the pyramidal layer domain extends caudally,
324 consistent with transcriptomic studies²⁴⁻²⁶ of SUB architecture. For instance, our PSd-po domains
325 (sections 31 and 30) strongly resemble the HGEA layer 4 found in Bienkowski et al. (2018)²⁴. Note that
326 differences may arise between panels in **Figure 3a** and **3c** because of sectioning variability and lack of
327 exact match between sections in ABC-MWB and the Ding et al. (2020) study. In addition to the
328 aforementioned regions we also observe high agreement in areas such as in the hippocampus-
329 amygdaloid transition area (HA) and ventral prosubiculum (PSv).



330

331 **Figure 3.** Comparison of CellTransformer domain sets identified in the Allen 1¹ dataset with $k=1300$ with a
 332 comprehensive region set found in Ding et al. (2020), reproduced with permission from authors. (a.) Representative
 333 images reproduced from Ding et al. of region boundaries in prosubiculum (PS), subiculum (SUB), and hippocampal-
 334 amygdala (HA) particularly along the dorsal-ventral axis. Polymorphic and pyramidal layers of the dorsal subiculum
 335 (SUBd) and ventral prosubiculum are indicated. (b.) Images from hippocampal formation across 4 sequential tissue
 336 sections (anterior to posterior) roughly aligned to sections presented in Ding et al. Each dot is a cell colored using
 337 domain labels with $k=1300$. (c.) Same as b. but only showing cells inside PS, SUB, and HA. Putative regional
 338 annotations are indicated and grouped by dorsal or ventral region within PS and SUB. (d.) Gene expression patterns
 339 visualized at the corresponding tissue section, where only cells within PS/SUB/HA are shown. Units are in \log_2 probe
 340 counts. (e.) Gene expression heatmap of identified subregions, with putative anatomical annotation. Dendrogram
 341 from hierarchical clustering in gene expression space is shown to the right. Genes visualized in d. are bolded and
 342 denoted with a pink asterisk. Colored brackets indicate the genes which are differentially expressed with respect to
 343 the domain (colors match those shown on the left of the heatmap). Two genes per domain are shown and each gene
 344 is expressed with at least log-fold change greater than 1 relative to the other domains. Abbreviations: PS-mo:
 345 prosubiculum molecular layer; PS-py: pyramidal layer of subiculum; SUBv-py: ventral subiculum, pyramidal layer; HA-
 346 py: hippocampal-amygdaloid transition area, pyramidal layer; HA-po: hippocampal-amygdaloid transition area,
 347 polymorphic layer; SUB-mo: subiculum, molecular layer; PSD-po: dorsal prosubiculum, polymorphic layer; PSV-po:
 348 ventral prosubiculum, polymorphic layer; PSD-py: dorsal prosubiculum, pyramidal layer; SUBd-py: dorsal subiculum,
 349 pyramidal layer; SUBv-py: ventral subiculum, pyramidal layer; PSV-po: ventral prosubiculum, polymorphic layer; PSD-
 350 po: dorsal prosubiculum, polymorphic layer.

351

352

353

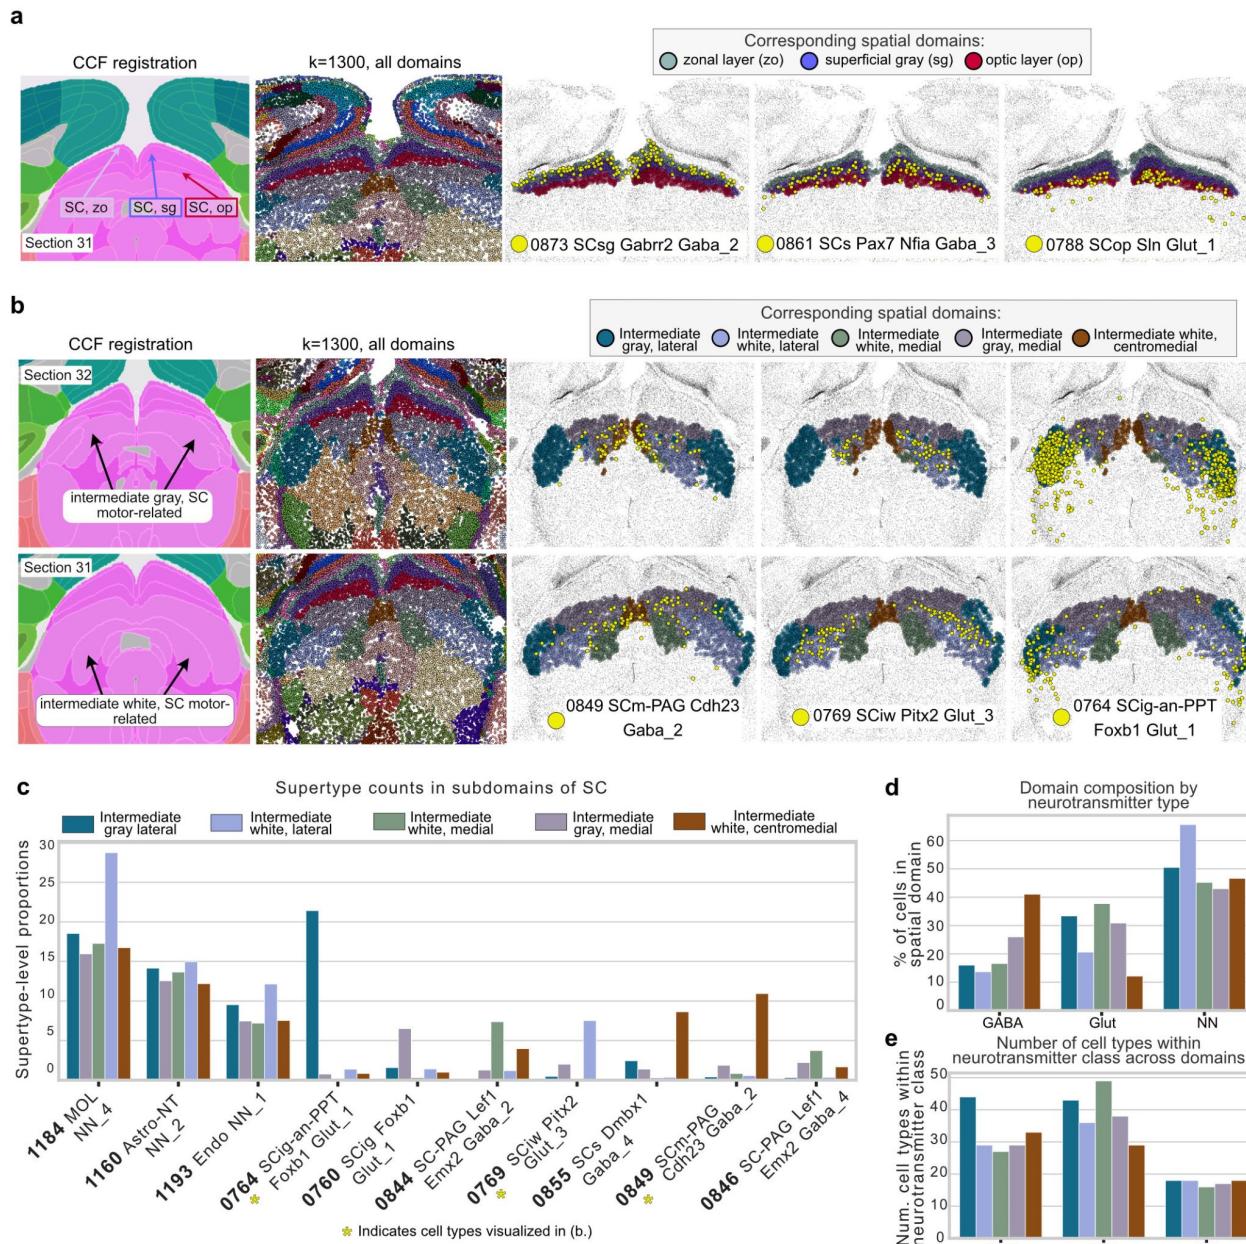
354 Ding et al. (2020) observed differential projection topology in dorsal subiculum versus ventral
 355 prosubiculum. Correspondingly, genes were found to form opposing gradients across the length of
 356 subiculum areas. Dorsolateral gene gradients appeared in SUBd and ventromedial gradients in PSV. Since
 357 CellTransformer domains appeared to correspond well with literature results, we explored gene
 358 expression patterns across domains to verify whether dorsal-ventral and medial-laterally varying gene
 359 patterns could be observed. We conducted differential expression analysis across our subiculum domains

360 (Figure 3e) that when visualized (Figure 3d) clearly reflected these gradients. Many genes expressed in
361 SUB and PS traverse their long axis as reported previously²⁵. The identification of spatial domains which
362 subdivided specific layers of PS and SUB similarly to the results in Ding et al., and featured similar types
363 of gene expression gradients as existing literature, suggests that our pipeline was successful in learning
364 neuroanatomically useful information. Importantly, while results in Ding et al. and related works were
365 enabled by significant neuroanatomical and experimental expertise, CellTransformer allows identification
366 of granular tissue structure in a data-driven fashion. Encouraged by this result, we continued our
367 investigation of CellTransformer correspondence with known literature with a comparison in superior
368 colliculus.

369 CellTransformer allows for quantification of laminar and columnar 370 organization in superior colliculus

371 Recent studies using systematic mapping of cortico-tectal fibers in superior colliculus (SC) have
372 identified distinct laminar and columnar structure³¹, suggestive of the complex role SC plays in integration
373 of sensory information and coordination of signals. Therefore, SC presented an excellent opportunity to
374 identify transcriptomic and cellular correlates of connectomic variation. We observed a strong
375 correspondence of three of our spatial clusters ($k=1300$) in the Allen 1 dataset with known layers of
376 superior colliculus, sensory area, particularly the zonal (zo), superficial gray (sg), and optic (op) layers
377 across a set of tissue sections spanning ~600 μ m from anterior to posterior (rows of **Figure 4a** and
378 **Supplementary Figure 9a**). CellCharter was unable to identify these structures (**Supplementary Figure**
379 **7**) and only identifies two layers in SC, which does not conform with existing results.

380 By visualizing the cell type composition within the top-ten most abundant types for these three
381 spatial domains (**Supplementary Figure 9a, 9b**), we were able to identify cell types that were highly
382 selective for our data-driven SC layers: types 0873 SCsg Gabrr2 Gaba_2, 0861 SCs Pax7 Nfia Gaba_3,
383 and 0788 SCop SIn Glut_1. Crucially, the cell types, which have already been annotated as being
384 associated with one of the zonal, optic, or superficial gray, are identified automatically by
385 CellTransformer. We chose the supertype level to allow inspection of abundant cell types without being
386 difficult to visualize. Supertype-level visualizations also show that even with granular cell types (1201
387 types in Yao et al.) CellTransformer domains are often marked by spatially specific cell type patterning;
388 we note that we do not filter cells outside of our putative superior colliculus layers for visualization. Next
389 we visualized the percentage of cells in each domain (**Supplementary Figure 9c**), grouping them by
390 neurotransmitter class (GABA-ergic, glutamatergic, and non-neuronal). The superficial gray layer showed
391 the higher proportion of GABA-ergic neurons, while the optic layer had the highest proportion of
392 glutamatergic neurons. To further explore these relationships, we calculated the number of distinct cell
393 types (supertype level) within each neurotransmitter class and domain. A clear dorsal-ventral organization
394 was evident (**Supplementary Figure 9d**) with the number of GABA-ergic and glutamatergic neuron types
395 increasing with layer depth, suggesting CellTransformer's ability in capturing complex patterns of cellular
396 spatial organization.



