

Neurotransmitter Classification from Electron Microscopy Images at Synaptic Sites in *Drosophila*

Nils Eckstein^{1,2}, Alexander S. Bates³, Michelle Du¹, Volker Hartenstein⁵, Gregory S.X.E. Jefferis^{3,4}, and Jan Funke¹

¹HHMI Janelia Research Campus, Ashburn, USA

²Institute of Neuroinformatics UZH/ETHZ, Zurich, Switzerland

³Neurobiology Division, MRC Laboratory of Molecular Biology, Cambridge, United Kingdom

⁴Department of Zoology, University of Cambridge, Cambridge, United Kingdom

⁵Department of Molecular Cell and Developmental Biology, University of California Los Angeles, Los Angeles, USA



Abstract—High-resolution electron microscopy (EM) of nervous systems enables the reconstruction of neural circuits at the level of individual synaptic connections. However, for invertebrates, such as *Drosophila melanogaster*, it has so far been unclear whether the phenotype of neurons or synapses alone is sufficient to predict specific functional properties such as neurotransmitter identity. Here, we show that in *Drosophila melanogaster* artificial convolutional neural networks can confidently predict the type of neurotransmitter released at a synaptic site from EM images alone. The network successfully discriminates between six types of neurotransmitters (GABA, glutamate, acetylcholine, serotonin, dopamine, and octopamine) with an average accuracy of 87% for individual synapses and 94% for entire neurons, assuming each neuron expresses only one neurotransmitter. This result is surprising as there are often no obvious cues in the EM images that human observers can use to predict neurotransmitter identity. We apply the proposed method to quantify whether, similar to the ventral nervous system (VNS), all hemilineages in the *Drosophila melanogaster* brain express only one fast acting transmitter within their neurons. To test this principle, we predict the neurotransmitter identity of all identified synapses in 89 hemilineages in the *Drosophila melanogaster* adult brain. While the majority of our predictions show homogeneity of fast-acting neurotransmitter identity within a single hemilineage, we identify a set of hemilineages that express two fast-acting neurotransmitters with high statistical significance. As a result, our predictions are inconsistent with the hypothesis that all neurons within a hemilineage express the same fast-acting neurotransmitter in the brain of *Drosophila melanogaster*.

1 INTRODUCTION

In recent years, advances in imaging technology enabled high resolution electron microscopy (EM) imaging of whole brain data sets (Zheng et al., 2018; Ryan et al., 2016; Cook et al., 2019; Ohya et al., 2015), opening up the possibility of generating cellular level wiring diagrams (connectomes) of nervous systems. Generating a connectome entails identifying all neurons and the synapses that connect them. Due to the size of these data sets manual tracing of all neurons and synapses is not feasible even for comparatively small organisms such as *Drosophila melanogaster*. However, recent advances in automated methods for segmenting neurons (Funke et al., 2018; Januszewski et al., 2018; Lee et al., 2019) and detecting synapses (Kreshuk et al., 2015; Staffler et al., 2017; Buhmann et al., 2019) greatly reduce the time of human involvement in these tasks and have just recently been applied

to generate the connectome for a large part of the *Drosophila melanogaster* brain (Xu et al., 2020).

However, EM data does not directly give us information about gene expression and as a result quantities such as neurotransmitter identity, while crucial to determine the function of any given synapse, are unknown for a majority of the cells in the connectome of *Drosophila melanogaster*. The action a neuron has on its downstream targets is determined by the neurotransmitters it releases into the extracellular space. Before release, neurotransmitters are packaged into different types of vesicles at synaptic sites. The so-called 'classical', fast-acting transmitters GABA, acetylcholine and glutamate are contained in small, clear vesicles, while monoamines such as dopamine, norepinephrine, octopamine and serotonin are packaged into pleomorphic clear-core or small dense-core vesicles (Goyal and Chaudhury, 2013). The large number of various neuropeptides such as cholecystokinin, galanin, neurokinin and oxytocin are contained in large dense-core vesicles. In mammals it is generally possible for humans to distinguish between different clear-core and dense-core vesicles and thus identify the neurotransmitter of any given synaptic site (Goyal and Chaudhury, 2013). In invertebrates such as *Drosophila melanogaster* it is so far unknown whether synaptic phenotype, as seen in EM, is sufficient to consistently determine neurotransmitter identity, especially different varieties of clear-core vesicles.

As a result, adding neurotransmitter identities to connectomic data requires light microscopy (LM) pipelines, in which gene transcripts or proteins involved in the pathway of interest have been made visible using fluorescent probes. Common methods for neurotransmitter detection include sequencing transcriptomics (RNAseq) (Henry et al., 2012; Konstantinides et al., 2015; Davie et al., 2018; Davis et al., 2020), immunolabeling (Hyatt and Wise, 2001) and fluorescent *in situ* hybridization (FISH) (Long et al., 2017; Meissner et al., 2019). Subsequent morphological matching of these neurons to reconstructed neurons in the EM data can then be performed using NBLAST (Costa et al., 2016),

providing neurotransmitter identity to connectome data (Bates et al., 2019b,a).

However, this approach is very difficult to scale to an entire connectome of *Drosophila melanogaster* comprising ~150,000 cells. Although imaging expression patterns for multiple neurotransmitters on a brain scale can be done in a matter of minutes to days depending on the required resolution¹ (Meissner et al., 2019), bridging the gap between LM and EM remains challenging: In addition to imaging expression patterns of neurotransmitter related proteins, it requires a well-characterized, sparse genetic driver line in order to perform accurate morphological matching to EM tracings using tools such as NBLAST (Costa et al., 2016). As a result, transmitter identity is known for only a few hundred types of neurons (Bates et al., 2019a).

Here, we show that it is possible to determine the primary neurotransmitter of a given neuron in the *Drosophila melanogaster* brain from the phenotype of its synaptic sites in EM alone. For that, we train a deep learning classifier to predict the neurotransmitter of a $640 \times 640 \times 640\text{nm}^3$ 3D EM volume with a synaptic site at the center. We find that this method is able to classify the neurotransmitter of any given synapse with 87% accuracy on average. Furthermore, we show on a large test set that the classifier generalizes across neurons with different developmental histories (*i.e.*, that derive from different 'hemilineages') and brain regions, indicating that the influence of the neurotransmitter on the phenotype of a synaptic site is largely conserved across cells. We use our method to predict the neurotransmitter identity of over 1000 neurons in 89 hemilineages with so far unknown neurotransmitter identities in the *Drosophila melanogaster* brain. In contrast to recent findings in the ventral nervous system (VNS) (Lacin et al., 2019), our results suggest that the neurotransmitter identity of neurons within hemilineages in the brain is not limited to one fast-acting transmitter.

Given that the relation of synaptic phenotype and neurotransmitter identity is not fully understood in *Drosophila melanogaster*, it is an exciting research question whether we are able to extract the set of rules and features the classifier uses to predict neurotransmitter identity. However, due to the blackbox nature of deep neural network classifiers and the complexity of synaptic sites, this is a challenging problem and beyond the scope of this paper. We plan to investigate this question in future work.

In summary, our method circumvents a major bottleneck in neurotransmitter identification, matching LM expression patterns to EM tracings, and is able to assign neurotransmitter identity to individually traced neurons in a matter of seconds. Combined with automated synapse detection methods (Buhmann et al., 2018; Kreshuk et al., 2015; Staffler et al., 2017; Buhmann et al., 2019), this opens up the possibility of generating a comprehensive neurotransmitter atlas for the connectome of *Drosophila melanogaster*.

2 METHODS

We learn a mapping $f : x \rightarrow y$, where x is a local 3D EM volume with a synaptic site at the center and y the neurotransmitter of the corresponding neuron (see Fig. 1). To this end, we need to generate a training dataset of pairs (x,y) . This involves light microscopy of genetically tagged neurons to determine their neurotransmitter expression, neuron tracing and synapse annotation in the corresponding EM dataset and matching of the LM neuron morphology to a traced neuron in EM.

2.1 Data Acquisition

We acquire the majority of neurotransmitter to neuron assignments used for training and evaluation from published reconstructions in the full adult fly brain (FAFB) dataset (Bates et al., 2020; Dolan et al., 2018; Felsenberg et al., 2018; Frechter et al., 2019; Huoviala et al., 2018; Dolan et al., 2019; Marin et al., 2020; Sayin et al., 2019; Turner-Evans et al., 2019; Zheng et al., 2018), as well as unpublished but identified neuron reconstructions offered by the FAFB community (see acknowledgements). In these studies, the authors had already linked some of their reconstructed cell types to immunohistochemical data (Aso et al., 2014; Bräcker et al., 2013; Busch et al., 2009; Davis et al., 2018; Dolan et al., 2019; Ito et al., 2013; Lai et al., 2008; Okada et al., 2009; Shinomiya et al., 2015; Tanaka et al., 2012; Wilson and Laurent, 2005). Stainings were typically performed on neurons visualized by GFP expression in a GAL4/split-GAL4 line. Dissected brains are incubated with primary antibodies (*e.g.*, anti-VGluT, anti-GABA, anti-ChAT), followed by secondary antibodies which have a fluorescent tag to visualise the primary antibody. The transcripts/proteins related to certain transmitter expressions are thus labelled across the brain and if they colocalize with the GFP signal for the GAL4/split-GAL4 line of interest, those neurons are considered to express that transmitter³. For RNA transcripts, usually the neuron's soma is examined. Other methods involve RNA sequencing and include TAPIN-seq (Davis et al., 2018). Note that individual studies often only test single transmitters and do not show negative staining. As a result, there is limited data for cotransmission of multiple neurotransmitters in a single neuron and we therefore assume no cotransmission of neurotransmitters within one neuron.

We use manually reconstructed neuron annotations from the FAFB community CATMAID⁴ (Saalfeld et al., 2009; Schneider-Mizell et al., 2016) database⁵. Synapses were annotated at presynaptic sites, defined by T-bars, vesicles and a thick dark active zone by a synaptic cleft (Prokop and Meinertzhagen, 2006). In total, the assembled dataset contains 153,593 cholinergic synapses (679 neurons), 33881 GABAergic synapses (181 neurons), 7953 glutamatergic synapses (49 neurons), 9526 dopaminergic synapses (89 neurons), 2924 octopaminergic synapses (7 neurons) and 4732 serotonergic synapses (5 neurons).