397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412

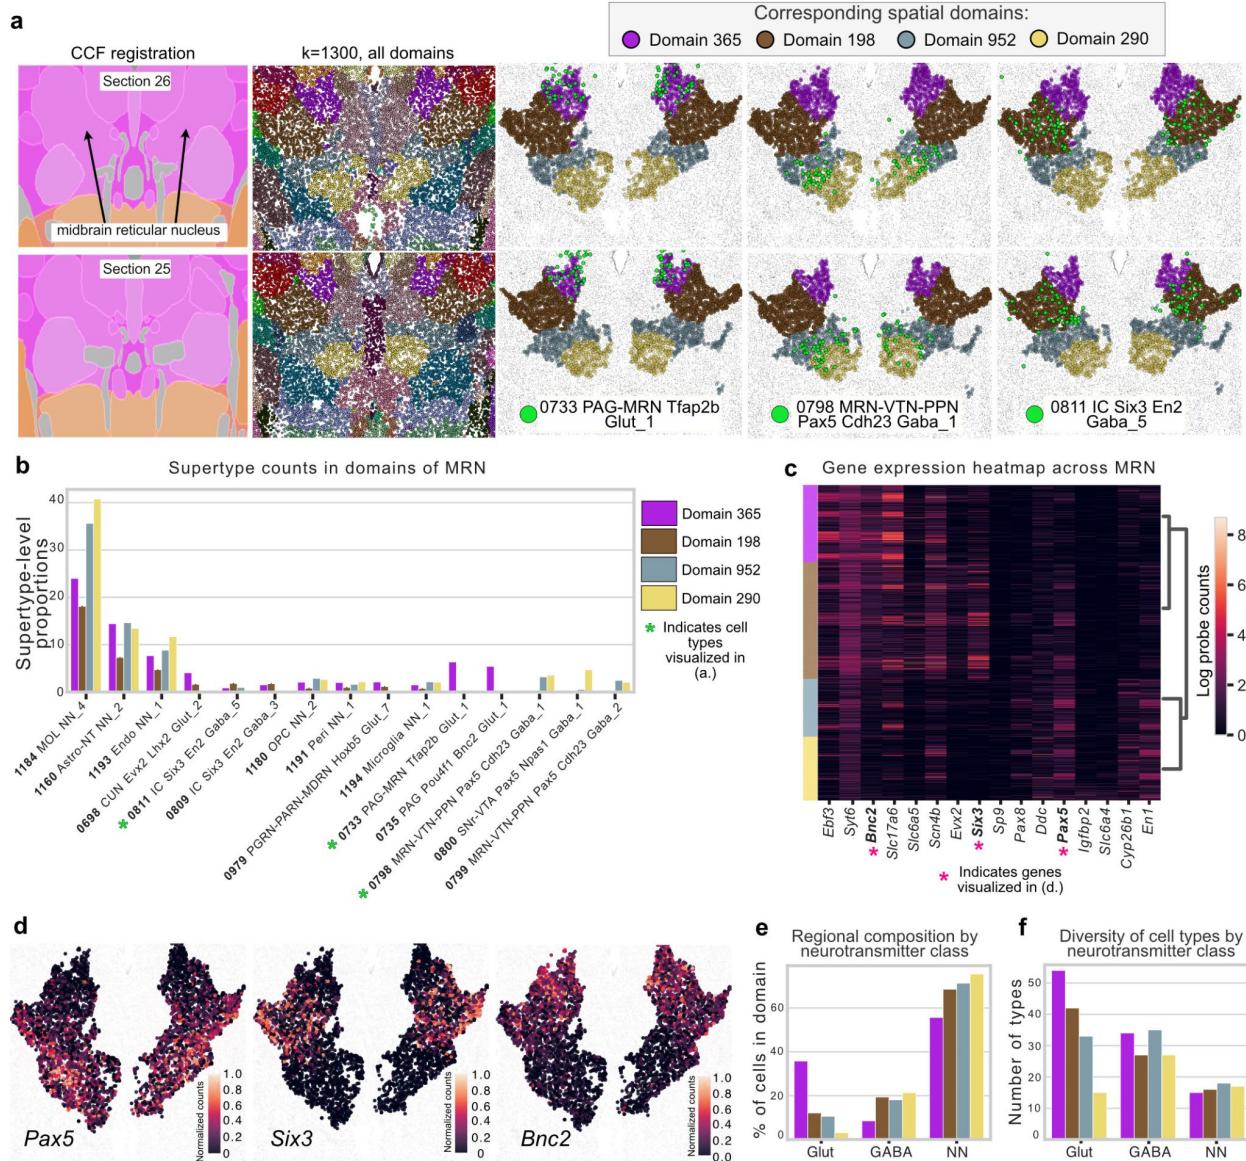
Figure 4. Examination of putative dorsal and ventral subregions of superior colliculus identified in the Allen 1¹ dataset using CellTransformer. (a.) Putative subregions of sensory layers of superior colliculus in tissue section 32 identified with $k=1300$. CellTransformer domains. CCF registration is in the first column, with zonal (zo), superficial gray (sg), and optic (op) layers labeled by the color of their CellTransformer domain in third, fourth, and fifth columns. The second column shows all cells with color labels from their spatial domain from CellTransformer at $k=1300$. The third, fourth, and fifth columns show the putative zo (gray-green), sg (purple), and op (red) domains. These columns also show the spatial distribution of one supertype level cell type in yellow across the section. (b.) Sequential tissue sections (32: anterior, 31: posterior) shown similarly to a., but visualizing subregions of the intermediate gray and intermediate white layers, which are indicated with black arrows in the CCF registered annotation image. (c.) Proportions of different supertype level cell types for top-ten most abundant types in different spatial domains. Colors refer to the same spatial domain label in (a.) and (b.). Cell types visualized in (b.) are denoted with a yellow asterisk. (d.) Barplot of the percentage of cells of a given neurotransmitter class found in a given region (GABA - GABAergic; Glut - Glutamatergic; NN - non-neuronal). (e.) Number of unique cell types (at supertype level) found in each domain, grouped by neurotransmitter class.

413 Encouraged by these findings, we also investigated subregions of the intermediate gray and
414 intermediate white areas of the motor-related areas in SC (**Figure 4a-b**), where we identify consistent
415 regions across two consecutive sections that are not annotated in the CCF (rows of **Figure 4b**). We
416 define subregions of intermediate gray (ig) and white (iw), noting a medial-lateral structure similar to that
417 in Benavidez et al. (2021)³¹, which exhaustively cataloged projection zones in superior colliculus. Notably
418 unlike in superior colliculus sensory, a significant number of non-neuronal cell types are found in very
419 similar proportions across the intermediate white and gray layers (**Figure 4c**), and instead differences in
420 regions may be attributable to varying proportions of rare cell types. Encouragingly, even in these fine-
421 grained areas, cell types that are highly specific for our data-driven layers can be readily identified
422 (columns of **Figure 4b**). These rare domain-enriched cell types include: 0849 SCm-PAG Cdh23 Gaba_2
423 (enriched in the medial intermediate white layer, shown in dark green), 0769 SCig SCiw Pitx2 Glut_3
424 (enriched in lateral intermediate white, shown in light blue), and 0764 SCig-an-PTT Foxb1 Glut_1
425 (enriched in medial intermediate gray, shown in dark blue). The identification of *Pitx2*-expressing neurons
426 also supports our assertions that CellTransformer identifies biologically relevant domains, with previous
427 studies using *Pitx2* expression specifically as an intermediate layer marker in superior colliculus^{16,28}.

428 We observed complex cell type abundance gradients when visualizing the percentage of cells in
429 a given domain by their neurotransmitter type (**Figure 4**). We used supertype level to confirm that
430 spatially-varying cell distribution patterns persisted when using more granular cell type annotations.
431 Lateral domains such as intermediate gray, lateral (shown in dark blue) and intermediate white, lateral
432 (shown in light blue) featured a smaller proportion of GABA-ergic neurons than medial domains but were
433 enriched for glutamatergic neurons and non-neurons (**Figure 4d**). Despite the low proportion of GABA-
434 ergic neurons, the lateral domain of intermediate gray possessed the highest number of GABA-ergic
435 neuron cell types (**Figure 4e**). Conversely, non-neurons featured the same number of types across the
436 laminae.

437 A medial-lateral gradient of inhibitory neurons in the midbrain 438 reticular nucleus

439 Next, we investigated the midbrain reticular nucleus (MRN), a subcortical structure with few
440 anatomical annotations in CCF. MRN is highly enriched for interneurons in a dense array of connections
441 and appears to play a complex role in movement initiation and release^{29,30}. CellTransformer identifies four
442 subregions of the MRN, which are not included in the existing CCF annotation (**Figure 5a**). Plotting cell
443 type proportions across the MRN, we identified cell types which are enriched for these putatively
444 uncharacterized areas, although all domains were predominantly glial (e.g., 1184 MOL NN_4 supertype is
445 abundant in all regions, **Figure 5b**). By visualizing differentially expressed genes across the domains
446 (**Figure 5c**), we identified genes which were subregion selective (selected genes shown in **Figure 5d**)
447 and form dorsal-ventral expression gradients. Hierarchical clustering showed that the two dorsal domains
448 (purple and brown) group together with the two ventral ones (gold and gray).



449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465

Figure 5. Subregions of midbrain reticular nucleus discovered in the Allen 1¹ dataset using CellTransformer. **(a.)** Sequential tissue sections (26 and 25, anterior to posterior). First column: CCF registration with approximate location of midbrain reticular nucleus shown with arrows. Note that registration is not exact and can differ across hemispheres. Second column: all cells in field of view, with color from spatial domains determined with $k=1300$. The rest of the columns show only cells located in the MRN, and each column shows a different supertype level cell type in green. **(b.)** Supertype level cell type proportions for top fifteen most abundant types across the MRN subregions. Cell types visualized in **(a.)** are denoted with green asterisks. **(c.)** Selected 4 differentially expressed genes across regions. Each gene is expressed at least log fold change greater than 1 relative to the other domains. MERFISH probe distributions for select genes indicated with pink asterisks are shown in **(d.)**. **(d.)** Gene expression gradients across tissue section 25 for *Pax5*, *Six3*, and *Bnc2*, showing specificity for each of the putative MRN subregions. Intensity of color is 0-1 normalized after log scaling raw probe counts. Each dot is a cell, and the color shows the relative transcript count. We show only cells within the subregions to make it visually easier to distinguish the relevant cells. **(e.)** Bar plot of the percentage of cells for a given neurotransmitter type found in each domain, (GABA - GABAergic; Glut - Glutamatergic; NN - non-neuronal). **(f.)** Number of unique cell types (at supertype level) found in each domain, grouped by neurotransmitter class..

466 We again visualized the neurotransmitter composition and number of unique cell types of given
467 neurotransmitter classes. We observed that the number of types of excitatory neurons was spatially
468 graded. Dorsal domains of MRN (domain 365 shown in purple and domain 198 shown in brown, (**Figure**
469 **5e**) featured the highest proportion of glutamatergic neurons, and the proportion of glutamatergic neurons
470 in MRN domains decreased with increasing depth. Nonneuronal cells are also organized along this
471 gradient, but in the opposite direction, the ventral areas featuring the highest proportion of glia and the
472 dorsal the lowest. Interestingly, MRN domains composed of a higher proportion of glutamatergic neurons
473 were also the ones with the greatest number of glutamatergic neuron types, also following a dorsal-
474 ventral gradient (Pearson correlation $r = 0.89$). This relationship was observed for the nonneuronal cells
475 (Pearson correlation $r = 0.81$, **Figure 5f**), but not for GABA-ergic neurons (Pearson correlation $r = -0.64$).
476 This suggests that CellTransformer can identify plausible structures even in historically difficult to
477 characterize areas.

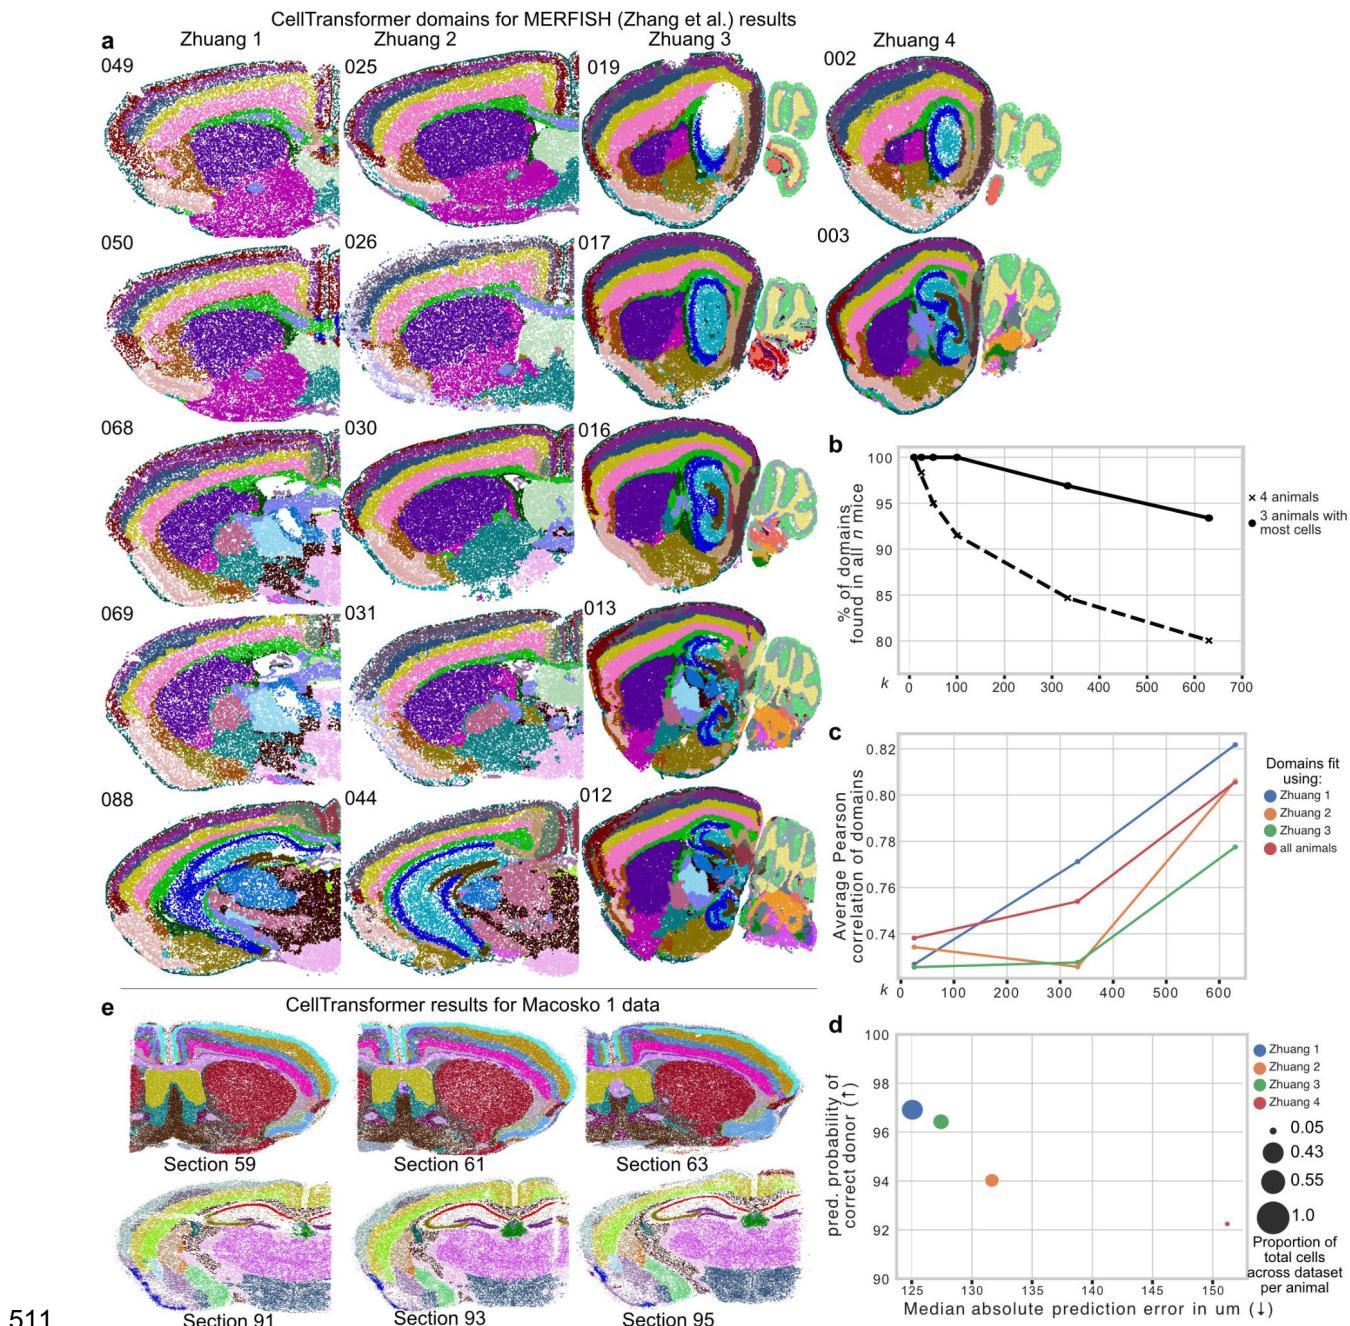
478 **CellTransformer enables scaling up to multi-animal, million cell**
479 **datasets and generalizes to other spatial assays**

480 In order to investigate CellTransformer's ability to integrate across animals, we trained a new
481 model from scratch on the Zhang et al. (2023)¹¹ MERFISH data, which uses an 1129 gene panel and is
482 split over four animals, with both coronal (Zhuang 1 and 2) and sagittal sections (Zhuang 3 and 4). We
483 computed embeddings for each neighborhood as in the previous analysis and performed k -means,
484 concatenating representations for all mice and sections. This provided an opportunity to examine whether
485 CellTransformer could adapt to a multi-animal case in addition to finding spatial domains across tissue
486 sections of the same animal.

487 Spatial domains in sequential tissue sections appeared highly concordant across all four mice
488 (**Figure 6a**) at the 50-domain resolution. We used 50 domains to facilitate clear visualization of the
489 domains across animals with relatively few colors. Coronal and sagittal sections across mice clearly
490 corresponded anatomically. Cortical layers were highly consistent across animal and section orientation.
491 Structures that appear in the coronal view can be readily identified in the sagittal sections. For example
492 the hippocampal formation (blue) is well delineated in sections 088 for Zhuang 1, section 044 for Zhuang
493 2, and across displayed sections of Zhuang 3 and Zhuang 4. Despite a relatively low number of cells in
494 mouse 4 (162,579 cells versus more than 1.5 million for each of the other animals), nearly all spatial
495 domains observed for Zhuang 4 are present in other animals. Note that sections from this animal only
496 cover a section of the lateral portion of the brain and do not span the entirety of the sagittal plane.

497 We quantified the robustness of CellTransformer domains in a multi-animal context across and
498 within Zhuang 1-4 datasets. We ran clustering and identified domains at the three values of k : 25, 333,
499 630. These k values correspond to three CCF resolution levels reported by registration in Zhang et al.
500 (2023) (note the number of domains differs due to registration differences). For each k value, we counted
501 the number of domains observed in all four animals. We also repeated this analysis without data for
502 Zhuang 4, which contains far fewer cells than the datasets from other animals (**Figure 6b**). We find that
503 even at high resolution (630 domains), 93.3% domains were found in each mouse, showing high
504 consistency of CellTransformer domains across datasets. With the Zhuang 4 included, at 630 domains,
505 80.0% domains were found in every animal. To verify that domain consistency across animals was not
506 related to loss of domain spatial coherence, we repeated the neighborhood smoothness analysis we
507 developed for analysis of the Allen 1 dataset on the combined Zhang et al. (2023) data. Spatial
508 smoothness was similar to that of Allen 1 (**Supplementary Figure 10**), indicating CellTransformer can
509 discover spatially coherent domains that are robustly integrated across animals.

510



512 **Figure 6.** Investigation into performance of CellTransformer on the Zhuang 1-4 datasets (239 sections, both
513 coronal and sagittal, with a 1129 gene MERFISH panel¹¹). (a.) Representative images of all four mice arranged by
514 column. The section number for each mouse is shown in the upper left of each image. Note that Zhuang 4 only had
515 three sections. For each image, each dot is a cell neighborhood and colors come from a spatial clustering with $k=50$
516 (number of CCF regions at structure level), fit by concatenating embeddings across mice. (b.) Quantification of
517 number of per-mouse specific spatial clusters, computed by clustering at different k and computing the number of
518 clusters found for all mice (4 animals) and for the three mice with the most cells per mouse (Zhuang 1, 2, and 3).
519 Note that because serial sections were collected at a higher frequency (100 μm versus 200 μm), different areas of the
520 brain will have marginally higher coverage in one brain or another. (c.) Average correlation of the cell type
521 composition of brain regions computed CellTransformer to CCF regions, computed using the linear-sum assignment
522 matching algorithm (exclusively matching regions from one set to the other). Dotted lines with "o" marker indicate
523 results when fitting using all three mice with >1M cells together. Solid lines with "x" marker indicate results when

524 computing spatial clustering on each mouse in isolation. (d.) Quantification of subject-level information present in
525 embeddings using linear regression. The median absolute prediction error (x-axis) quantifies accuracy in predicting
526 the (x, y, z) coordinates of a neighborhood from its embeddings. The y-axis quantifies accuracy when predicting
527 mouse identity from embeddings using logistic regression. Values are averaged across cells per mouse. (e.) Results
528 of domain discovery ($k=50$) on a Slide-SeqV2 whole mouse brain (Macosko 1) dataset¹². Two sets of three sequential
529 sections are shown.

530
531
532

533 We next quantified the similarity of CellTransformer domains to CCF regions. Similarly to our
534 analysis of the Yao et al. (2023) dataset, we computed average similarity of cell type composition vectors
535 from CCF and CellTransformer. In domain discovery across all animals, we found cell-type composition
536 vectors that correspond strongly to CCF (Pearson correlation = 0.805, red line in **Figure 6b**). We also
537 evaluated whether clustering only on embeddings from one animal would significantly affect similarity to
538 CCF. Correspondence between CellTransformer domains and CCF is high even when domains are fit
539 with a subset of the dataset (Pearson correlations > 0.7 for all comparisons across resolutions and
540 domain source, **Figure 6c**). This demonstrates CellTransformer can reproduce a consistent
541 neuroanatomical structure even with a small number of observations. Results were highly similar overall
542 to CCF in the Zhang dataset and Yao dataset (Pearson correlations greater than 0.6 for all comparisons),
543 indicating CellTransformer's robustness to changes in gene panel and preprocessing choices.

544 To further investigate how donor metadata was encoded in the embeddings, we employed linear
545 probing strategies commonly used in interpretation of deep learning model embeddings. We regressed
546 CCF-registered (x, y, z) coordinate position across all embeddings and used logistic regression to classify
547 animal identities. Neighborhood level prediction of donor identity was very accurate ($>94\%$ for all animals,
548 **Figure 6d**) and median absolute prediction error was accurate within 151 μm . Decoding accuracy for both
549 metrics was strongly associated with the number of cells (Pearson correlation -0.92 for coordinate error
550 and 0.92 for predicted donor probability). The observation that mouse donor identity is easily predicted
551 from per-neighborhood embeddings while still maintaining cross-animal and cross-section coherence is
552 another demonstration of the richness of the representation learned by our approach.

553 Finally to demonstrate the applicability of our strategy to a different spatial transcriptomics
554 modality, we analyzed a whole mouse brain Slide-SeqV2 dataset (“Macosko 1”), collected in Langlieb et
555 al. (2023)¹². Slide-SeqV2 provides whole transcriptome coverage in a spatial context by tiling tissue slices
556 with 10 μm by 10 μm squares. As each square may contain more than one cell or a partial cell, we fit our
557 model to the deconvoluted single-cell data computed using the RCTD³² method provided. This produced
558 4,783,976 cells across 101 slices. We also filtered the dataset for low-quality cells and infrequently
559 expressed genes (see **Methods**). We found that increasing the size of our model (from 4 encoder layers
560 to 10) was necessary to identify spatially coherent domains, perhaps driven by the much larger number of
561 genes detected (5019 versus 500 or 1129 in the two MERFISH datasets; see **Methods**). We plot three
562 sequential sections from domain discovery at $k=50$ (**Figure 6e**). We show that CellTransformer robustly
563 identifies cortical layers across sections and in known structures such as the midbrain and the piriform
564 areas. Domain discovery with greater than 50 regions did not produce adequate integration across
565 sections, possibly because of variable cellular density and single-cell read depth across sections.

566 Overall, we found that the CellTransformer workflow successfully identifies interpretable domains
567 across different spatial transcriptomics modalities, and that the resolution of cross-section and cross-
568 dataset domains depends on the specific spatial transcriptomics method and the quality of the datasets.

569 Discussion

570 In this study, we present a transformer-based pipeline to combine scRNA-seq and spatially resolved
571 transcriptomic atlases to perform accurate organ level domain discovery. We employed a novel
572 representation learning workflow and implemented a computationally efficient pipeline that readily allows
573 scaling to multi-million cell, multi-animal datasets. The representations learned in our model can be
574 clustered to identify progressively finer-scale spatial domains directly from local cellular and molecular
575 information alone, without predefined spatial labels. We show these regions can be interpreted at the
576 gene or cell level and recapitulate a variety of existing findings in the neuroscience literature where many
577 existing methods cannot. Our pipeline allows extraction of a very high number of domains, which retain
578 high correspondence to existing brain region ontology (CCF). These domains are also highly spatially
579 consistent both within and across tissue sections and even over multiple animals. To our knowledge this
580 is the first demonstration of fine-grained domain detection beyond human annotation (CCF) using data-
581 driven methods in transcriptomic data. Not only can CellTransformer discover this fine structure, but it can
582 reliably find it across animals even with hundreds of regions. This capability is intrinsic to our model, and
583 is learned despite any conditional modeling for donor or section-level covariates, indicating the
584 robustness of learned features.

585 We demonstrated the robustness of our model in uncovering biologically relevant models and
586 characterized our pipeline's ability to reproduce known neuroanatomy in the hippocampal formation and
587 superior colliculus. Detected domains were concordant with previous comprehensive transcriptomic and
588 connectivity studies of these areas but were identified in a scalable and data-driven way. Not only were
589 we able to detect known and novel regions, we found that CellTransformer domains can recapitulate and
590 extend known spatial cell type enrichment patterns and gene expression gradients.

591 We highlight several advantages of our architecture and approach. Although the use of graph-
592 structured architecture or self-attention to model cells in a neighborhood graph is not novel^{14,21,33}, our
593 approach combines novel self-supervised training objective based on spatial correlation between a cell
594 and its neighbors that facilitates learning of a fixed representation for a cellular neighborhood
595 (**Supplementary Note 2**). The intersection of graph neural networks, transformers, and representation
596 learning research is a rich and rapidly moving research area. Methods for spatial-graph structured data
597 such as CellTransformer will benefit immensely from implementing more effective ways of encoding the
598 data and its metadata such as better position encoding mechanisms³⁴, rotationally-invariant
599 architectures³⁵, or arbitrary numbers of genes³⁶. There are also significant opportunities to extend
600 CellTransformer's local representation framework to include other data modalities. Using a transformer
601 than graph neural network facilitates inclusion of arbitrary contextual data such as cell-level (e.g.
602 neurophysiology^{37,38}) and pixel-level data (e.g. mesoscale axonal connectivity³⁹, or magnetic resonance
603 imaging⁴⁰) which can be tokenized and included in our framework.

604 Caveats of our approach include the necessity of a user-specified spatial radius (for
605 neighborhood computation) and choice of k for spatial cluster detection. The stability of detected spatial
606 domains at a given radius or k poses an interesting future angle from which to study anatomical
607 hierarchies in the brain. Users must also have access to GPUs (to allow for timely model fitting), which
608 reduces overall accessibility, although the hardware requirements are still much less intensive than for
609 many existing models such as spaGCN and scENVI⁴.

610 CellTransformer advances the state of the art for automated domain detection by allowing
611 identification of granular and biologically relevant spatial domains that is extensible to both very large,
612 multi-animal spatial transcriptomic datasets. As spatially resolved transcriptomic and multi-omic studies
613 of the brain become more prevalent, tools such as CellTransformer provide avenues to transform data
614 into refined anatomical maps of the brain and other complex organs and pave the way towards tissue-
615 level structure-function mapping.