2.2 Train and Test Datasets

For each neurotransmitter $y \in \{\text{GABA, ACh, GLUT, SER, OCT, DOP}\}$, we divide the data in test, train and validation set by

1. Throughput estimated at around one neuron per minute at sufficient spatial resolution for colocalization with single cell labeling - personal communication with authors.

2. See Footnote 5

3. A commonly used, full step-by-step protocol can be found at <https://www.janelia.org/project-team/flylight/protocols>.

4. <http://www.catmaid.org>

5. <https://neuropil.janelia.org/tracing/fafb>

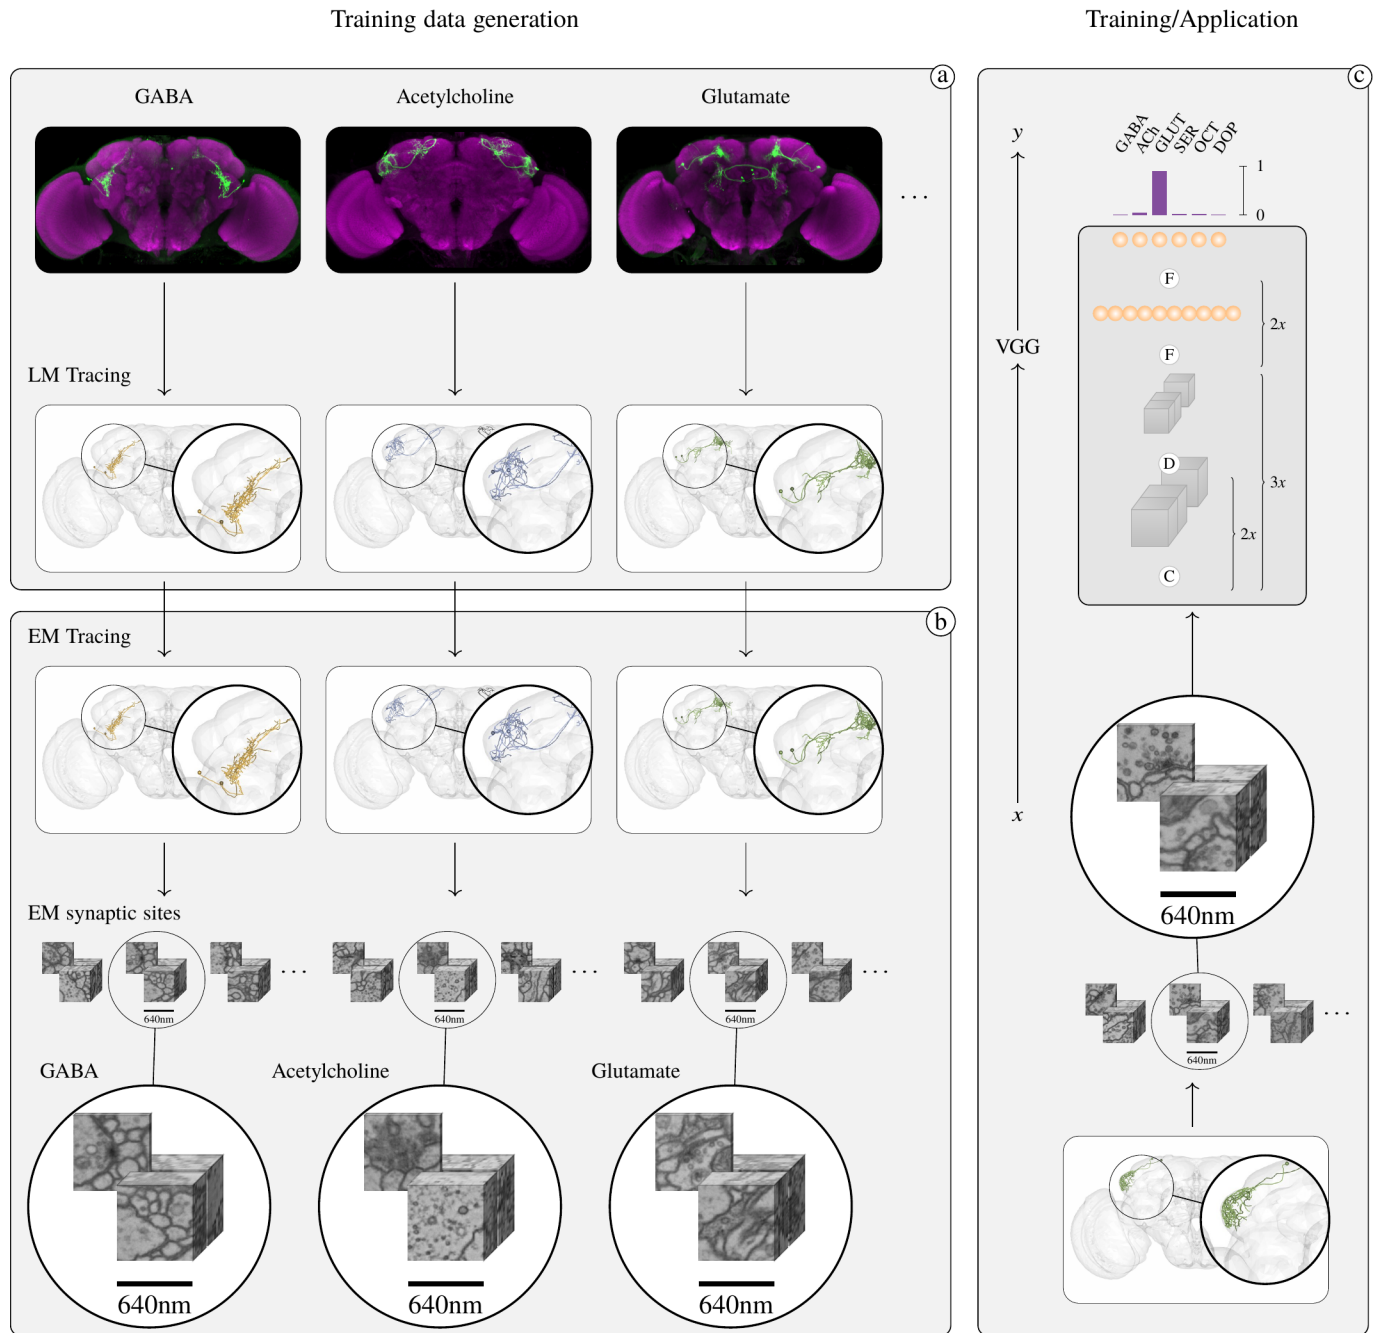


Figure 1: Method Overview. We assemble a dataset of neurons with known neurotransmitter in the *Drosophila* whole brain EM dataset (FAB) (Zheng et al., 2018) from the literature and retrieve corresponding synaptic locations from the subset of skeletons that have been annotated in the CATMAID (Saalfeld et al., 2009; Schneider-Mizell et al., 2016) FAB collaboration database². (a) Typically, neurons are genetically tagged to identify their neurotransmitter identity and to reconstruct their coarse morphology using light microscopy. (b) Light microscopy tracings of neurons are then matched to corresponding neuron traces in the FAB dataset, and synaptic locations are annotated, resulting in a data set of EM volumes of synaptic sites with known neurotransmitter identity. (c) We use the resulting pair (x, y) , with x a 3D EM volume of a synaptic site and $y \in \{\text{GABA, ACh, GLUT, SER, OCT, DOP}\}$, the neurotransmitter of that synaptic site, to train a 3D VGG-style (Simonyan and Zisserman, 2014) deep neural network to assign a given synaptic site x to one of the six considered neurotransmitters. We use the trained network to predict the neurotransmitter identity of synapses from neurons with so far unknown neurotransmitter identity in the *Drosophila* FAB dataset. C, D, and F denote convolution, downsampling, and fully connected layers respectively.

	Neuron split	Hemilineage split	Brain Region split
Train	140,565	140,868	138,982
Test	40,104	40,703	39,715
Validation	20,084	19,182	19,858
Avg. Synapse Accuracy	87%	75%	88%
Avg. Neuron Accuracy	94%	92%	95%

Table 1: Overview of the three data splits used for evaluation of the classifier. Shown are the number of synapses for training, testing and validation as well as average synapse and neuron classification accuracy on the test set for each data split.

randomly assigning entire neurons, each containing multiple synapses. We refer to this split as *Neuron split* in the following. We use 70% of neurons for training, 10% for validation and the remaining 20% for testing. Splitting the dataset by entire neurons, instead of randomly sampling synapses, mirrors the real world use case in which we typically know the neurotransmitter of an entire neuron and are interested in the neurotransmitter of an entirely different neuron.

In order to test how well the proposed methods generalizes across morphologically distinct cells and regions, and to exclude potentially confounding variables, we also generate two additional splits that separate the data by hemilineage (*Hemilineage split* - neurons in a hemilineage are lineally related, see section 4 for further details) and brain region (*Brain Region split*) respectively. To this end, we find the optimal split between entire hemilineages and brain regions, such that the fraction of synapses for every neurotransmitter in the train set is close to 80% of all synapses of that neurotransmitter. We further use randomly selected 12.5% of the training synapses (10% of the entire dataset) for validation.

2.3 Network Architecture and Training

We use a 3D VGG-style (Simonyan and Zisserman, 2014) network to predict the neurotransmitter identity from a 3D EM input cube of edge length 640 nm with a synaptic site at its center. The network consists of four functional blocks, each consisting of two 3D convolution operations, batch normalization, ReLU non-linearities and subsequent max pooling with a downsample factor of ($z=1, y=2, x=2$) for the first three blocks and ($z=2, y=2, x=2$) for the last block and is followed by three fully connected layers with dropout ($p=0.5$) applied after the first two. We train the network to minimize cross entropy loss over the six classes (GABA, ACh, GLUT, SER, OCT and DOP), using the Adam optimizer (Kingma and Ba, 2014). We train for a total of 500,000 iterations in batches of size eight and select the iteration with highest validation accuracy for testing. A full specification of the network architecture and training pipeline, including data augmentations, can be found in the appendix. For an illustration of the used network architecture see Fig. 1c.