616 Methods

617 Allen Brain Cell Mouse Whole Brain (ABC-MWB) dataset processing

618
619 **Allen Institute for Brain Science dataset preprocessing**
620 We downloaded the log-transformed MERFISH probe counts and metadata for the Allen Institute for
621 Brain Science animal ("Allen 1") from the Allen Institute public release
622 (https://alleninstitute.github.io/abc_atlas_access/intro.html) access for ABC-MWB. The Allen 1 dataset is
623 composed of 53 coronal sections. The MERFISH probe set included 500 genes. Serial sections were
624 collected at 200 μm intervals. We used the taxonomy from the "20231215" data release. Allen 1 is
625 composed of 3,737,550 cells. We transformed the (x, y) coordinates of each cell into microns instead of
626 mm as provided. Otherwise the dataset was used as-is for neural network training.
627

628 Zhuang lab (Zhang et al.) dataset processing

629 Data were downloaded from the "20230830" data release from the Allen Institute ABC-MWB public data
630 release. Two animals ("Zhuang 1" and "Zhuang 2") were collected with coronal sections. The other two
631 animals ("Zhuang 3" and "Zhuang 4") were collected sagittally. Serial sections for Zhuang 1 (female) were
632 collected at 100 μm intervals, while serial sections for the other animals (all male) were collected at 200
633 μm intervals. The size of the MERFISH probe set included 1129 genes. Zhuang 1 and Zhuang 2 consist
634 of 2,846,909 cells and 1,227,409 cells, respectively. Zhuang 3 and Zhuang 4 consist of 1,585,844 cells
635 and 162,579 cells, respectively. We transformed the (x, y) coordinates of each cell into microns instead of
636 mm as provided. Otherwise, the data were used as-is for neural network training.
637

638 Cellular neighborhood construction

639 We consider cells in the same neighborhood as a reference cell if the distance between them is within a
640 box of fixed size. For all MERFISH datasets we used a box width of 85 μm .
641

642 CellTransformer architecture

643 We construct a CellTransformer to generate a latent representation from a cellular neighborhood where
644 this representation is composed of both molecular and cell type information. We represent cells as nodes
645 in an undirected graph, $\mathcal{G} = \mathcal{V}, \mathcal{E}$ where \mathcal{V} indexes the nodes in the graph (cells) and we add an edge
646 (v_i, v_j) to the edge set \mathcal{E} if $d_{i,j} < r$, with r a user-specified distance in microns. We assume also that for
647 each node we have access to $\mathbf{x}_i \in \mathbb{R}^g$, a g -dimensional vector of MERFISH probe or cell deconvoluted
648 transcript counts. We also assume we are given class labels $\mathbf{c} = \{c_i \in \{1, \dots, C\}\}_{i=1}^N$ for each of the N
649 cells. The user must also specify an embedding dimension and number of transformer encoder and
650 decoder layers; in all experiments in this paper using MERFISH data we use an embedding dimension of
651 384, 4 encoder layers, and 4 decoder layers. For Slide-SeqV2 analysis we used 10 encoder layers and 4
652 decoder layers.
653

654 To generate a neighborhood embedding, we identify a particular cell which we call a reference
655 cell. Its first degree neighbors are extracted from \mathcal{G} . We first apply a shallow encoder (two layer
656 perceptron with GELU nonlinearity) function $f_\theta : \mathbb{R}^g \rightarrow \mathbb{R}^{192}$ which maps the gene expression into
657 embedding space. We likewise construct and apply the function $g_\theta : \mathbb{R}^C \rightarrow \mathbb{R}^{192}$ to map one-hot
658 encoded cell type labels to embedding space, here a simple lookup into a learnable embedding table.
659 These representations are concatenated into a single 384-dimensional representation. We apply these
660 transformations to each cell in the neighborhood, not including the reference cell. We note that at this
point, all operations have been per-cell without interactions. In addition to these cell tokens, we also

661 instantiate for each neighborhood a register token which we use to accumulate global information across
662 the neighborhood. We refer to this token as a <cls>-like token in keeping with previous literature.

663 We then apply a transformer encoder to the cells, only allowing cells within the same
664 neighborhood and their <cls>-like tokens to attend to each other. We use 8 attention heads with GELU
665 activations and layer norm prior to attention and MLP projection. We note that including a bias term in the
666 key, query, and value MLPs is important to stabilize training, while not noting any significant differences in
667 models fit with and without bias terms for the rest of the encoder and decoder layer MLPs. Following the
668 transformer encoder, we use attention pooling to aggregate the cell and <cls> representations for each
669 neighborhood into a single token with embedding dimension 384. We refer to these as the neighborhood
670 representations.

671 We then instantiate a new token from each reference cell that is a learned embedding for each
672 cell type (separately from the encoder cell type embedding). These are concatenated to the neighborhood
673 representations. We then apply a transformer decoder to the tokens, allowing only the neighborhood
674 token and masked cell embedding to attend to each other if they are from the same cellular
675 neighborhood. This decoder embedding dimension was 384 with 8 attention heads.

676 During training, we extract only the masked reference cell tokens. We then use separate linear
677 projections to output the mean, dispersion, scale, and zero inflation logit parameters for zero-inflated
678 negative binomial regression. We optimize the model by minimizing the log likelihood of a negative
679 binomial distribution using observed cells' MERFISH probe counts. We trained all versions of
680 CellTransformer on a system with 2 NVIDIA A6000 GPUs with effective batch size 256.

681
682 **Spatial domain detection**
683 Once trained, we apply CellTransformer to a given dataset and instead of extracting reference cell tokens
684 we extract the neighborhood representation. We then cluster this representation using k -means. We use
685 the cuml library to perform this operation on GPU (`cuml.KMeans`), with arguments `n_init=3`,
686 `oversampling_factor=3`, and `max_iter=1000`.
687

688 **Optional smoothing of embeddings**
689 We observe spatial domains are spatially smooth. However in the case that there is a high-frequency
690 signal that the end-user would like to filter, we optionally introduce a step prior to k -means where we
691 smooth the embeddings using a Gaussian filter. For all comparisons except those in **Supplementary**
692 **Figure 12**, smoothing was performed with a Gaussian filter with 40 micron full-width at half maxima
693 (sigma of 12.01 microns).
694

695 **Model fitting on the Allen 1 dataset**
696 We used an 80%-20% train-test split proportion (random splitting across the entire dataset) and the
697 ADAM optimizer over 40 epochs. We perform a linear warmup for 500 steps to a peak learning rate of
698 0.001 and use an inverse-square root learning rate scheduler to decay the learning rate continuously. We
699 use a weight decay value of 0.00005 which we do not warm up.
700

701 **Model fitting on Zhuang datasets**
702 We perform training from scratch without transfer. We trained for 40 epochs with the same settings as for
703 the Allen 1 with the exception of adapting projections to 1129 genes instead of 500.
704

705 **Computation of stability criterion**
706 We follow Wu et al. (2016) in using an Amari-type distance to compare clustering solutions. Briefly, we
707 compute several replicates (20 in this work) of k -means at a given choice of k with different random seed
708 as $D_{k,i}$, with i indexing the different centroids for a given solution. We then measure the stability of a

709 given choice of k by comparing the similarity of all pairs of \mathbf{D}_k . Define \mathbf{C} the Pearson correlation matrix
710 between pairs \mathbf{D} and \mathbf{D}' . Then we use this dissimilarity metric:
711

$$712 \text{diss}(\mathbf{D}, \mathbf{D}') = \frac{1}{2K} \left(2K - \sum_{i=1}^K \max_{1 \leq k \leq K} \mathbf{C}_{kj} - \sum_{j=1}^K \max_{1 \leq k \leq K} \mathbf{C}_{kj} \right).$$

713 to identify the choice of k which is most stable.
714

716 **Regional matching with CCF computation**

717 To quantify overall similarity of regions extracted using CellTransformer with CCF, we first extract cell
718 type composition vectors for each region at a given level of the hierarchy. For all comparisons in **Figure 2**,
719 we use the subclass level (338 cell types), resulting in k -region by 338 matrices. For each region
720 derived from one of the tested models, we compute two quantities: the best match (maximum value of
721 Pearson correlation, non-exclusively) to any CCF (**Figure 2d**) or an exclusive match (using the linear sum
722 assignment algorithm) to pair the regions from either set one-to-one (**Figure 2e**). We then computed the
723 average Pearson correlation across the paired matches as the metric. We use `scipy.optimize`'s
724 implementation to solve the linear sum problem.
725

726 **CellCharter**

727 To run CellCharter we first generated scVI embeddings using the default settings for depth and width of
728 the network and with the tissue section labels as conditional batch variables. We trained for 50 epochs
729 using the `early_stopping=True` setting. We then aggregated across 3 (default settings), 6, 9 layers
730 using the `cellcharter.gr.aggregate_neighbors` function. We then applied CellCharter's
731 Gaussian mixture model implementation at various choices of the number of Gaussians. We could not
732 run the mixture model with our hardware (A6000, 48GB GPU memory) for more than 9 layers, which was
733 also the number which produced the highest correspondence with CCF and is reported in **Figure 2e**.
734

735 **SPIRAL**

736 To run SPIRAL we generated edge sets for 40um, 85um, and 170um neighborhood radii. SPIRAL
737 requires supervision on single-cell types so for this we use the subclass cell type levels. We trained
738 models across neighborhood sizes for 1 epoch and then chose the neighborhood size with best
739 performance (170um) and trained this model to saturation (10 epochs). SPIRAL uses four objective
740 functions so to assess saturation we averaged them. We note that SPIRAL does not use a training and
741 testing set split in their training, making it difficult to assess an optimal stopping point. For the $k=354$ and
742 $k=670$ domain discovery analyses the SPIRAL clustering pipeline produced an out-of-memory error and
743 we instead used our own pipeline with k -means on SPIRAL embeddings.
744

745 **Nearest-neighbor smoothness computation**

746 To quantify smoothness of the spatial domains, we use a nearest-neighbor approach. We extract
747 approximate spatial neighbors for each cell using `cuml.NearestNeighbor` with 100 neighbors,
748 restricting neighbors to be within the same tissue section. For a given domain set, either from CCF,
749 CellTransformer, or CellCharter, we extract the spatial domain label of the given cell and count the
750 proportion of times that cell is observed in the 100 neighbors. These proportions are averaged across all
751 cells and tissues.
752

753 **Linear probing experiments**

754 We extract neighborhood representations for each of the cells in the Zhuang lab datasets. First, we
755 regress these embeddings on the (x, y, z) coordinates. We then computed the absolute prediction error in
756 terms of the coordinates and then reported the average. We also fit a multi-class logistic regression using
757 the mouse donor identity. For the logistic regression we use `cuml.LogisticRegression` with default
758 settings in `cuml`. For the cell position regression we fit simple least squares using PyTorch via QR
759 decomposition.

760

761 **Quantification of spatial contribution to gene expression**

762 We interpret the accuracy of the gene expression predictions for a given cell as an index of correlation of
763 an instance of a particular cell type with its surrounding neighbors. To do this, we compute a simple
764 baseline model which predicts average gene expression (computed across the entire Allen Institute for
765 Brain Science mouse dataset) for each cell. We compute the average Pearson correlation for each
766 instance of a given cell type and average across instances to obtain an average Pearson correlation. We
767 then compute a Pearson correlation between each cell's observed gene expression and the
768 `CellTransformer` predictions, averaging similarly across instances of a given cell type. The difference
769 between the baseline and model predictions is displayed, per cell type, and grouped across
770 neurotransmitter types in **Supplementary Figure 12**.

771

772 **Zhuang lab dataset per-animal CCF comparison**

773 We contrast two methods of extracting spatial domains from the four animals in the Zhuang lab dataset¹¹.
774 We first fix k , the number of desired spatial domains. Then we fit one k -means model on all of the
775 neighborhood embeddings for all four (Zhuang 1, 2, 3, and 4) mice together. We also fit a k -means model
776 to the embeddings of the mice separately. We then compute the similarity of these region sets using the
777 same method used to quantify differences between `CellCharter` and `CellTransformer` by comparing their
778 regional cell type composition vectors.

779

780 **Slide-seqV2 analysis**

781 Initial results with a direct transfer of hyperparameters to the Langlieb et al. (2023) dataset¹² did not
782 produce spatially coherent domains. We therefore implemented two quality control procedures on the raw
783 data. After filtering for coding genes and non-mitochondrial genes, we additionally used only genes that
784 were expressed in >10% of cells in the dataset. At the cellular level we identified cells with >20%
785 mitochondrial genes and those within the 10th percentile of read depth across each section. We also
786 removed these cells. This left 5,019 genes and 4,783,456 cells. We noted that a successful segmentation
787 in the Langlieb et al. (2023) dataset required a larger model than the MERFISH ones, using 10 encoder
788 layers rather than 4, which we attributed to the 10X higher number of genes in this dataset versus the 500
789 in the ABC-MWB Allen 1 dataset. We used a neighborhood size of 50um to reduce memory footprint,
790 reasoning the higher cell density in this dataset and higher number of genes would provide enough
791 information for representational richness.

792

793 **Other software**

794 Principal software used in this work includes PyTorch⁴¹, numpy⁴², scikit-learn⁴³, scipy⁴⁴, scanpy⁴⁵, `cuml`⁴⁶,
795 matplotlib⁴⁷, and seaborn⁴⁸.

796

797 **Code and data availability**

798 Code will be publicly available at <https://github.com/abbasilab/celltransformer>. All data used for this
799 publication is available either at the Allen Institute ABC-MWB data portal ([https://portal.brain-
map.org/atlas-and-data/bkp/abc-atlas](https://portal.brain-
800 map.org/atlas-and-data/bkp/abc-atlas)) or the CZI cellxgene portal
801 (<https://cellxgene.cziscience.com/datasets>).

802 Acknowledgments

803 The authors would like to acknowledge Forrest Collman for help with Neuroglancer data visualization
804 of spatial domains. We acknowledge Patrick Xian for helpful discussions. AJL acknowledges
805 Katharine Z. Yu, Chang Kim, Parker Grosjean, Lee Rao, and Tom Nowakowski for feedback on
806 neuroscientific and machine learning contexts of paper. RA, BT, AJL, and HZ would like to
807 acknowledge support from the Weill Neurohub through the Weill Neurohub's Next Great Ideas
808 Award. RA would like to acknowledge support from the National Institute of Mental Health of the
809 National Institutes of Health under award number RF1MH128672 and Sandler Program for
810 Breakthrough Biomedical Research, which is partially funded by the Sandler Foundation. AJL
811 acknowledges support from the UCSF Discovery Fellowship and the UCSF Graduate Research
812 Mentorship Fellowship. Sharing of Allen Institute for Brain Science data through Allen Brain Cell
813 Atlas (and related tools) and registration of the AIBS MERFISH brain to the CCFv3 was funded
814 through 1U24MH130918-01 to L.N.

815 References

- 816 1. Yao, Z. *et al.* A high-resolution transcriptomic and spatial atlas of cell types in the whole mouse brain.
817 *Nature* **624**, 317–332 (2023).
- 818 2. Consortium, T. Mic. *et al.* Functional connectomics spanning multiple areas of mouse visual cortex.
819 2021.07.28.454025 Preprint at <https://doi.org/10.1101/2021.07.28.454025> (2023).
- 820 3. Chen, X. *et al.* Whole-cortex *in situ* sequencing reveals input-dependent area identity. *Nature* 1–10
821 (2024) doi:10.1038/s41586-024-07221-6.
- 822 4. Haviv, D. *et al.* The covariance environment defines cellular niches for spatial inference. *Nat.*
823 *Biotechnol.* 1–12 (2024) doi:10.1038/s41587-024-02193-4.
- 824 5. Hu, J. *et al.* SpaGCN: Integrating gene expression, spatial location and histology to identify spatial
825 domains and spatially variable genes by graph convolutional network. *Nat. Methods* **18**, 1342–1351
826 (2021).
- 827 6. Zhou, X., Dong, K. & Zhang, S. Integrating spatial transcriptomics data across different conditions,
828 technologies and developmental stages. *Nat. Comput. Sci.* **3**, 894–906 (2023).
- 829 7. Townes, F. W. & Engelhardt, B. E. Nonnegative spatial factorization applied to spatial genomics. *Nat.*
830 *Methods* 1–10 (2022) doi:10.1038/s41592-022-01687-w.
- 831 8. Birk, S. *et al.* Quantitative characterization of cell niches in spatial atlases. 2024.02.21.581428
832 Preprint at <https://doi.org/10.1101/2024.02.21.581428> (2024).

833 9. Shi, H. *et al.* Spatial atlas of the mouse central nervous system at molecular resolution. *Nature* **622**,
834 552–561 (2023).

835 10. Wang, Q. *et al.* The Allen Mouse Brain Common Coordinate Framework: A 3D Reference Atlas. *Cell*
836 **181**, 936–953.e20 (2020).

837 11. Zhang, M. *et al.* Molecularly defined and spatially resolved cell atlas of the whole mouse brain.
838 *Nature* **624**, 343–354 (2023).

839 12. Langlieb, J. *et al.* The molecular cytoarchitecture of the adult mouse brain. *Nature* **624**, 333–342
840 (2023).

841 13. Veličković, P. *et al.* Graph Attention Networks. Preprint at <https://doi.org/10.48550/arXiv.1710.10903>
842 (2018).

843 14. Fischer, D. S., Schaar, A. C. & Theis, F. J. Modeling intercellular communication in tissues using
844 spatial graphs of cells. *Nat. Biotechnol.* **41**, 332–336 (2023).

845 15. Yao, Z. *et al.* A taxonomy of transcriptomic cell types across the isocortex and hippocampal
846 formation. *Cell* **184**, 3222–3241.e26 (2021).

847 16. Bakken, T. E. *et al.* Comparative cellular analysis of motor cortex in human, marmoset and mouse.
848 *Nature* **598**, 111–119 (2021).

849 17. Hintiryan, H. *et al.* The mouse cortico-striatal projectome. *Nat. Neurosci.* **19**, 1100–1114 (2016).

850 18. Varrone, M., Tavernari, D., Santamaría-Martínez, A., Walsh, L. A. & Ciriello, G. CellCharter reveals
851 spatial cell niches associated with tissue remodeling and cell plasticity. *Nat. Genet.* **56**, 74–84 (2024).

852 19. Guo, T. *et al.* SPIRAL: integrating and aligning spatially resolved transcriptomics data across different
853 experiments, conditions, and technologies. *Genome Biol.* **24**, 241 (2023).

854 20. Zhang, X., Wang, X., Shivashankar, G. V. & Uhler, C. Graph-based autoencoder integrates spatial
855 transcriptomics with chromatin images and identifies joint biomarkers for Alzheimer's disease. *Nat.*
856 *Commun.* **13**, 7480 (2022).

857 21. Dong, K. & Zhang, S. Deciphering spatial domains from spatially resolved transcriptomics with an
858 adaptive graph attention auto-encoder. *Nat. Commun.* **13**, 1739 (2022).

859 22. Spatially informed clustering, integration, and deconvolution of spatial transcriptomics with GraphST |
860 Nature Communications. <https://www.nature.com/articles/s41467-023-36796-3>.

861 23. Wu, S. *et al.* Stability-driven nonnegative matrix factorization to interpret spatial gene expression and
862 build local gene networks. *Proc. Natl. Acad. Sci.* **113**, 4290–4295 (2016).

863 24. Bienkowski, M. S. *et al.* Integration of gene expression and brain-wide connectivity reveals the
864 multiscale organization of mouse hippocampal networks. *Nat. Neurosci.* **21**, 1628–1643 (2018).

865 25. Cembrowski, M. S. *et al.* The subiculum is a patchwork of discrete subregions. *eLife* **7**, e37701
866 (2018).

867 26. Bienkowski, M. S. *et al.* Homologous laminar organization of the mouse and human subiculum. *Sci.
868 Rep.* **11**, 3729 (2021).

869 27. Ding, S.-L. *et al.* Distinct Transcriptomic Cell Types and Neural Circuits of the Subiculum and
870 Prosubiculum along the Dorsal-Ventral Axis. *Cell Rep.* **31**, 107648 (2020).

871 28. Cang, J., Chen, C., Li, C. & Liu, Y. Genetically defined neuron types underlying visuomotor
872 transformation in the superior colliculus. *Nat. Rev. Neurosci.* **25**, 726–739 (2024).

873 29. Inagaki, H. K. *et al.* A midbrain-thalamus-cortex circuit reorganizes cortical dynamics to initiate
874 movement. *Cell* **185**, 1065-1081.e23 (2022).

875 30. McElvain, L. E. *et al.* Specific populations of basal ganglia output neurons target distinct brain stem
876 areas while collateralizing throughout the diencephalon. *Neuron* **109**, 1721-1738.e4 (2021).

877 31. Benavidez, N. L. *et al.* Organization of the inputs and outputs of the mouse superior colliculus. *Nat.
878 Commun.* **12**, 4004 (2021).

879 32. Cable, D. M. *et al.* Robust decomposition of cell type mixtures in spatial transcriptomics. *Nat.
880 Biotechnol.* **40**, 517–526 (2022).

881 33. Brbić, M. *et al.* Annotation of spatially resolved single-cell data with STELLAR. *Nat. Methods* **19**,
882 1411–1418 (2022).

883 34. Maskey, S. *et al.* Generalized Laplacian Positional Encoding for Graph Representation Learning.
884 *arXiv.org* <https://arxiv.org/abs/2210.15956v2> (2022).

885 35. Burgess, J., Nirschl, J. J., Zanellati, M.-C., Cohen, S. & Yeung, S. Learning orientation-invariant
886 representations enables accurate and robust morphologic profiling of cells and organelles.
887 2022.12.08.519671 Preprint at <https://doi.org/10.1101/2022.12.08.519671> (2022).

888 36. Connell, W., Khan, U. & Keiser, M. J. A single-cell gene expression language model. *arXiv.org*

889 https://arxiv.org/abs/2210.14330v1 (2022).

890 37. Laboratory, I. B. *et al.* A Brain-Wide Map of Neural Activity during Complex Behaviour.

891 2023.07.04.547681 Preprint at <https://doi.org/10.1101/2023.07.04.547681> (2023).

892 38. de Vries, S. E., Siegle, J. H. & Koch, C. Sharing neurophysiology data from the Allen Brain

893 Observatory. *eLife* **12**, e85550 (2023).

894 39. Oh, S. W. *et al.* A mesoscale connectome of the mouse brain. *Nature* **508**, 207–214 (2014).

895 40. Hike, D. *et al.* High-resolution awake mouse fMRI at 14 Tesla. *eLife* **13**, (2024).

896 41. Paszke, A. *et al.* PyTorch: An Imperative Style, High-Performance Deep Learning Library. Preprint at

897 <https://doi.org/10.48550/arXiv.1912.01703> (2019).

898 42. Array programming with NumPy | Nature. <https://www.nature.com/articles/s41586-020-2649-2>.

899 43. Pedregosa, F. *et al.* Scikit-learn: Machine Learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830

900 (2011).

901 44. Virtanen, P. *et al.* SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nat.*

902 *Methods* **17**, 261–272 (2020).

903 45. Wolf, F. A., Angerer, P. & Theis, F. J. SCANPY: large-scale single-cell gene expression data

904 analysis. *Genome Biol.* **19**, 15 (2018).

905 46. Raschka, S., Patterson, J. & Nolet, C. Machine Learning in Python: Main developments and

906 technology trends in data science, machine learning, and artificial intelligence. *ArXiv Prepr.*

907 *ArXiv200204803* (2020).

908 47. Hunter, J. D. Matplotlib: A 2D graphics environment. *Comput. Sci. Eng.* **9**, 90–95 (2007).

909 48. Waskom, M. L. seaborn: statistical data visualization. *J. Open Source Softw.* **6**, 3021 (2021).

910 49. Ollivier, M. *et al.* Ctrym-positive striatal astrocytes gate perseverative behaviour. *Nature* **627**, 358–366

911 (2024).

912 **Supplementary Note 1: Effect of smoothing and analysis on** 913 **striatal glial populations**

914 When scaling up the number of regions past 500 in the Allen 1 dataset, we observed that almost
915 all spatial clusters were spatially smooth except for a recurring pattern in the striatum. We plot
916 (**Supplementary Figure 11a-b**) six sequential sections where we identified an irregular (which we define
917 here as a broadly non-convex shape that does not form relatively singular connected component) pattern
918 of cells in the striatum and only the striatum (note the spatial uniformity of areas surrounding striatum in
919 cortex and endopiriform area, nucleus accumbens etc.). We identified cells in these areas and found they
920 were mostly non-neuronal, with astrocytic types (such as 1163 **Astro-TE NN_3, Supplementary Figure**
921 **11**) forming a large proportion of cells.

922 We sought to understand whether these spatial clusters might be biologically relevant or
923 somehow related to noise. A recent paper, Ollivier et al. (2024)⁴⁹ identified a novel population of *Crym*⁺
924 astrocytes in a similar spatial distribution as observed in our regions, specifically in a dorsoventral and
925 lateromedial distribution (see **Supplementary Figure 11e** for reproduction from Ollivier et al., granted
926 with permission). As *Crym* was included in the MERFISH panel, we quantified *Crym* expression in
927 astrocytes within these areas, finding that all but two of these spatially irregular domains had very high
928 levels of *Crym* expression. Notably, the two groups with lower expression, spatial clusters 457 and 758,
929 were the most dorsolateral, and are distributed where *Crym*⁺ astrocytes were not observed in Ollivier et
930 al. We reasoned that these spatial clusters may have biological relevance.

931 However, to simplify downstream analyses and conform with neuroanatomical conventions we
932 applied a simple smoothing operation (see **Methods**), which removed this spatial cluster in successive
933 clustering operations. We used a very small smoothing window (12 micron sigma, or 40 micron full-width
934 at half-maxima) and found the order of ranked methods and their relative performance changes were not
935 significantly affected.

936 **Supplementary Note 2: Interpreting the CellTransformer** 937 **objective as a measure of spatial dependence**

938 One interpretation of the CellTransformer architecture is learning two representations of cellular
939 gene expression. The first (the learned embedding for each cell type) is unconditional on the spatial
940 neighborhood information. The second is one that is conditional on the spatial neighborhood information
941 learned in the encoder portion of the network and parameterized as a residual update. This residual
942 update can then conveniently be aggregated across layers and represented as a single update term on
943 the unconditional representation to produce the final output. We interpret the increase in accuracy from
944 neighborhood-conditional gene expression prediction as an index of spatial dependence.
945 A trivial or poorly fit optimization solution would produce a small value of this index. A similar idea has
946 been previously presented in a number of works, most recently the NCEM approach¹⁴.

947 When analyzing the Allen 1 dataset we observe increases in the predictive accuracy (mean 0.10
948 +/- 0.0701 in correlation, averaged across cell types) across the dataset. Moreover, there are few cell
949 types for which there is a decrease in predictive accuracy, indicating that our model has nontrivially
950 learned the objective (**Supplementary Figure 12a**). Those that are poorly predicted are often only
951 present in the dataset at very low abundances. Conditioned on cell types with more than 10^2 - 10^3 cells,
952 accuracy has only a mild correlation with cellular density (**Supplementary Figure 12b**) or on number of
953 cells of given type in the dataset. Immature neurons (IMN) are the class which benefits the most from
954 conditional prediction, suggestive of their complex migratory dynamics¹. Note that we compute accuracies
955 at the subclass level (338 types). Increase in accuracy does correlate strongly with log-number of