3 CLASSIFIER ACCURACY

We tested the classifier on our held out test sets. For the *Neuron split*, the test set consists of a total of 40,104 synapses from 185 neurons that the network was not trained on. We achieve an average, per transmitter accuracy of 87% for the neurotransmitter prediction of individual synapses. Since we assume that each neuron expresses the same neurotransmitters at its synaptic sites we can additionally quantify the per neuron accuracy of the presented method. To this end we assign each neuron with more




Figure 2: Illustration of (a) the progression of a Type I neuroblast from third-instar (L3) larva into the adult, GMC, ganglion mother cell and (b) breakdown of a single secondary lineage, LH12 (also known as DPLal2) into its two hemilineages. Neuronal reconstruction data from the FAFB project shown, which covers half the neurons in this lineage.

than 30 synapses in the test set a neurotransmitter by majority vote of its synapses, leading to an average accuracy of 94% for the neurotransmitter prediction per neuron. For the *Hemilineage split*, we find an accuracy of 75% for individual synapses and 92% for entire neurons. The *Brain Region split* evaluates to 88% synapse classification accuracy and 95% neuron classification accuracy. A per class overview can be seen in Fig. 3, for a summary of the results and data splits see Table 1.

4 TRANSMITTER PREDICTION FOR HEMILINEAGES

Similar to neurons, which release the same set of neurotransmitters at their synaptic sites (Eccles, 1976; Dale, 1934), it has been found that sets of lineally related neurons in the *Drosophila melanogaster* ventral nervous system (VNS), so-called hemilineages, also show homogeneous neurotransmitter expression patterns (Lacin et al., 2019). If a similar principle holds for the *Drosophila melanogaster* brain, it would enable us to assign neurotransmitter identity to large groups of neurons simultaneously. With the presented method we are able to verify to what extent such a principle holds.

The about 45,000 neurons of the central brain of *Drosophila melanogaster* (Croset et al., 2018) are generated by a set of stem cells known as neuroblasts. During division neuroblasts generate two cells, one additional stem cell and one cell that further divides into two sibling neurons. In only one of these siblings the so called Notch pathway is activated, leading to two different “hemilineages” of neurons within the lineage (Kumar et al., 2009; Sen, 2019; Lacin et al., 2019). Lacin et al. (2019) showed that each hemilineage in the VNS expresses just one of the fast-acting transmitters acetylcholine, glutamate and GABA, even though mRNA transcripts for combinations of these can appear in the nucleus (Lacin et al., 2019). This raises the question whether the same holds true in the adult brain. Using the presented classifier, we predict the neurotransmitter identity of all identified neurons within 89 out of a total of ~150 identified hemilineages in the *Drosophila melanogaster* brain. The majority of our predictions show homogeneity of neurotransmitter identity within a single hemilineage, in line with findings in the VNS. However, we identify a set of hemilineages which express two fast acting neurotransmitters with high statistical significance. We find no hemilineage that expresses all three. As a result, our predictions are inconsistent with the hypothesis that all hemilineages express the same fast-acting neurotransmitter.

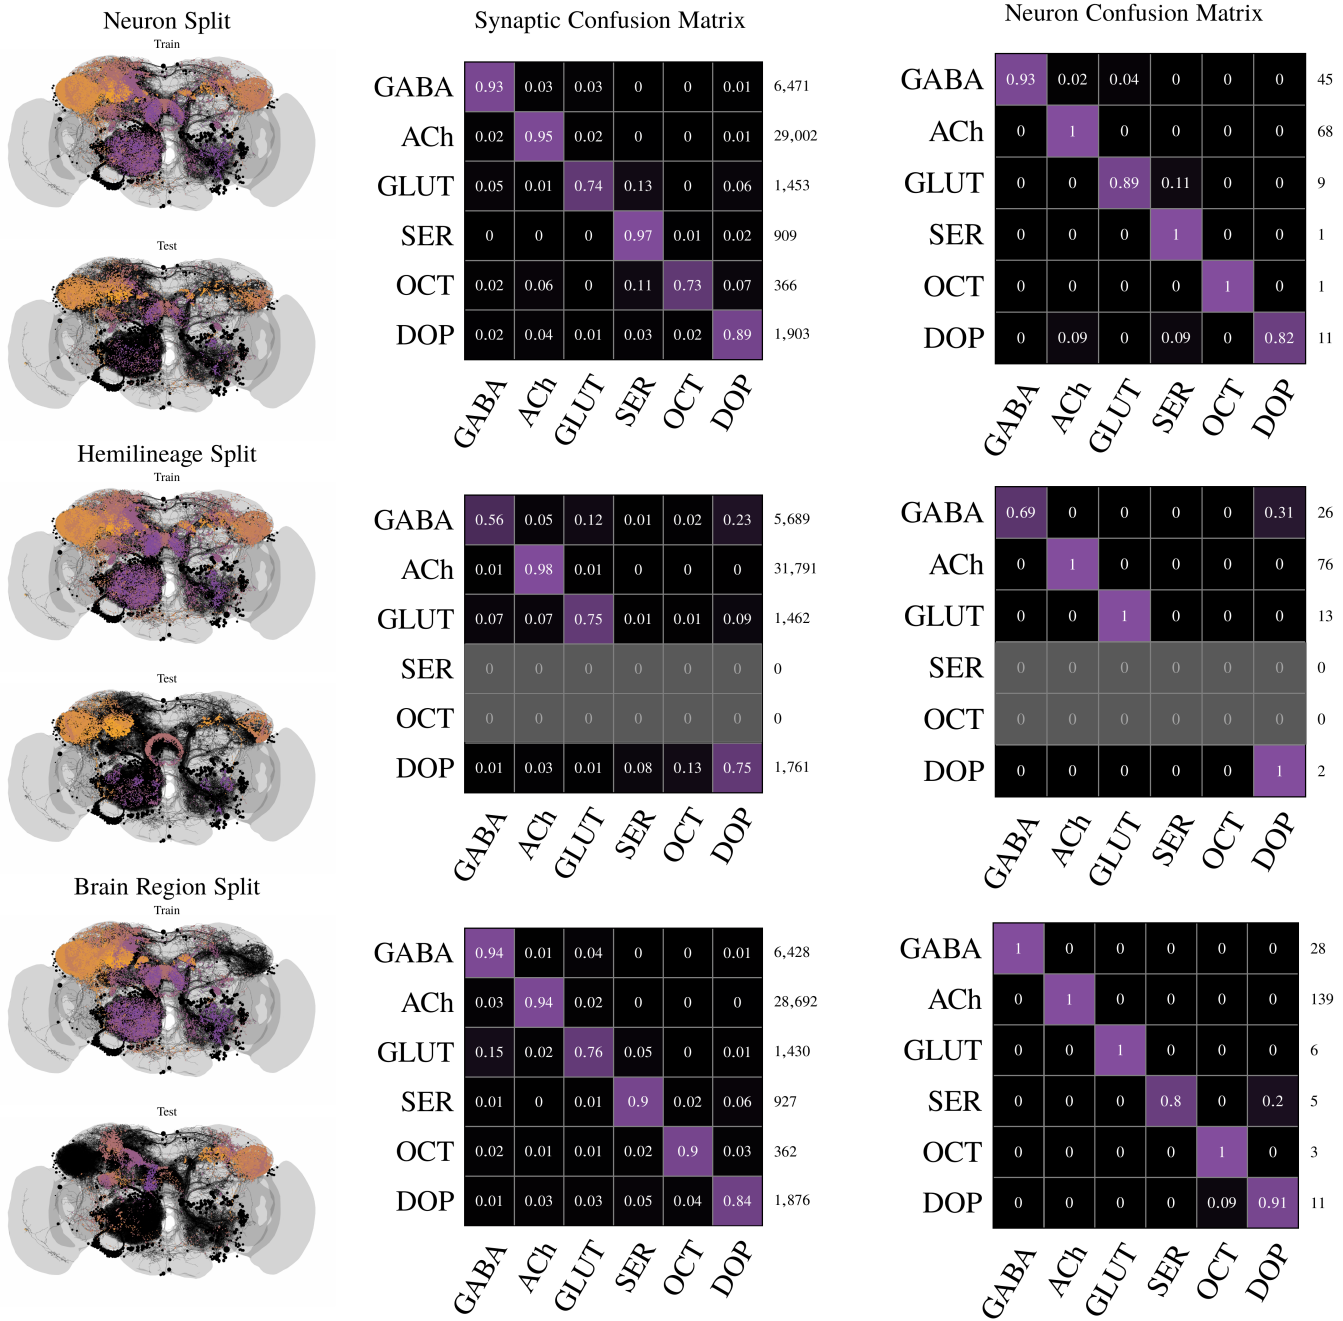


Figure 3: Illustration of the spatial distribution of synapses in the considered data splits (left column) and corresponding confusion matrices for synapses (middle column) and entire neurons (right column). **Datasets:** For each split we visualize the synaptic locations used for training (top) and for testing (bottom). Synapse locations are color coded according to their z-depth (perpendicular to viewing plane). **Confusion Matrices:** Rows show labels and columns the predicted neurotransmitter. The total number of test set ground truth synapses and neurons respectively are shown next to each row. In order to be able to have a meaningful majority vote we only consider neurons with more than 30 synapses for the neuron confusion matrices. Note that for the hemilineage split it was not possible to generate a fully balanced split and as a result there are no SER and OCT neurons in the test set, as indicated by grayed out rows.

4.1 Hemilineage assignments in *Drosophila*

Cell body fiber tracts for identified hemilineages had previously been identified using TrakEM2 (Cardona et al., 2012) in a light-level atlas for a *Drosophila melanogaster* brain, stained with an antibody against neurotactin (BP104) (Lovick et al., 2013). We extracted these expertly identified tracts and registered them into a common template brain, JFRC2, using CMTK (Rohlfing and Maurer, 2003), and then into FAFB space (Bates et al., 2019b). We could now identify cell body fibre tracts in this ssTEM dataset in the vicinity of the transformed hemilineage tracts.

4.2 Predictions

We retrain the classifier on 90% of the entire dataset and use the remaining 10% to select the best performing iteration. We predict the neurotransmitter identity of 180,675 synapses within 1,164 neurons with so far unknown neurotransmitter identity. These neurons come from a total of 89 hemilineages, of which 20 have more than one neuron with genetically identified neurotransmitter. Fig. 4 shows ground truth neurotransmitter annotations for the subset of neurons N_{gt}^h that have known neurotransmitters and our predictions for the remaining neurons N_{pred}^h ($N_{gt}^h \cap N_{pred}^h = \emptyset$) in the hemilineage for five selected hemilineages. In the following, we analyse the results by quantifying how neurotransmitter predictions are distributed over neurons and synapses within hemilineages.

4.2.1 Neuron Level Entropy

In order to quantify multimodality of neurotransmitter predictions on neuron level within a hemilineage we calculate the entropy H of the neurotransmitter distribution over neurons in the following way: Let $n \in N_h$ be a neuron in hemilineage h and $\hat{y}_n \in Y = \{\text{GABA, ACh, GLUT, SER, OCT, DOP}\}$ the predicted neurotransmitter of neuron n . Then

$$H(N_h) = - \sum_{y \in Y} p_h(y) \log_6 p_h(y), \text{ with} \quad (1)$$

$$p_h(y) = \frac{1}{|N_h|} \sum_{n \in N_h} \delta(\hat{y}_n = y) \quad (2)$$

A value of $H(N_h) = 0$ (minimal entropy) then means that all neurons within hemilineage h have the same predicted neurotransmitter, while a value of $H(N_h) = 1$ (maximal entropy) means that within hemilineage h all predicted neurotransmitters are equally common.

4.2.2 Synapse Level Entropy

Similarly we can quantify the average multimodality over synapses within neurons of a given hemilineage: Let $s \in S_n$ be the synapses in neuron $n \in N_h$ of hemilineage h and \hat{y}_s the predicted neurotransmitter. The entropy of predicted synaptic neurotransmitters $H(s_n)$ in neuron n is then given by:

$$H(S_n) = - \sum_{y \in Y} p_n(y) \log_6 p_n(y), \text{ with} \quad (3)$$

$$p_n(y) = \frac{1}{|S_n|} \sum_{s \in S_n} \delta(\hat{y}_s = y) \quad (4)$$

With this, the average synaptic entropy over all neurons within hemilineage h is given by:

$$H(S_h) = \frac{1}{|N_h|} \sum_{n \in N_h} H(S_n) \quad (5)$$

A value of $H(S_h) = 0$ (minimal entropy) then means that all synapses of all neurons in hemilineage h have the same predicted neurotransmitter, while a value of $H(S_h) = 1$ (maximal entropy) means that in all neurons within hemilineage h all synaptic neurotransmitter predictions are equally common. Fig. 5 shows the distribution of $H(N_h)$ and $H(S_h)$ of all predicted hemilineages with more than ten neurons that have more than 30 synapses each.

On the population level we find relatively lower values of $H(S_h)$ (Synapse level entropy) than $H(N_h)$ (Neuron level entropy). 75% of hemilineages show a synapse level entropy below $q_{75}(H(S_h)) = 0.19$ as compared to $q_{75}(H(N_h)) = 0.34$. This is reassuring as it suggests less variation of neurotransmitter identity predictions within individual neurons compared to variations of neurotransmitter identity of neurons within a hemilineage. However, we also find cases with a high level of synaptic entropy, such as hemilineage 16 and 30. For these hemilineages it is unclear whether neuron level multimodality is only an artifact of uncertain, multimodal predictions on synapse level of individual neurons. In contrast to 16 and 30 hemilineages 29, 27 and 42 show high neuron level entropy $H(N_h) \geq q_{75}$ but low synapse level entropy $H(S_h) \leq q_{25}$, suggesting clear neuron level segregation of predicted neurotransmitters within those hemilineages. Hemilineages such as 24 and 40 with $H(S_h) < q_{25}$ and $H(S_n) < q_{25}$ appear homogeneous within each neuron and within the entire hemilineage.

4.3 Number of distinct, fast-acting neurotransmitters in hemilineages of the *Drosophila melanogaster* adult brain

We can now ask the question how likely it is to observe a given prediction of neurotransmitters in a hemilineage under some error rate given by the confusion matrix on the test set, and the assumption that all neurons in the hemilineage have the same underlying neurotransmitter. We can then compare this likelihood to the alternative hypothesis that a hemilineage consists of neurons with more than one neurotransmitter. Out of 26 investigated hemilineages with a sufficient number of predicted neurotransmitters, up to five show strong evidence for expressing two distinct, fast-acting neurotransmitters (Bayes factor $K \geq 10^2$, *decisive*). We find none that expresses all three.

4.3.1 Probability to observe neurotransmitter predictions \hat{y}

Given a neuron has true neurotransmitter $y \in Y$, the probability that we predict neurotransmitter $\hat{y} \in Y$ (assuming that each prediction is independent and identically distributed) is given by the categorical distribution

$$p(\hat{y}|y) = C_{y,\hat{y}} \quad (6)$$

where C is the neuron confusion matrix obtained on the test data set (see Fig. 3).

Let m be the number of different neurotransmitters in hemilineage h . We model the probability $p(\hat{y}|m)$ of observing neurotransmitter predictions $\hat{y} = \{\hat{y}_0, \hat{y}_1, \dots, \hat{y}_n\}$ under the assumption that hemilineage h contains m different neurotransmitters. Here, \hat{y}_j




Figure 4: Neurotransmitter barcode plots and corresponding renderings of skeletons and synapses (color coded according to their neurotransmitter identity) of five selected hemilineages, for which a subset of neurons N_{gt}^h have genetically determined, known neurotransmitters (left) and our predictions on the remaining neurons N_{pred}^h in the same hemilineage (right). Each column in the neurotransmitter barcode represents one neuron. For each neuron (column), the relative number of synapses with neurotransmitter $y, \hat{y} \in Y = \{\text{GABA, ACh, GLUT, SER, OCT, DOP}\}$ is represented by the color intensity of the respective row. The total number of synapses in each neuron is shown above each row. Note that $N_{\text{gt}}^h \cap N_{\text{pred}}^h = \emptyset$. Neuron classes, shown above the inset, are given for the most numerous cells in our training data, for each hemilineage that we show (Dolan et al., 2019; Otto et al., 2020; Aso et al., 2014). For an overview of all hemilineages that have partially known neurotransmitter identities and our associated predictions, see Fig. 10.

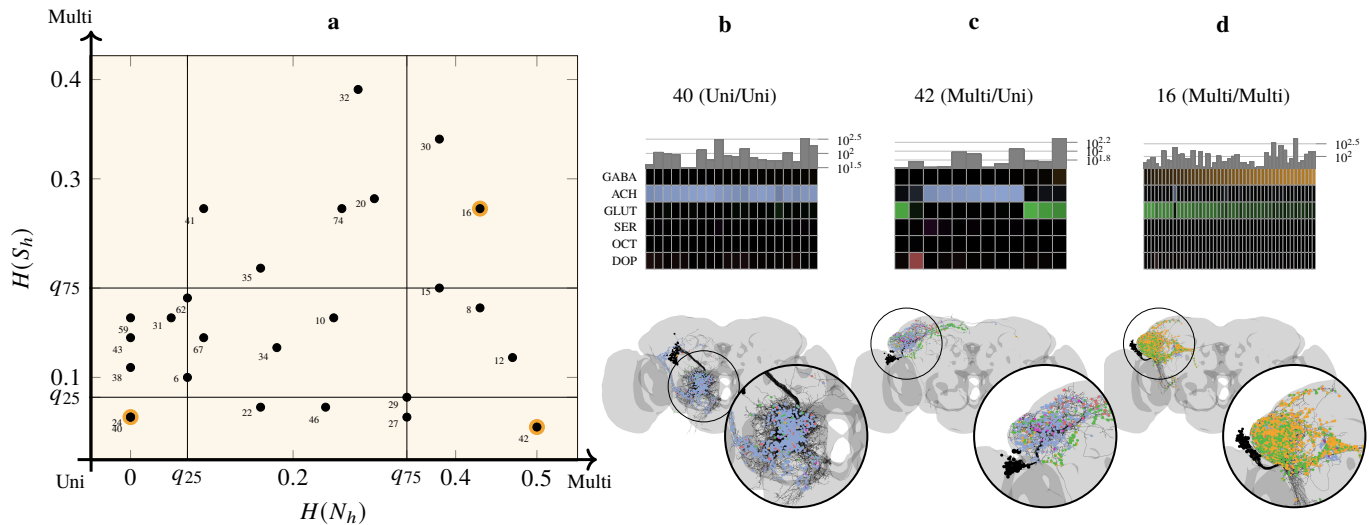


Figure 5: **(a)** Neuron level entropy vs average synapse level entropy for all predicted hemilineages with more than 10 neurons and more than 30 synapses per neuron. $q_{25}(H(N_h)) = 0.07$, $q_{25}(H(S_h)) = 0.08$ and $q_{75}(H(N_h)) = 0.34$, $q_{75}(H(S_h)) = 0.19$ are indicating 25% and 75% percentiles respectively. Highlighted hemilineages show the extreme points of the entropy distribution: **(b)** Hemilineage 40 shows low neuron level entropy $H(N_h)$ and low synapse level entropy $H(S_h)$, caused by a unimodal distribution of neurotransmitters on synapse and neuron level in this hemilineage. **(c)** Hemilineage 42 shows high neuron level entropy $H(N_h)$ but low synapse level entropy $H(S_h)$, caused by neurotransmitter predictions that are unimodal within each neuron but multimodal across neurons. **(d)** Hemilineage 16 shows high neuron level entropy $H(N_h)$ and high synapse level entropy $H(S_h)$, as a result of a bimodal distribution of neurotransmitter predictions within most neurons of this hemilineage. For a mapping of hemilineage ID to hemilineage name see Table 2.