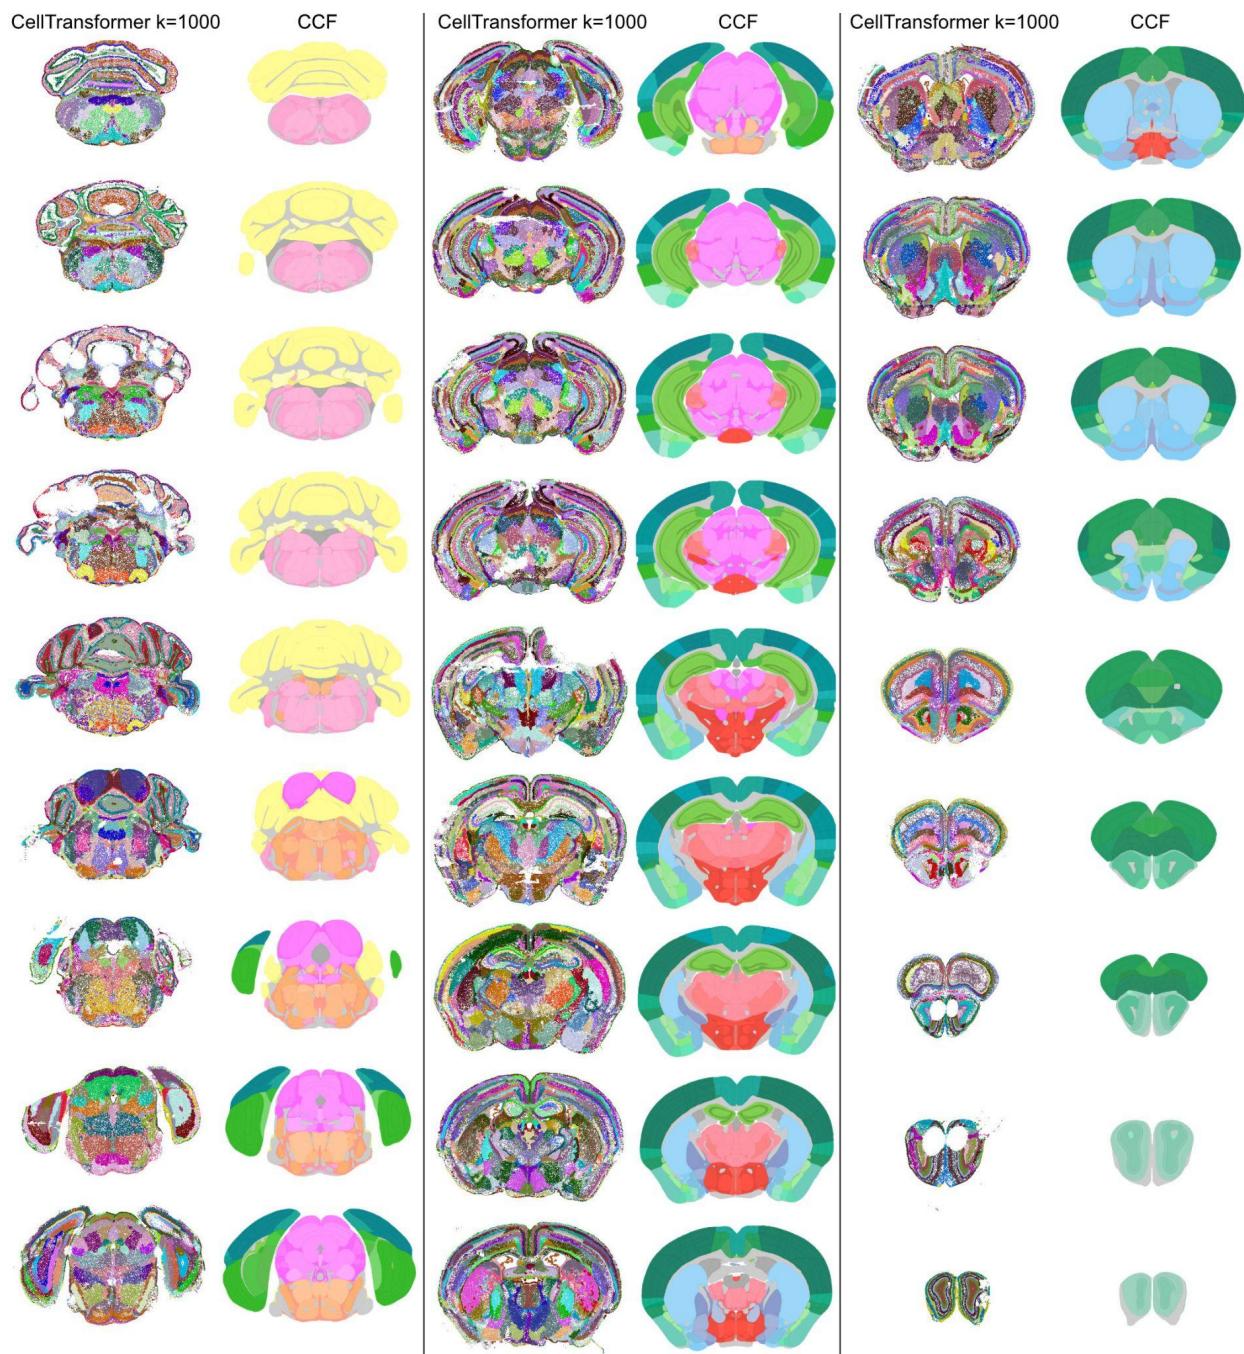
956 observations per cell type, (Pearson correlation of 0.71, **Supplementary Figure 12c**), possibly indicating
957 inefficiencies in pretraining.

958 **Supplementary Figures**

959



960
961 **Supplementary Figure 1.** CellTransformer spatial domains (left) and the corresponding CCF annotations (right)
962 organized in 3 columns for roughly half of the sections in the Allen 1 dataset¹, approximately every other section.
963 CellTransformer domains were calculated at $k=25$ clusters.





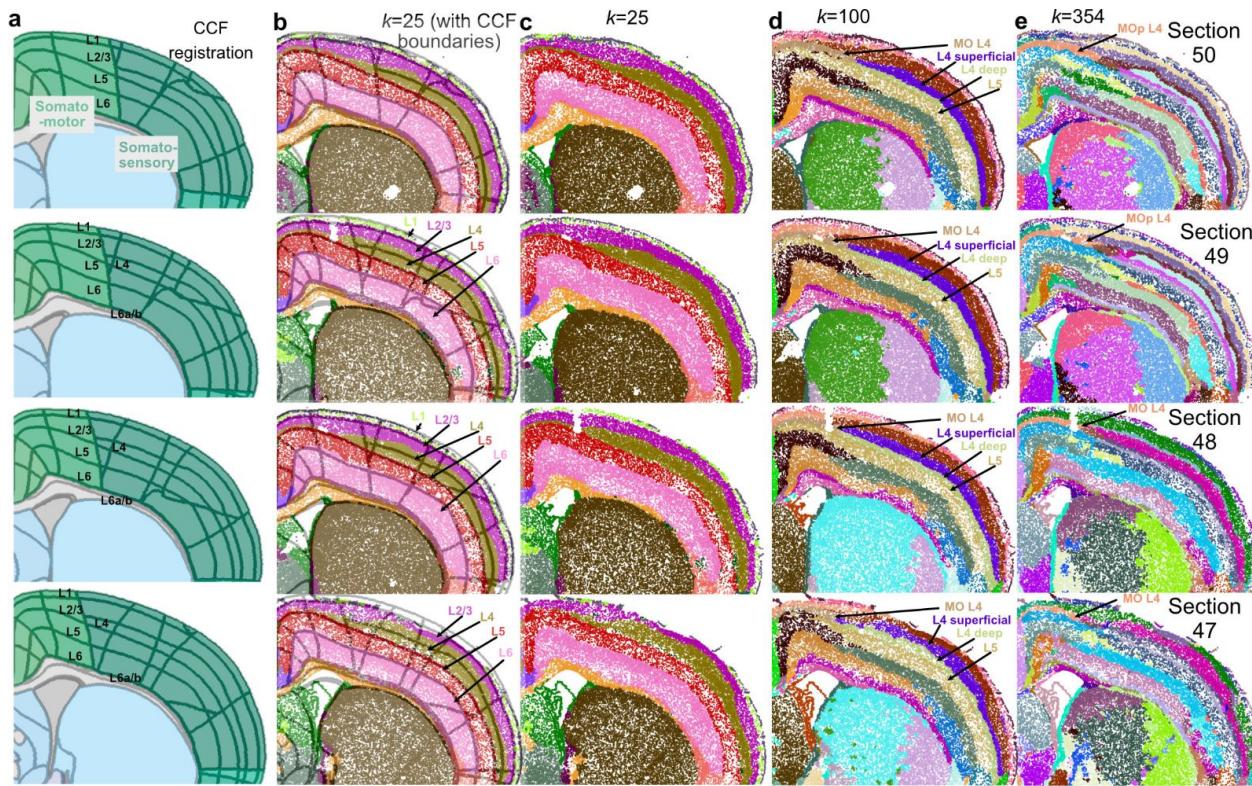
968

969

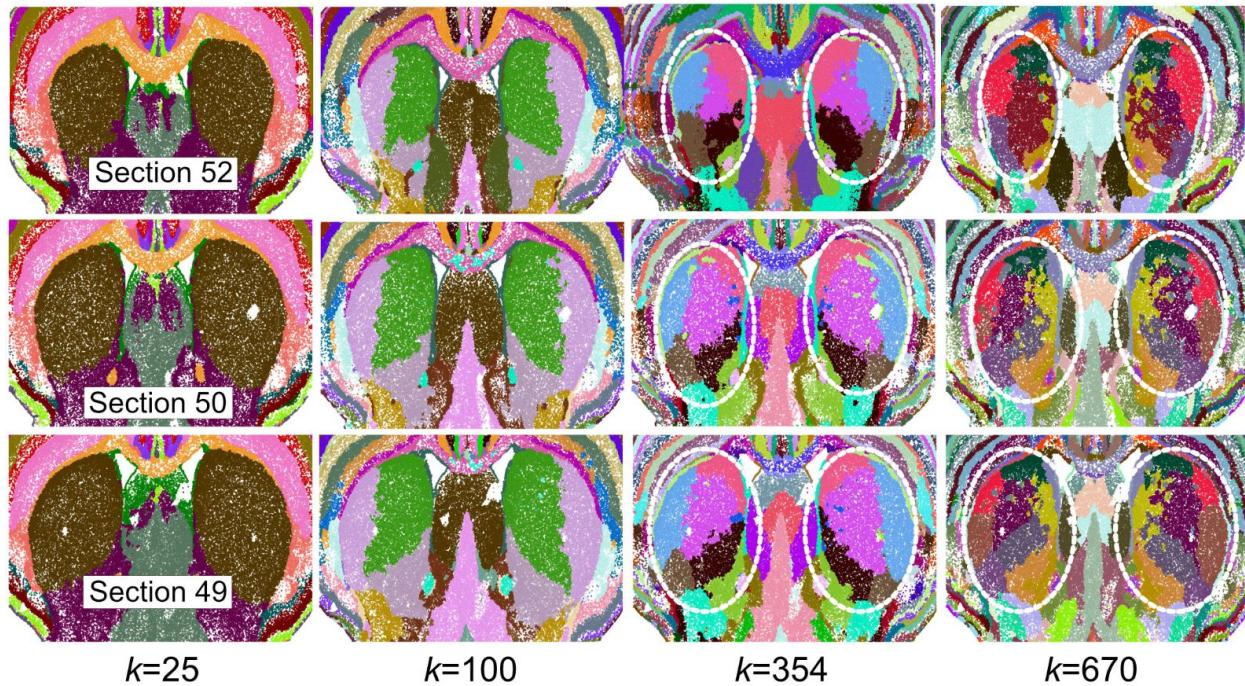
970

971

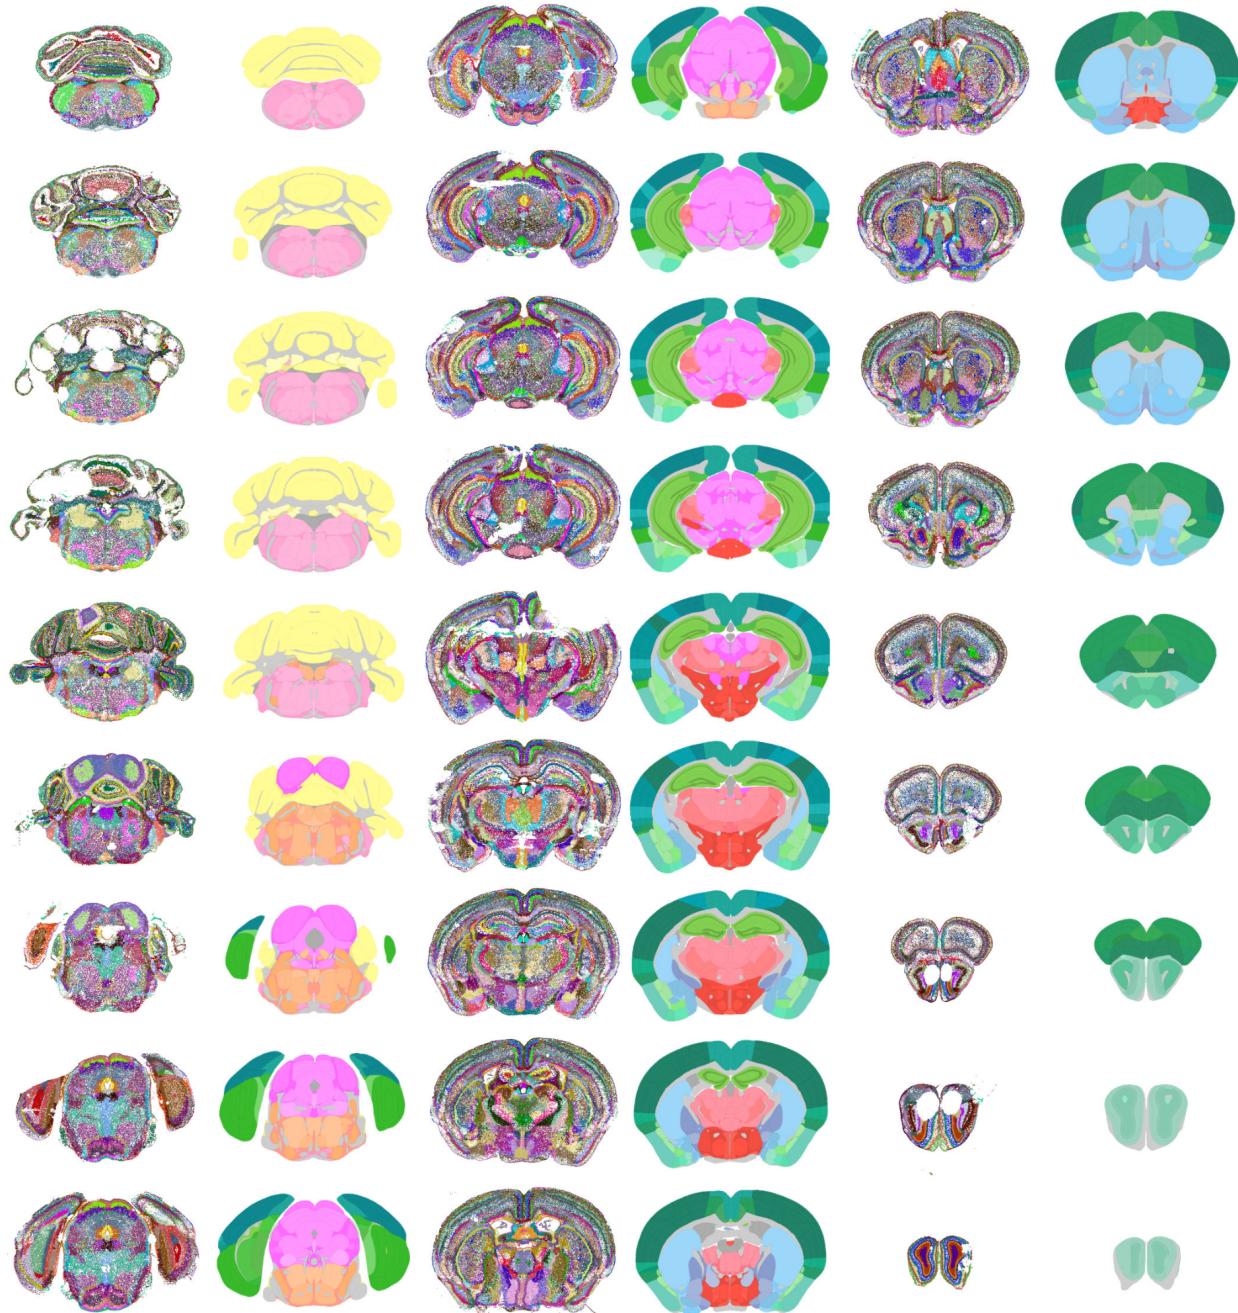
Supplementary Figure 3. CellTransformer spatial domains (left) and the corresponding CCF annotations (right) organized in 3 columns for roughly half of the sections in the Allen 1 dataset¹, approximately every other section. CellTransformer domains were calculated at $k=1300$ clusters.