is the predicted neurotransmitter of neuron j in hemilineage h with n neurons total. Let $\mathbb{P}_c(Y)$ be the set of subsets of true neurotransmitters Y with cardinality c , then:

$$p(\hat{y}|m) = \sum_{S \in \mathbb{P}_m(Y)} p(\hat{y}|S) \cdot p(S|m), \quad (7)$$

where $p(\hat{y}|S)$ is the probability to observe predictions \hat{y} if the hemilineage has true underlying neurotransmitters $y \in S$ and $p(S|m)$ is the probability for the set of true neurotransmitters S given the hemilineage contains m different neurotransmitters. Since we assume i.i.d. predictions \hat{y} , $p(\hat{y}|S)$ factorizes as follows:

$$p(\hat{y}|S) = \prod_j p(\hat{y}_j|S) \quad (8)$$

and marginalizing over $y \in S$ yields:

$$p(\hat{y}|S) = \prod_j \sum_{y \in S} p(\hat{y}_j|y) \cdot p(y|S) \quad (9)$$

$$= \prod_j \sum_{y \in S} C_{y, \hat{y}_j} \cdot p(y|S) \quad (10)$$

Regarding $p(S|m)$ and $p(y|S)$ we assume a flat prior, *i.e.*:

$$p(S|m) = \binom{|Y|}{m}^{-1} \quad (11)$$

$$p(y|S) = \frac{1}{|S|} = \frac{1}{m} \quad (12)$$

With this, the probability of observing predictions \hat{y} given m different neurotransmitters becomes:

$$p(\hat{y}|m) = \binom{|Y|}{m}^{-1} \sum_{S \in \mathbb{P}_m(Y)} \left(\prod_j \sum_{y \in S} C_{y, \hat{y}_j} \cdot \frac{1}{|S|} \right) \quad (13)$$

4.3.2 Bayes Factor

With this formalism in place, we can compare hypotheses about the number of true neurotransmitters m in a given hemilineage by using the Bayes Factor $K = \frac{p(D|M_1)}{p(D|M_2)}$, where D is our observed data (predicted neurotransmitters) and M_1 , M_2 are two models about the underlying true neurotransmitters that we wish to compare. The Bayes factor for a model M_1 with m_1 true neurotransmitters per hemilineage and model M_2 with m_2 different neurotransmitters is given by:

$$K = \frac{p(\hat{y}|m_1)}{p(\hat{y}|m_2)} \quad (14)$$

$$= \frac{\binom{|Y|}{m_1}^{-1} \sum_{S \in \mathbb{P}_{m_1}(Y)} \left(\prod_j \sum_{y \in S} C_{y, \hat{y}_j} \cdot \frac{1}{m_1} \right)}{\binom{|Y|}{m_2}^{-1} \sum_{S \in \mathbb{P}_{m_2}(Y)} \left(\prod_j \sum_{y \in S} C_{y, \hat{y}_j} \cdot \frac{1}{m_2} \right)} \quad (15)$$

So far, we assumed that $p(\hat{y}_j|y) = C_{y, \hat{y}_j}$, *i.e.*, we estimate this distribution on the test dataset. However, because our test set is finite we can not expect that the estimated error rates perfectly transfer to other datasets. In order to relax our assumptions about this distribution we simulate additional errors, by incorporating additive smoothing on the counts of neurons $N_{y, \hat{y}}$ that have true neurotransmitter y and were predicted as neurotransmitter \hat{y} , *i.e.*:

$$\tilde{C}_{y, \hat{y}} = \frac{N_{y, \hat{y}} + \beta}{\sum_{\hat{y}} N_{y, \hat{y}} + 6\beta}, \quad (16)$$

where $\beta \in \mathbb{N}_0$ is the smoothing parameter. With $C_{y, \hat{y}} = \frac{N_{y, \hat{y}}}{\sum_{\hat{y}} N_{y, \hat{y}}}$ we then have

$$\tilde{C}_{y, \hat{y}} = \frac{C_{y, \hat{y}} + \frac{\beta}{\sum_{\hat{y}} N_{y, \hat{y}}}}{1 + 6 \frac{\beta}{\sum_{\hat{y}} N_{y, \hat{y}}}} = \frac{C_{y, \hat{y}} + \alpha}{1 + 6\alpha} \quad (17)$$

and $\alpha \in \mathbb{R}_{\geq 0}$ the count normalized smoothing parameter. In the limit of $\alpha \rightarrow \infty$, $\tilde{C}_{y,\hat{y}}$ approaches the uniform distribution with probability $1/6$ for each neurotransmitter, whereas a value of $\alpha = 0$ means we recover the observed confusion matrix C . With this, our distributions are now parametrized by α and the bayes factor becomes:

$$K = \frac{\int_{\alpha} p(\hat{y}, \alpha | m_1) p(\alpha) d\alpha}{\int_{\alpha} p(\hat{y}, \alpha | m_2) p(\alpha) d\alpha} \quad (18)$$

$$= \frac{\tilde{p}(\hat{y} | m_1)}{\tilde{p}(\hat{y} | m_2)} \quad (19)$$

$$(20)$$

where $\tilde{p}(\hat{y} | m)$ is as defined in (13) but with C_{y,\hat{y}_j} replaced with its expected value $\mathbb{E}_{p(\alpha)}[\tilde{C}_{y,\hat{y}_j}]$.

The prior distribution on α , $p(\alpha)$, allows us to encode our prior knowledge about α and use it to weight the likelihood of the corresponding model. Given the data, a value of $\alpha = \epsilon$ with epsilon small ($0 < \epsilon \ll 1$), should be most probable, while the probability of values $\alpha > \epsilon$ should monotonically decrease as we deviate more from the observed confusion matrix. Values of $\alpha < \epsilon$ should have probability zero, because they correspond to the unsmoothed confusion matrix with zero entries, *i.e.*, a probability of zero for missclassification of certain neurotransmitters. While these probabilities may be small, they are likely greater than zero and an artifact caused by the finite test set. Many distributions fulfill these criteria, in particular the family of exponential distributions with rate parameter λ :

$$p(\alpha) = \begin{cases} \lambda e^{-\lambda(\alpha-\epsilon)} & \alpha \geq \epsilon \\ 0 & \alpha < \epsilon \end{cases}$$

Thus, λ controls the weight for smoothing parameter α in the integral $\mathbb{E}_{p(\alpha|M)}[\tilde{C}(\alpha)_{y,\hat{y}_j}] = \int_{\alpha} \tilde{C}_{y,\hat{y}_j} p(\alpha) d\alpha$. For $\lambda \rightarrow 0$, the expected confusion matrix approaches the unweighted average of all $C(\alpha)$ in the integration range. For $\lambda \rightarrow \infty$, the expected confusion matrix approaches the ϵ -smoothed confusion matrix $\tilde{C}_{y,\hat{y}} = \frac{C_{y,\hat{y}} + \epsilon}{1+6\epsilon}$. The rate parameter λ can also be understood via its influence on the expected average accuracy $\tilde{c}_{\text{exp}} = \frac{1}{6} \sum_i \mathbb{E}_{p(\alpha|M)}[\tilde{C}]_{i,i}$. For values of $\lambda \rightarrow 0$, the expected accuracy approaches chance level while for values of $\lambda \rightarrow \infty$, the expected accuracy approaches the ϵ -smoothed, observed accuracy on the test set.

4.3.3 Evidence for two fast acting transmitters in a single hemilineage

We calculate the bayes factor $K_{2,1} = \frac{p(\hat{y}|m=2)}{p(\hat{y}|m=1)}$ and $K_{3,2} = \frac{p(\hat{y}|m=3)}{p(\hat{y}|m=2)}$ for the set of three classical neurotransmitters $Y_{\text{cl}} = \{\text{GABA}, \text{ACh}, \text{GLUT}\}$ for those hemilineages that have more than ten annotated neurons and 30 annotated synapses each with neurotransmitter predictions in the set Y_{cl} . For this analysis, we ignore all other neurons with predicted neurotransmitter identity $\hat{y} \notin Y_{\text{cl}}$. Fig. 6 shows $K_{2,1}$ and $K_{3,2}$ for a range of rate parameters λ and corresponding expected average accuracy \tilde{c}_{exp} . For hemilineages 30, 16, 27, 20, and 42 there is decisive evidence ($K_{2,1} \geq 10^2$) for the presence of two distinct fast acting neurotransmitters for a large range of expected accuracies \tilde{c}_{exp} . However, note that hemilineage 30 and 16 show high synaptic entropy $H(S_h)$ (see Fig. 5), indicating that individual neurons within the hemilineage already show bimodal neurotransmitter predictions. As such,

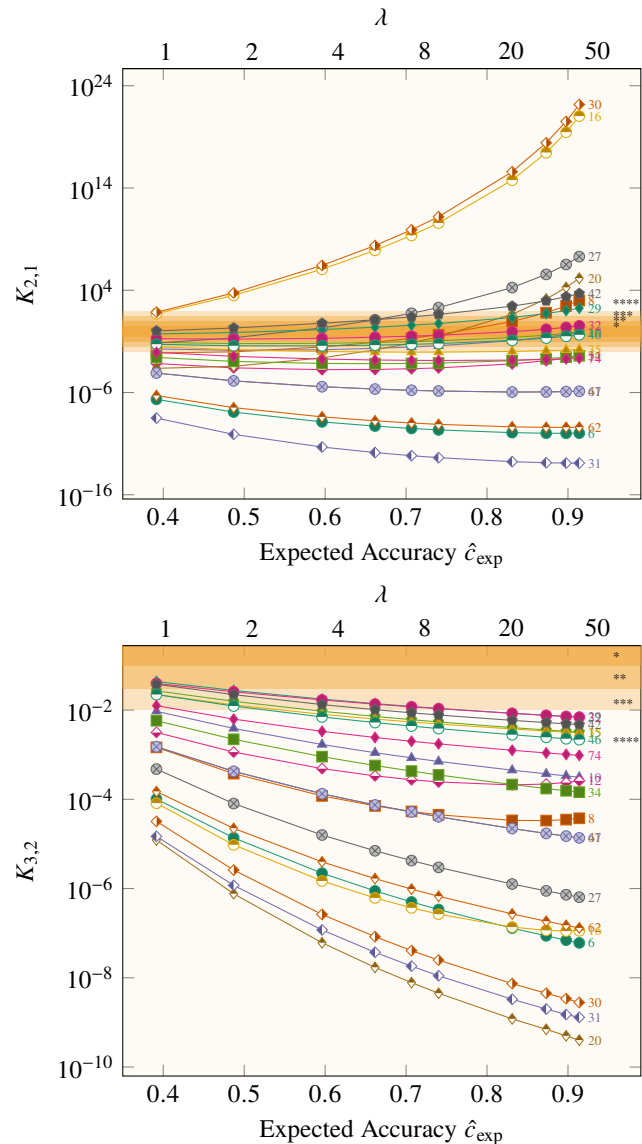


Figure 6: Bayes factor K over a range of expected average predictor accuracy \tilde{c}_{exp} . Shown are hemilineages with more than ten neurons that have more than 30 synapses each and more than one predicted, fast-acting neurotransmitter. Stars indicate regions of evidence for model M_1 ($K > 1$) or model M_2 ($K < 1$), respectively: * - substantial, ** - strong, *** - very strong, **** - decisive (Jeffreys, 1998). **Top:** $M_1: m_1 = 2$ and $M_2: m_2 = 1$. **Bottom:** $M_1: m_1 = 3$ and $M_2: m_2 = 2$. For a mapping of hemilineage ID to hemilineage name see Table 2.