972
973 **Supplementary Figure 4.** Four sequential sections of the Allen 1 dataset (200 μ m sampling interval between
974 sections¹) displayed with CellTransformer labels at varying resolution alongside CCF registration results. MO: motor
975 cortex. (a.) CCF registration of four sequential sections shown in **Figure 2**. cortical layers are marked based on CCF
976 annotations. (b.) $k=25$ spatial domains with CellTransformer shown with regional boundaries from CCF in light gray.
977 Putative cortical layers are annotated, showing CellTransformer replicates known cortical layers. (c.) 25 domains
978 shown without CCF annotations to facilitate visualization. (d.) Same sections now shown with 100 domains to help
979 show the transition from coarse (25 domains) to fine (100 domains). Sublayers of cortex are identified including layer
980 4 in motor cortex which transcriptomic studies have verified but has been difficult to identify using histological
981 approaches. (e.) 354 domain zoom in on the same sections, showing consistency of layer 4 motor cortex detection as
982 well as an anterior-posterior subdivision across motor and somatosensory cortical layers and clear distinction of
983 cortical layers that lie within motor and somatosensory areas.

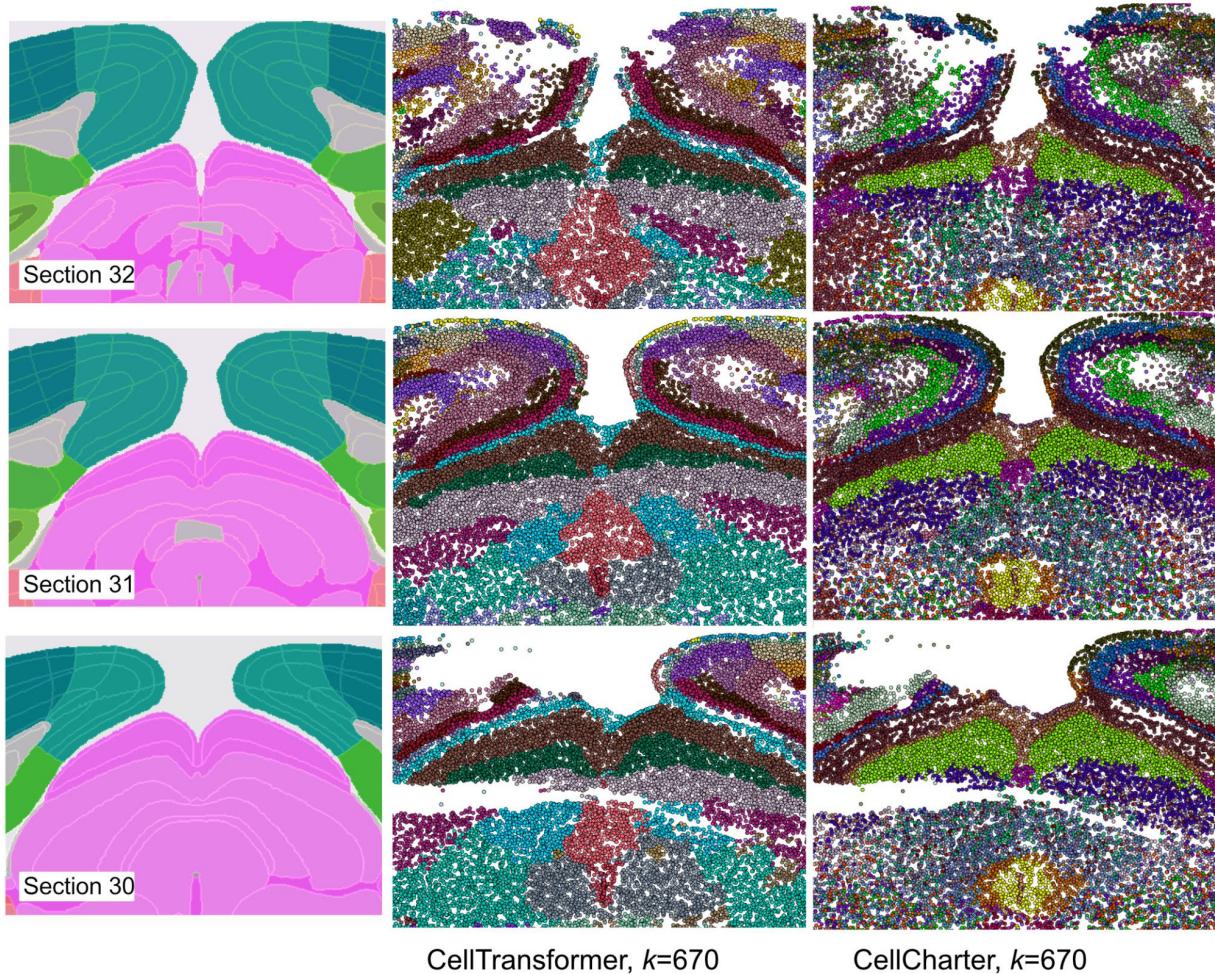


984
985 **Supplementary Figure 5.** CellTransformer domains identified in the Allen 1 dataset¹ at varied values of k , colored in
986 three sequential sections (200 μ m sampling interval between sections with consecutive numbers, top to bottom
987 corresponds to rostral to caudal). Caudoputamen is roughly highlighted by dotted circles in $k=354$ and $k=670$ to assist
988 visualization.



989
990
991
992

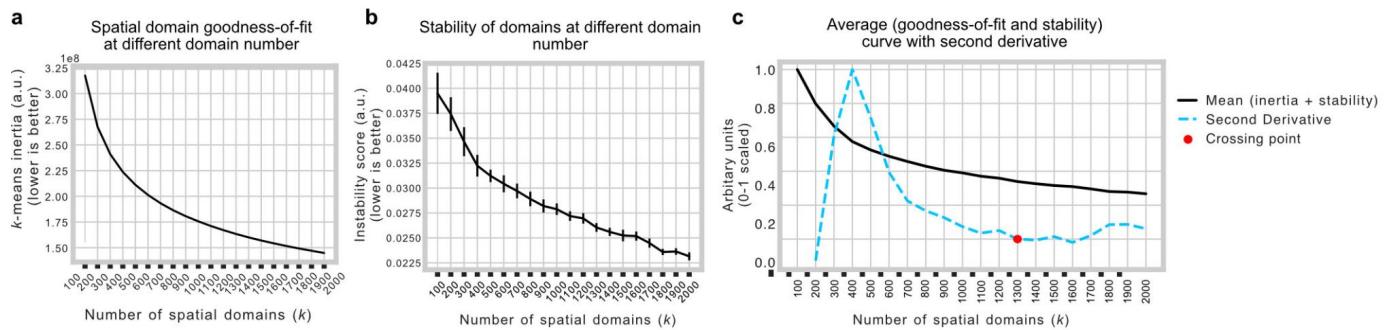
Supplementary Figure 6. CellCharter results (left) and the corresponding CCF annotations (right) organized in 3 columns for roughly half of the sections in the ABC-MWB (Allen 1¹) dataset, approximately every other section. The color labels for CellCharter correspond to its Gaussian mixture model implementation with $k=670$ clusters.



993
994
995
996
997
998
999

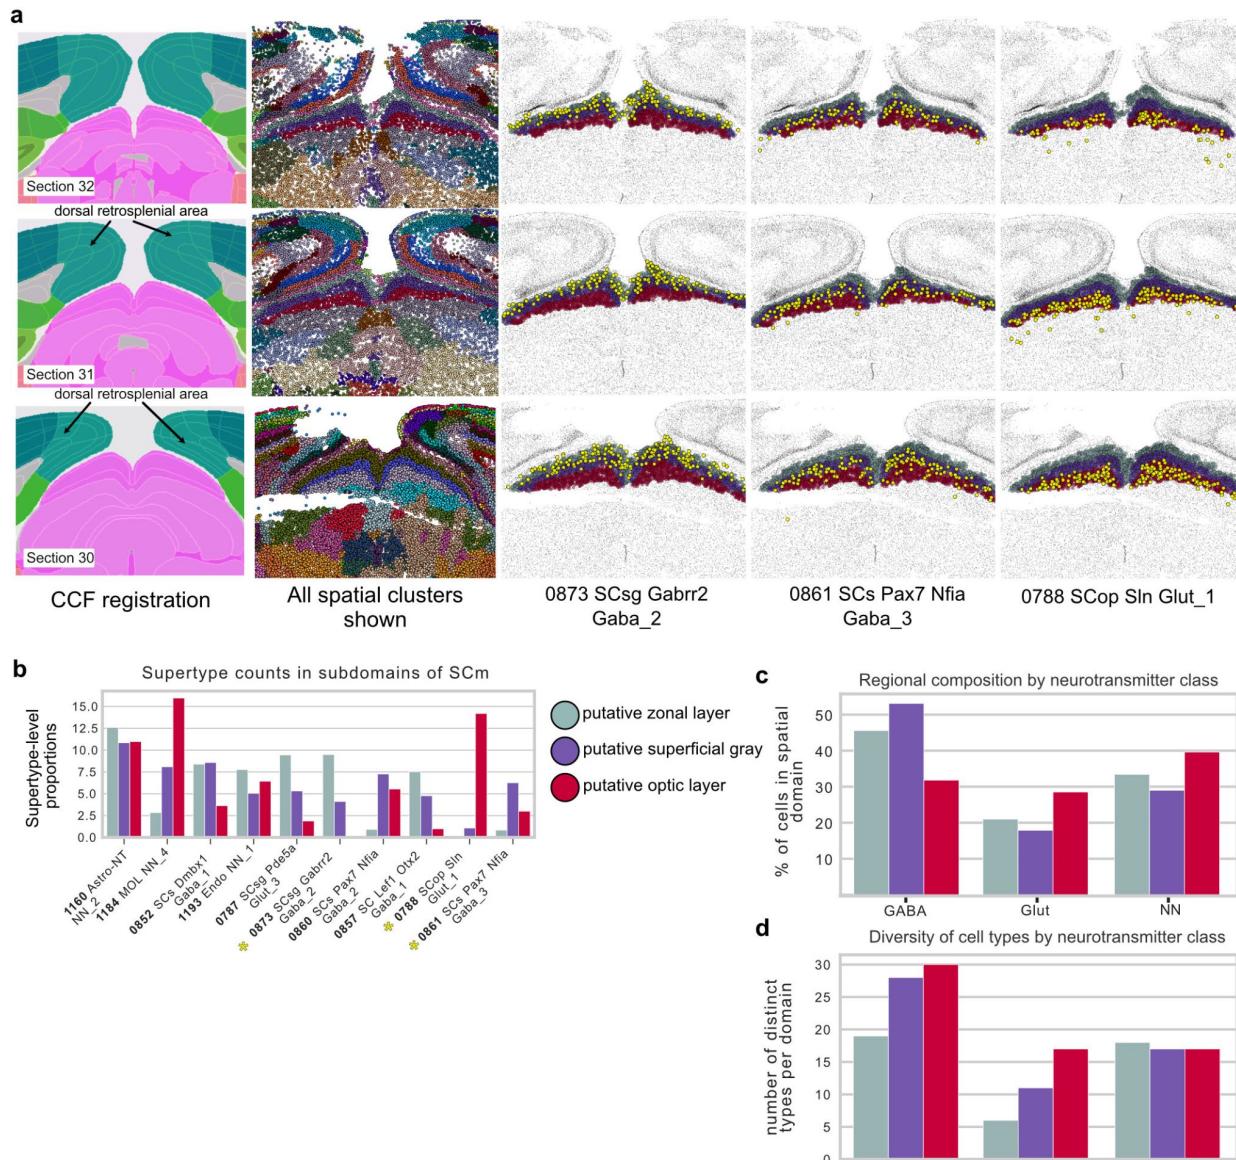
Supplementary Figure 7. Comparison of spatial domains in midbrain for CellTransformer and CellCharter discovered in the Allen 1 dataset¹. Left column shows approximate CCF registration. Middle column shows CellTransformer domains at $k=670$ and the right column shows CellCharter domains with 670 Gaussians. The general performance in outlining cortical layers is similar, however in the midbrain, even at half the number of clusters, CellCharter loses spatial coherence compared with CellTransformer. For example, CellCharter only identifies two layers of superior colliculus, whereas multiple layers are defined by CellTransformer.