strong bimodality on neuron level is at least partially explained by uncertain predictions on individual synaptic level. This is in contrast to hemilineage 27 and 42, who show synaptic entropies below the 25% percentile. In these hemilineages, large Bayes factor values $K_{2,1}$ directly stem from neuron level segregation of the predicted neurotransmitters within the hemilineage. See Fig. 7 for a rendering of the neurotransmitter predictions of these hemilineages and corresponding NBLAST dendrograms, indicating that the two fast acting neurotransmitters in some of these hemilineages are divided between morphologically distinct neurons. The remaining 13 hemilineages show no strong evidence

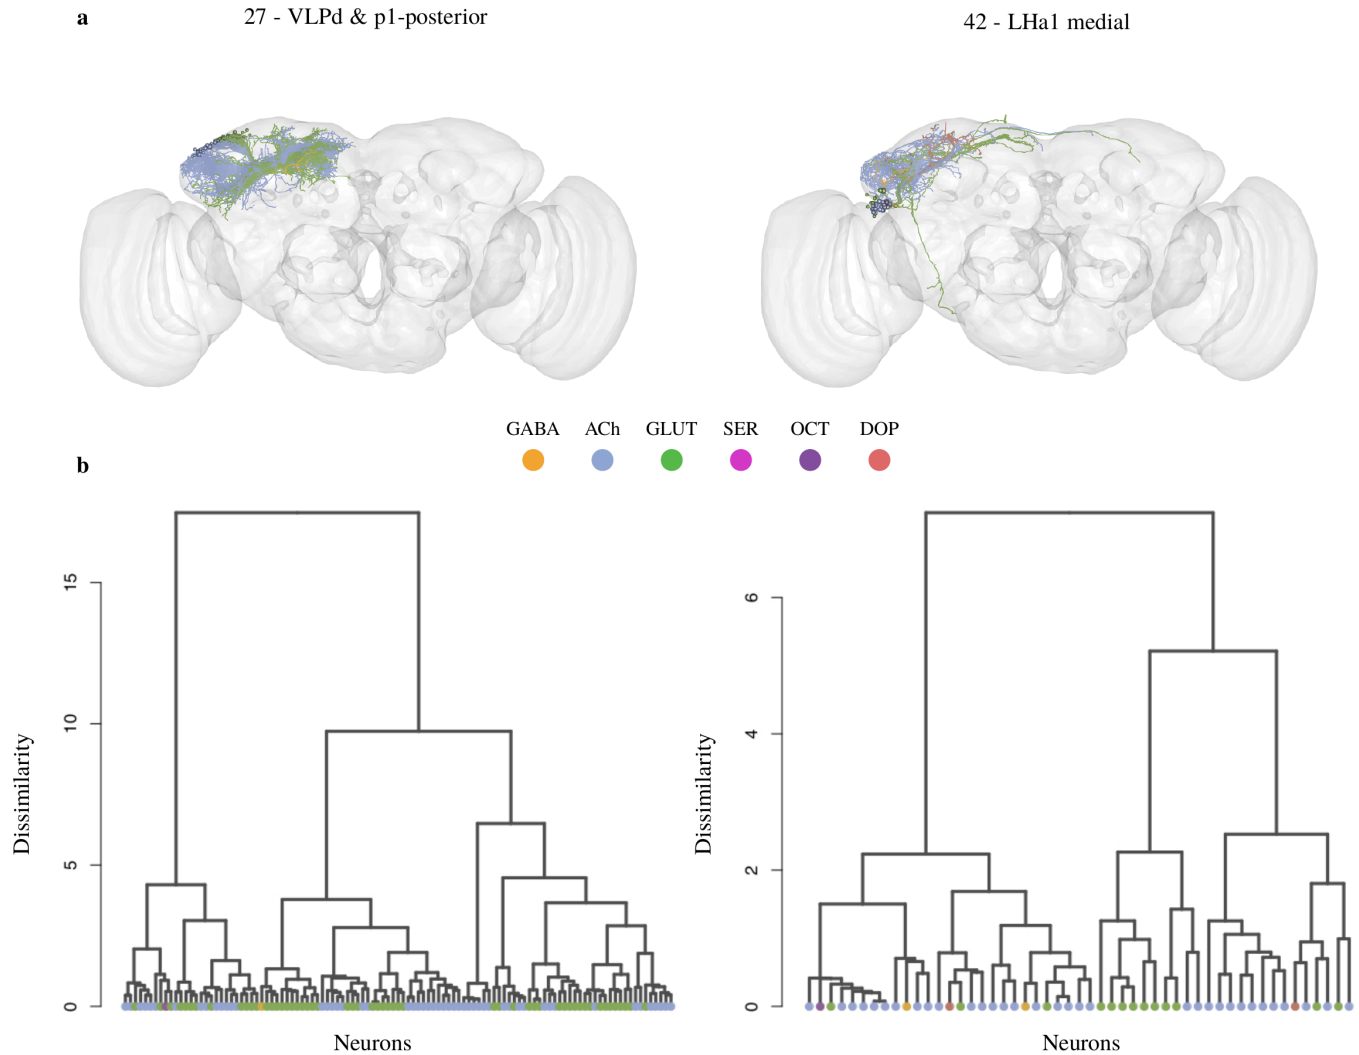


Figure 7: Renderings of neurotransmitter predictions of all neurons within two selected hemilineages that show a high bayes factor $K_{2,1}$ (see Fig. 6) in addition to low synaptic entropy $H(S_h)$ (see Fig. 5) (a), and corresponding NBLAST dendrograms (b). Y-axis shows the morphological dissimilarity between clusters, based on Ward's method. Each neuron is color coded according to its predicted, majority vote neurotransmitter. The dendrograms show that neurotransmitter predictions correlate strongly with neuron morphology for hemilineage 27 and to a lesser degree for hemilineage 42. For renderings of all hemilineages and corresponding barcode plots see Fig. 9 and Fig. 8.

for either hypothesis ($K_{2,1} \approx 1$, $n=5$) or favor the hypothesis of expressing only one fast acting neurotransmitter ($K_{2,3} \leq 10^{-2}$, $n=8$). No hemilineage shows evidence for expressing all three fast acting neurotransmitters ($K_{2,3} < 10^{-2}$).

5 DISCUSSION

5.1 Results

We presented a classifier which is able to predict the neurotransmitter identity of a synapse from a local 3D EM volume with high accuracy. We showed that the method generalizes across neurons, brain regions and hemilineages. Furthermore we predicted the neurotransmitter identity of 180,675 synapses within 1,164 neurons from 89 hemilineages with so far unknown neurotransmitter identity. We analyzed the neurotransmitter distribution of 26 hemilineages that have a sufficient amount of annotated neurons and synapses and showed that most of them homogeneously express one fast acting neurotransmitter. However, we also identified a

set of five hemilineages that, according to our predictions, express two distinct fast-acting neurotransmitters with high statistical significance. Two of those five, 27 and 42, also show low synaptic entropy $H(S_h)$, indicating that the observed effect is a result of neuron-level neurotransmitter segregation within the hemilineage.

5.2 Limitations

A potential source of neurotransmitter missclassification is the possibility that a given neuron releases more than one neurotransmitter at its synaptic sites. Due to a lack of known and annotated neurons with cotransmission of the considered neurotransmitters our current model ignores this possibility. However, single cell transcriptomic data of the *Drosophila melanogaster* brain shows that neurotransmitter gene expression is largely exclusive for the fast acting transmitters ACh, GABA and GLUT (Croset et al., 2018), excluding widespread cotransmission of these transmitters. For the considered monoamines (SER, OCT,

and DOP), coexpression with another fast acting transmitter is more probable. In particular Croset et al. (2018) suggests that a large fraction of octopaminergic neurons likely corelease GLUT, while SER and DOP show less evidence for coexpression with fast acting transmitters. If a particular neuron in the dataset were to corelease a fast acting transmitter and a monoamine the presented classifier would predict only one of the two. However, this is not a fundamental limitation of the presented approach and could be remedied if coexpression training data becomes available.

Another current limitation is the fact that we only consider the set of six neurotransmitters {GABA, ACh, GLUT, DOP, OCT, SER} and due to our use of a softmax normalization at the network output layer, the model is forced to select one of the six classes, even if there is no evidence for either of them. As a result, the current model is not able to identify synapses or neurons that are not part of the considered neurotransmitters, notably histamine (Nässel, 1999), tyramine and a vast number of neuropeptides (Croset et al., 2018). Similar to coexpression, we expect an extension to further neurotransmitters to be possible with available training data.

Regarding our analysis of the number of distinct fast-acting neurotransmitters in a hemilineage, a potential source of error is misassignments of neurons to hemilineages. If neurons are erroneously assigned to a particular hemilineage any observed effect of multimodal neurotransmitter distributions on neuron-level within a hemilineage could be an artifact. Furthermore, for hemilineages 16 and 30 the high synaptic entropy $H(S_h)$ suggests that the phenotype of synapses is ambiguous in these hemilineages. Although coexpression of fast-acting neurotransmitters is unlikely (Croset et al., 2018), the neurotransmitter distribution would be consistent with cotransmission of GABA and glutamate in the neurons of these hemilineages.

5.3 Generalization

We showed that our network is able to generalize across brain regions and hemilineages. However, although the performance on the brain region split (88% average accuracy) outperforms even the baseline neuron split (87% average accuracy), the hemilineage split suffers a performance decrease of more than 10% (75% average accuracy), suggesting that the influence of the neurotransmitter on the phenotype of a synaptic site is influenced by its hemilineage. This is partially remedied when averaging over multiple synapses: Neuron-level neurotransmitter classification for the hemilineage split is robust with an average accuracy of 92%. Note that the presented data splits already exclude neuron identity, hemilineage identity and brain region as potential confounding variables for the prediction of neurotransmitter identity as performance is far above chance level.

5.4 Interpretability

Given that in invertebrates humans can not generally distinguish different neurotransmitter containing vesicles from EM alone, an interesting question is which features the classifier relies on. However, despite the fact that we have complete information about the artificial neural system that is able to assign images of synaptic sites to neurotransmitters, it is difficult to extract the rules under which it operates. It is, for example, possible to visualise areas of

importance for the neural network via use of various attribution methods (for an overview see Ancona et al. (2018)), but often the masked input image still contains a large amount of complexity that varies for each image of a given class (data not shown). We aim to explore this question in more depth in future work.

6 ACKNOWLEDGEMENTS

This work was supported by a Wellcome Trust Collaborative Award (203261/Z/16/Z), an ERC Consolidator grant (649111) and core support from the MRC (MC-U105188491) to G.S.X.E.J; a Boehringer Ingelheim Fonds PhD Fellowship and a Herchel Smith Studentship to A.S.B. The synapses in this study were manually identified in the FAFB dataset by tracers working on different projects across 17 laboratories. For these annotations, we thank Ruairi Roberts, Philipp Schlegel, Fiona Love, Lisa Marin, Amelia Edmondson-Stait, Xincheng Zhao, Jawaid Ali, Johann Schor, Imaan Tamimi, Arian Jamasb, Marisa Dreher, Markus Pleijzier, Robert Turnbull, Nadiya Sharifi, Steven Calle, Andrew Dacks, Konrad Heinz, Kimberly Meechan, Aidan Smith, Najla Masoodpanah, Serene Dhawan, Peter Gibb, Corey Fisher, Claire Peterson, Jason Polsky, Tansy Yang, Kathi Eichler, Joseph Hsu, Irene Varela, Lucia Kmecova, Istvan Taisz, Jacob Ratliff, Kaylynn Coates, Anna Li, Marta Costa, Tyler Paterson, Claire Managan, Adam Heath, Katie Stevens, Jack Mccarty, Nora Forknall, Laurin Bueld, Neha Rampally, Zane Mitrevica, Kelli Fairbanks, Stanley Tran, Shada Alghailani, Quinn Vanderbeck, Lauren Warner, Henrique Ludwig, Jeremy Johnson and Levi Helmick, each of whom have contributed over 1,000 synapses. We principally thank the Wellcome Trust UK and US Drosophila Connectomics, Jefferis, Janelia CATMAID annotation team, Bock, Preat, Wilson, Dacks, Hampel and Seeds groups for sharing their published and unpublished work in the FAFB dataset. Further, we thank Michael Reiser, Vivek Jayaraman, Arthur Zhao, Tatsuo Okubo, Jenny Lu and Kathi Eichler for identifying neuron matches which helped us build our ground truth.

REFERENCES

- Ancona, M., Ceolini, E., Öztireli, C., and Gross, M. (2018). Towards better understanding of gradient-based attribution methods for deep neural networks. In *International Conference on Learning Representations*.
- Aso, Y., Hattori, D., Yu, Y., Johnston, R. M., Nirmala, A., Ngo, T.-T., Dionne, H., Abbott, L. F., Axel, R., Tanimoto, H., and Rubin, G. M. (2014). The neuronal architecture of the mushroom body provides a logic for associative learning. *Elife*, pages 1–47.
- Bates, A. S., Janssens, J., Jefferis, G. S., and Aerts, S. (2019a). Neuronal cell types in the fly: single-cell anatomy meets single-cell genomics. *Curr. Opin. Neurobiol.*, 56:125–134.
- Bates, A. S., Manton, J. D., Jagannathan, S. R., Costa, M., Schlegel, P., Rohlfing, T., and Gregory S X (2019b). The natverse: a versatile computational toolbox to combine and analyse neuroanatomical data.
- Bates, A. S., Schlegel, P., Roberts, R. J. V., Drummond, N., Tamimi, I. F. M., Turnbull, R., Zhao, X., Marin, E. C., Popovici, P. D., Dhawan, S., Jamasb, A., Javier, A., Li, F., Rubin, G. M., Waddell, S., Bock, D. D., Costa, M., and G S X (2020). Complete connectomic reconstruction of olfactory projection neurons in the fly brain.
- Bräcker, L. B., Siju, K. P., Varela, N., Aso, Y., Zhang, M., Hein, I., Vasconcelos, M. L., and Grunwald Kadow, I. C. (2013). Essential role of the mushroom body in context-dependent CO2 avoidance in *Drosophila*. *Curr. Biol.*, 23(13):1228–1234.