1000



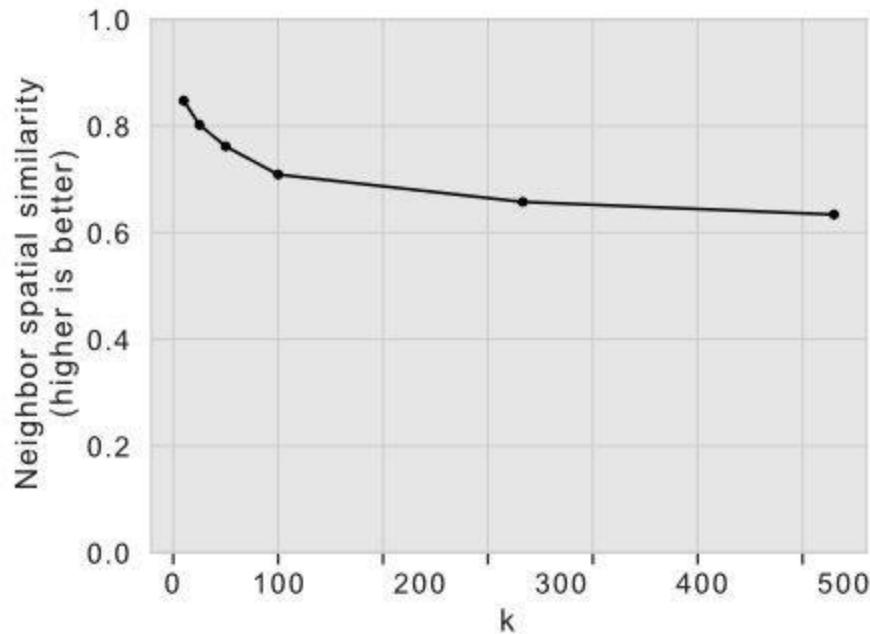
1001

1002 **Supplementary Figure 8.** Quantification of goodness-of-fit and stability of varying numbers of spatial domains. (a.)
1003 inertia (sum of squares errors for each cluster centroid) calculated for different clustering solutions when clustering
1004 embeddings generated using CellTransformer on the Allen 1 dataset¹. Error bars (standard deviation) are calculated
1005 but not visible due to scale. (b.) instability scores (see **Methods**) calculated for different clustering solutions using the
1006 Allen 1 dataset. Error bars are standard deviation. (c.) Average of inertia and stability curves (black line) and second
1007 derivative of same curve (blue dotted lines). Second derivative crossing point at $k=1300$ shown with red dot.



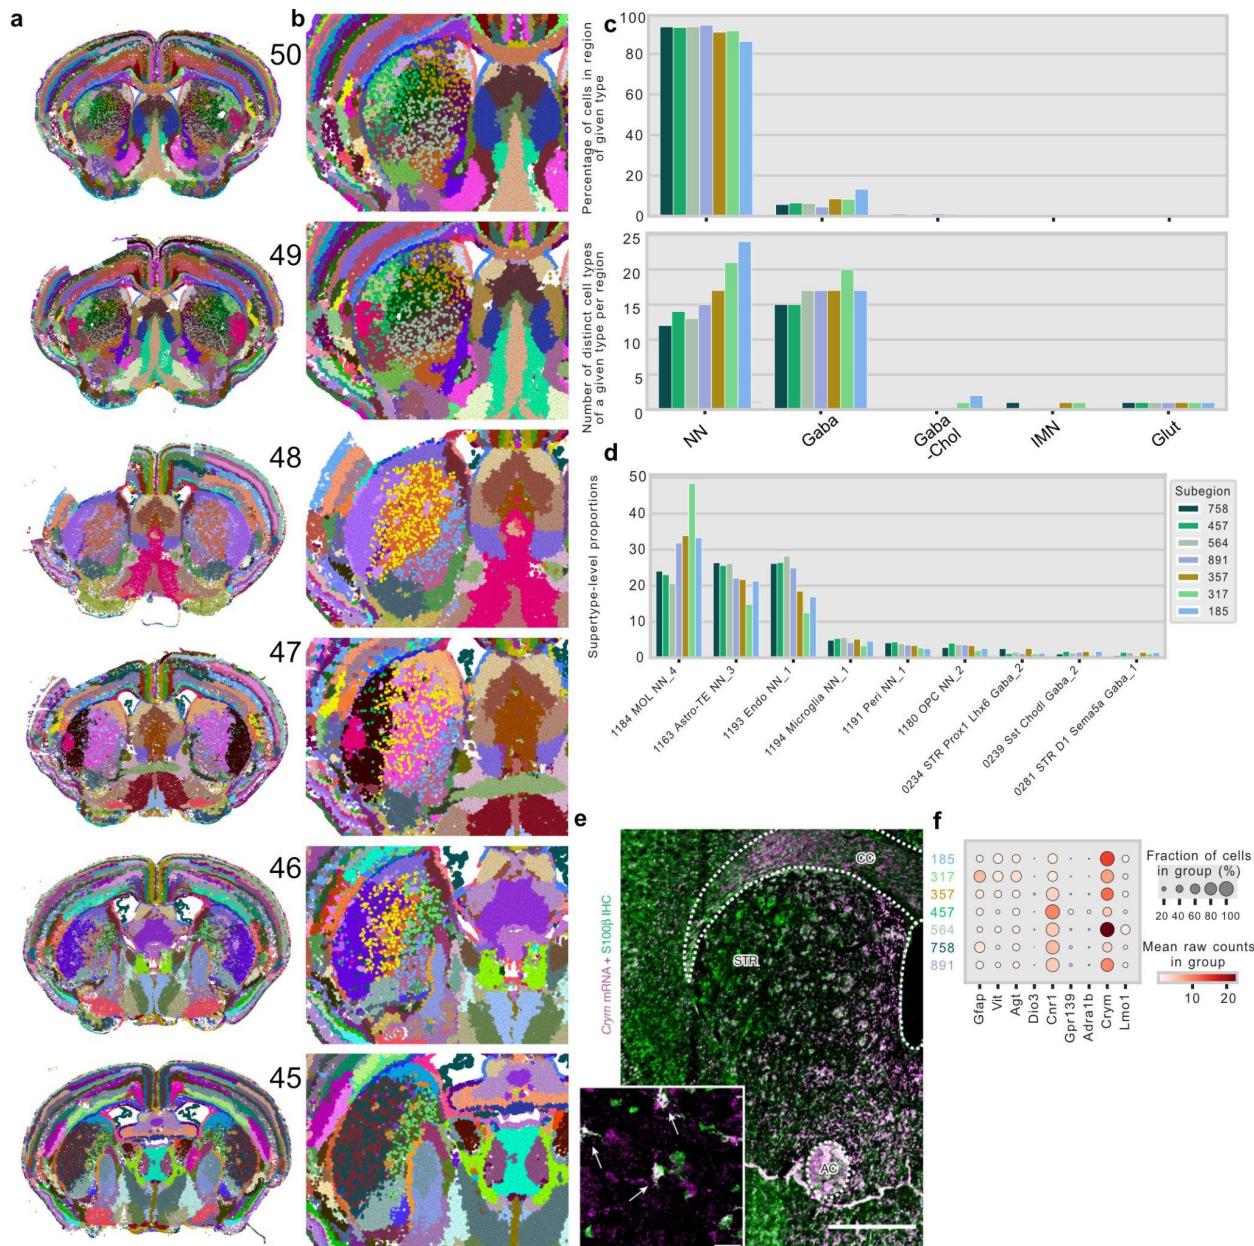
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018

Supplementary Figure 9. Comparison of best fit spatial domains from CellTransformer with layers of the superior colliculus, sensory related area. **(a.)** Sequential tissue sections (32, 31, 30, from anterior to posterior) showing in first column CCF registration and borders of relevant areas. Second column: all cells in field of view, colored by spatial domain from CellTransformer. Third column: only visualizing cells inside our putative matches for the zonal, superficial gray, and optic layers in the superior colliculus. The 0879 SCsg Pde5a Glut_1 cell type (supertype-level) in yellow. Fourth column: same as third but visualizing the 0865 SCs Pax7 Nfia Gaba_3 cell type. Fifth column: same as third and fourth but visualizing the 0882 SCop Sln Glut type. **(b.)** Bar chart of cell type abundance (as a percentage) for top ten most abundant types across the putative subregions. Cell types visualized in **(a.)** are marked with a yellow asterisk. **(c.)** Bar chart of per-region proportions of GABA-ergic and glutamatergic neurons and non-neuronal types. **(d.)** Bar chart of the number of distinct cell types at supertype level of the ABC-MWB taxonomy per domain.

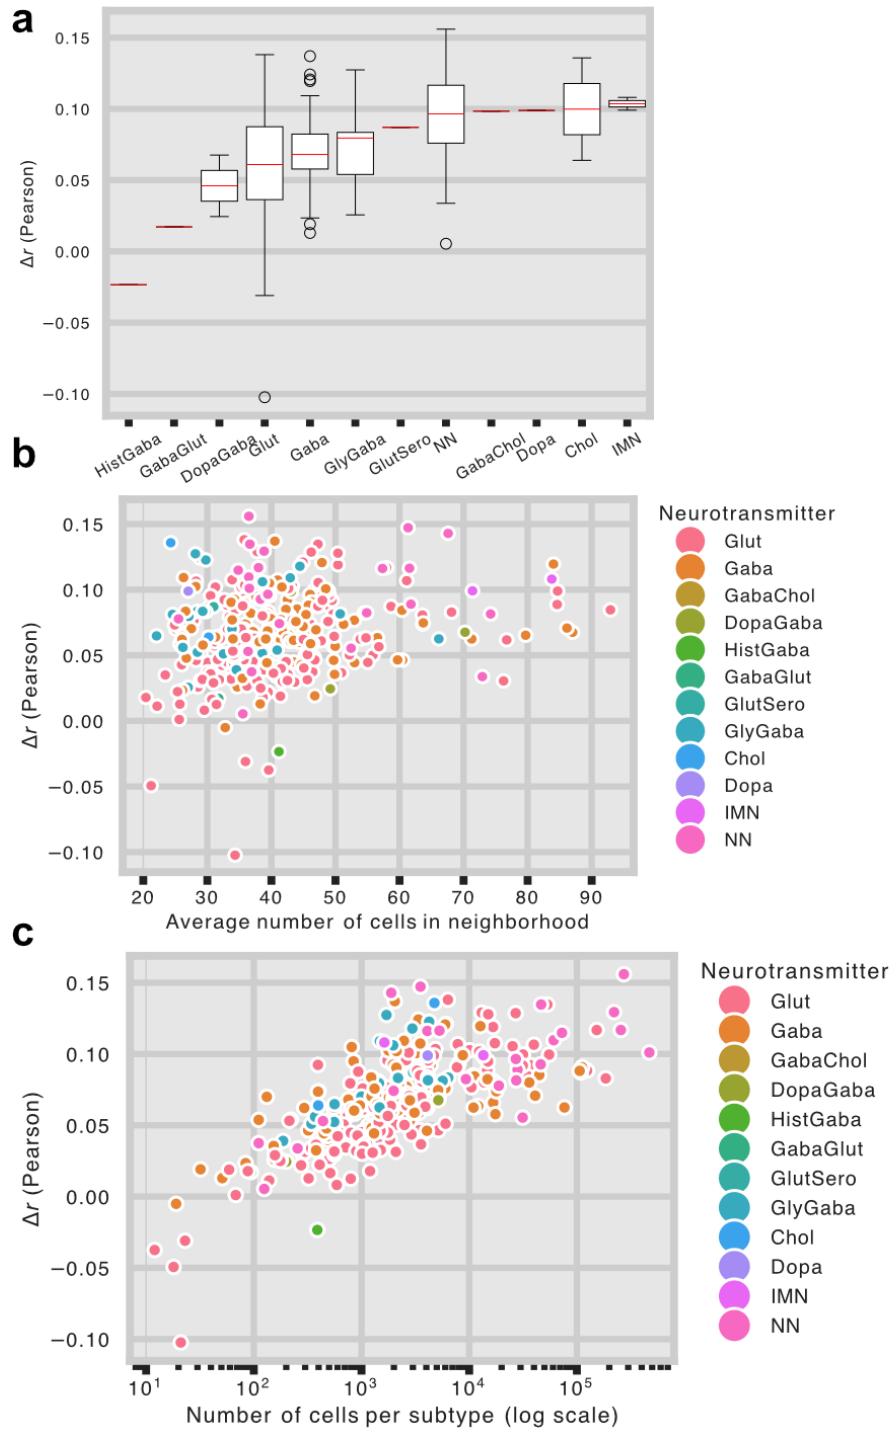


1019
1020
1021
1022
1023

Supplementary Figure 10. Results of quantitative comparison of CCF regions and CellTransformer regions at equivalent number of regions using the Zhuang 1-4 datasets¹¹. Spatial smoothness of spatial clusters as measured using a nearest-neighbors approach, computed by clustering the concatenated latent variables for neighborhoods in the Zhuang lab datasets.



1024
1025 **Supplementary Figure 11.** Representative images of spatial clustering from CellTransformer models with $k=1300$
1026 identified using the Allen 1 dataset¹. (a.) Sequential tissue sections (50 is most anterior) showing smoothness of
1027 spatial domains across and within tissue sections as well as consistent appearance of an irregular spatial pattern
1028 inside caudoputamen. (b.) Zoom in on the striatum for the same tissue sections. (c.) Plots showing percentage of cell
1029 types of different neurotransmitter for the non-uniform spatial clusters as well as the distribution of unique cell types of
1030 a given neurotransmitter type. (d.) Supertype-level counts in putative subpopulations of caudoputamen. (e.)
1031 Reproduction with permission of results from Ollivier et al. (2024). showing the distribution of *Crym* mRNA and its
1032 protein product (S100B), clearly identifying a medial population of *Crym*⁺ neurons which resembles the spatial pattern
1033 observed in clusters 758 and 457 (dorsomedial and *Crym*⁻). (f.) Dotplot of cell type expression proportions and mean
1034 counts per group (raw counts) in identified irregular caudoputamen areas.



1035
1036
1037
1038
1039
1040
1041

Supplementary Figure 12. Quantification of improved prediction accuracy as a result of CellTransformer's neighborhood-conditioned prediction. Results are over all cells in the Allen 1 dataset¹. (a.) Change in Pearson correlation from per-cell type (at subclass level) average expression. Red lines show medians per distribution. (b.) Scatterplot of increase in average Pearson correlation per subclass level cell type against average neighborhood size for reference cells of that type. c. Scatterplot of increase in average Pearson correlation per subclass level cell type vs the number of cells of that type in log scale.