- Buhmann, J., Krause, R., Ceballos Lentini, R., Eckstein, N., Cook, M., Turaga, S., and Funke, J. (2018). Synaptic partner prediction from point annotations in insect brains. *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*, 11071.
- Buhmann, J., Sheridan, A., Gerhard, S., Krause, R., Nguyen, T., Heinrich, L., Schlegel, P., Lee, W.-C. A., Wilson, R., Saalfeld, S., Jefferis, G., Bock, D., Turaga, S., Cook, M., and Funke, J. (2019). Automatic detection of synaptic partners in a whole-brain *Drosophila* em dataset. *bioRxiv*.
- Busch, S., Selcho, M., Ito, K., and Tanimoto, H. (2009). A map of octopaminergic neurons in the *Drosophila* brain. *J. Comp. Neurol.*, 513(6):643–667.
- Cardona, A., Saalfeld, S., Schindelin, J., Arganda-Carreras, I., Preibisch, S., Longair, M., Tomancak, P., Hartenstein, V., and Douglas, R. J. (2012). TrakEM2 software for neural circuit reconstruction. *PLoS One*, 7(6):e38011.
- Cook, S. J., Jarrell, T. A., Brittin, C. A., Wang, Y., Bloniarz, A. E., Yakovlev, M. A., Nguyen, K. C., Tang, L. T.-H., Bayer, E. A., Duerr, J. S., et al. (2019). Whole-animal connectomes of both *Caenorhabditis elegans* sexes. *Nature*, 571(7763):63–71.
- Costa, M., Manton, J. D., Ostrovsky, A. D., Prohaska, S., and Jefferis, G. S. X. E. (2016). NBLAST: Rapid, sensitive comparison of neuronal structure and construction of neuron family databases. *Neuron*, 91(2):293–311.
- Croset, V., Treiber, C., and Waddell, S. (2018). Cellular diversity in the *Drosophila* midbrain revealed by single-cell transcriptomics. *Elife*, 7:e34550.
- Dale, H. (1934). Pharmacology and nerve endings. *Proc. R. Soc. Med.*, 28:319–332.
- Davie, K., Janssens, J., Koldere, D., De Waegeneer, M., Pech, U., Kreft, L., Aibar, S., Makhzami, S., Christiaens, V., González-Blas, C. B., et al. (2018). A single-cell transcriptome atlas of the aging *Drosophila* brain. *Cell*, 174(4):982–998.
- Davis, F. P., Nern, A., Picard, S., Reiser, M. B., Rubin, G. M., Eddy, S. R., and Henry, G. L. (2018). A genetic, genomic, and computational resource for exploring neural circuit function.
- Davis, F. P., Nern, A., Picard, S., Reiser, M. B., Rubin, G. M., Eddy, S. R., and Henry, G. L. (2020). A genetic, genomic, and computational resource for exploring neural circuit function. *Elife*, 9.
- Dolan, M.-J., Belliard-Guérin, G., Bates, A. S., Frechter, S., Lampin-Saint-Amaux, A., Aso, Y., Roberts, R. J. V., Schlegel, P., Wong, A., Hammad, A., Bock, D., Rubin, G. M., Preat, T., Plaçais, P.-Y., and Jefferis, G. S. X. E. (2018). Communication from learned to innate olfactory processing centers is required for memory retrieval in *Drosophila*. *Neuron*.
- Dolan, M.-J., Frechter, S., Bates, A. S., Dan, C., Huoviala, P., Roberts, R. J., Schlegel, P., Dhawan, S., Tabano, R., Dionne, H., et al. (2019). Neurogenetic dissection of the *Drosophila* lateral horn reveals major outputs, diverse behavioural functions, and interactions with the mushroom body. *eLife*, 8:e43079.
- Eccles, J. C. (1976). From electrical to chemical transmission in the central nervous system: the closing address of the sir henry dale centennial symposium cambridge, 19 september 1975. *Notes and records of the Royal Society of London*, 30(2):219–230.
- Felsenberg, J., Jacob, P. F., Walker, T., Barnstedt, O., Edmondson-Stait, A. J., Pleijzier, M. W., Otto, N., Schlegel, P., Sharifi, N., Perisse, E., Smith, C. S., Lauritzen, J. S., Costa, M., Jefferis, G. S. X. E., Bock, D. D., and Waddell, S. (2018). Integration of parallel opposing memories underlies memory extinction. *Cell*, 175(3):709–722.e15.
- Frechter, S., Bates, A. S., Tootoonian, S., Dolan, M.-J., Manton, J. D., James, A. R., Kohl, J., Bock, D., and Jefferis, G. S. (2019). Functional and anatomical specificity in a higher olfactory centre. *Elife*, 8.
- Funke, J., Tschopp, F. D., Grisaitis, W., Sheridan, A., Singh, C., Saalfeld, S., and Turaga, S. C. (2018). Large scale image segmentation with structured loss based deep learning for connectome reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1.
- Goyal, R. K. and Chaudhury, A. (2013). Structure activity relationship of synaptic and junctional neurotransmission. *Autonomic Neuroscience*, 176(1-2):11–31.
- Henry, G. L., Davis, F. P., Picard, S., and Eddy, S. R. (2012). Cell type-specific genomics of *Drosophila* neurons. *Nucleic acids research*, 40(19):9691–9704.
- Huoviala, P., Dolan, M.-J., Love, F. M., Frechter, S., Roberts, R. J. V., Mitrevica, Z., Schlegel, P., Bates, A. S., Aso, Y., Rodrigues, T., Cornwall, H., Stensmyr, M., Bock, D., Rubin, G. M., Costa, M., and Gregory S X (2018). Neural circuit basis of aversive odour processing in *Drosophila* from sensory input to descending output.
- Hyatt, A. D. and Wise, T. G. (2001). Immunolabeling. In *Immunocytochemistry and In Situ Hybridization in the Biomedical Sciences*, pages 73–107. Springer.
- Ito, M., Masuda, N., Shinomiya, K., Endo, K., and Ito, K. (2013). Systematic analysis of neural projections reveals clonal composition of the *Drosophila* brain. *Curr. Biol.*, 23(8):644–655.
- Januszewski, M., Kornfeld, J., Li, P. H., Pope, A., Blakely, T., Lindsey, L., Maitin-Shepard, J., Tyka, M., Denk, W., and Jain, V. (2018). High-precision automated reconstruction of neurons with flood-filling networks. *Nature methods*, page 1.
- Jeffreys, H. (1998). *The theory of probability*. OUP Oxford.
- Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Konstantinides, N., Rossi, A. M., and Desplan, C. (2015). Common temporal identity factors regulate neuronal diversity in fly ventral nerve cord and mouse retina. *Neuron*, 85(3):447–449.
- Kreshuk, A., Funke, J., Cardona, A., and Hamprecht, F. A. (2015). Who is talking to whom: Synaptic partner detection in anisotropic volumes of insect brain. In Navab, N., Hornegger, J., Wells, W. M., and Frangi, A., editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 661–668, Cham. Springer International Publishing.
- Kumar, A., Bello, B., and Reichert, H. (2009). Lineage-specific cell death in postembryonic brain development of *Drosophila*. *Development*, 136(20):3433–3442.
- Lacin, H., Chen, H.-M., Long, X., Singer, R. H., Lee, T., and Truman, J. W. (2019). Neurotransmitter identity is acquired in a lineage-restricted manner in the *Drosophila* CNS. *Elife*, 8.
- Lai, S.-L., Awasaki, T., Ito, K., and Lee, T. (2008). Clonal analysis of *Drosophila* antennal lobe neurons: diverse neuronal architectures in the lateral neuroblast lineage. *Development*, 135(17):2883–2893.
- Lee, K., Lu, R., Luther, K., and Seung, H. S. (2019). Learning Dense Voxel Embeddings for 3D Neuron Reconstruction. *arXiv e-prints*, page arXiv:1909.09872.
- Long, X., Colonell, J., Wong, A. M., Singer, R. H., and Lionnet, T. (2017). Quantitative mRNA imaging throughout the entire *Drosophila* brain. *nature methods*, 14(7):703.
- Lovick, J. K., Ngo, K. T., Omoto, J. J., Wong, D. C., Nguyen, J. D., and Hartenstein, V. (2013). Postembryonic lineages of the *Drosophila* brain: I. development of the lineage-associated fiber tracts. *Dev. Biol.*, 384(2):228–257.
- Marin, E. C., Roberts, R. J. V., Büld, L., Theiss, M., Pleijzier, M. W., Sarkissian, T., Laursen, W. J., Turnbull, R., Schlegel, P., Bates, A. S., Li, F., Landgraf, M., Costa, M., Bock, D. D., Garrity, P. A., and Gregory S X (2020). Connectomics analysis reveals first, second, and third order thermosensory and hyposensory neurons in the adult *Drosophila* brain.
- Meissner, G. W., Nern, A., Singer, R. H., Wong, A. M., Malkesman, O., and Long, X. (2019). Mapping neurotransmitter identity in the whole-mount *Drosophila* brain using multiplex high-throughput fluorescence in situ hybridization. *Genetics*, 211(2):473–482.
- Nässel, D. R. (1999). Histamine in the brain of insects: a review. *Microscopy research and technique*, 44(2-3):121–136.
- Ohyama, T., Schneider-Mizell, C. M., Fetter, R. D., Aleman, J. V., Franconville, R., Rivera-Alba, M., Mensh, B. D., Branson, K. M., Simpson, J. H., Truman, J. W., et al. (2015). A multilevel multimodal circuit enhances action selection in *Drosophila*. *Nature*, 520(7549):633–639.
- Okada, R., Awasaki, T., and Ito, K. (2009). Gamma-aminobutyric acid

- (GABA)-mediated neural connections in the *Drosophila* antennal lobe. *J. Comp. Neurol.*, 514(1):74–91.
- Otto, N., Pleijzier, M. W., Morgan, I. C., Edmondson-Stait, A. J., Heinz, K. J., Stark, I., Dempsey, G., Ito, M., Kapoor, I., Hsu, J., et al. (2020). Input connectivity reveals additional heterogeneity of dopaminergic reinforcement in drosophila. *bioRxiv*.
- Prokop, A. and Meinertzhagen, I. A. (2006). Development and structure of synaptic contacts in *Drosophila*. *Semin. Cell Dev. Biol.*, 17(1):20–30.
- Rohlfing, T. and Maurer, C. R., J. (2003). Nonrigid image registration in shared-memory multiprocessor environments with application to brains, breasts, and bees. *IEEE Trans. Inf. Technol. Biomed.*, 7(1):16–25.
- Ryan, K., Lu, Z., and Meinertzhagen, I. A. (2016). The CNS connectome of a tadpole larva of *Ciona intestinalis* (L.) highlights sidedness in the brain of a chordate sibling. *Elife*, 5:e16962.
- Saalfeld, S., Cardona, A., Hartenstein, V., and Tomančák, P. (2009). Catmaid: collaborative annotation toolkit for massive amounts of image data. *Bioinformatics*, 25(15):1984–1986.
- Sayin, S., De Backer, J.-F., Siju, K. P., Wosniack, M. E., Lewis, L. P., Frisch, L.-M., Gansen, B., Schlegel, P., Edmondson-Stait, A., Sharifi, N., Fisher, C. B., Calle-Schuler, S. A., Lauritzen, J. S., Bock, D. D., Costa, M., Jefferis, G. S. X. E., Gjorgjieva, J., and Grunwald Kadow, I. C. (2019). A neural circuit arbitrates between persistence and withdrawal in hungry *Drosophila*. *Neuron*, 104(3):544–558.e6.
- Schneider-Mizell, C. M., Gerhard, S., Longair, M., Kazimiers, T., Li, F., Zwart, M. F., Champion, A., Midgley, F. M., Fetter, R. D., Saalfeld, S., et al. (2016). Quantitative neuroanatomy for connectomics in drosophila. *eLife*, 5:e12059.
- Sen, S. (2019). Neurotransmitter identity: A question of lineage. *eLife*, 8:e47162.
- Shinomiya, K., Takemura, S.-Y., Rivlin, P. K., Plaza, S. M., Scheffer, L. K., and Meinertzhagen, I. A. (2015). A common evolutionary origin for the ON- and OFF-edge motion detection pathways of the *Drosophila* visual system. *Front. Neural Circuits*, 9:33.
- Simonyan, K. and Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556.
- Staffler, B., Berning, M., Boergens, K. M., Gour, A., van der Smagt, P., and Helmstaedter, M. (2017). Synem, automated synapse detection for connectomics. *Elife*, 6:e26414.
- Tanaka, N. K., Endo, K., and Ito, K. (2012). Organization of antennal lobe-associated neurons in adult *Drosophila melanogaster* brain. *J. Comp. Neurol.*, 520(18):4067–4130.
- Turner-Evans, D. B., Jensen, K., Ali, S., Paterson, T., Sheridan, A., Ray, R. P., Lauritzen, S., Bock, D., and Jayaraman, V. (2019). The neuroanatomical ultrastructure and function of a biological ring attractor.
- Wilson, R. I. and Laurent, G. (2005). Role of GABAergic inhibition in shaping odor-evoked spatiotemporal patterns in the *Drosophila* antennal lobe. *J. Neurosci.*, 25(40):9069–9079.
- Xu, C. S., Januszewski, M., Lu, Z., Takemura, S.-y., Hayworth, K. J., Huang, G., Shinomiya, K., Maitin-Shepard, J., Ackerman, D., Berg, S., Blakely, T., Bogovic, J., Clements, J., Dolafi, T., Hubbard, P., Kainmueller, D., Katz, W., Kawase, T., Khairy, K. A., Leavitt, L., Li, P. H., Lindsey, L., Neubarth, N., Olbris, D. J., Otsuna, H., Troutman, E. T., Umayam, L., Zhao, T., Ito, M., Goldammer, J., Wolff, T., Svirskas, R., Schlegel, P., Neace, E. R., Knecht, C. J., Alvarado, C. X., Bailey, D. A., Ballinger, S., Borycz, J. A., Canino, B. S., Cheatham, N., Cook, M., Dreher, M., Duclos, O., Eubanks, B., Fairbanks, K., Finley, S., Forknall, N., Francis, A., Hopkins, G. P., Joyce, E. M., Kim, S., Kirk, N. A., Kovalyak, J., Lauchie, S. A., Lohff, A., Maldonado, C., Manley, E. A., McLin, S., Mooney, C., Ndama, M., Ogundeyi, O., Okeoma, N., Ordish, C., Padilla, N., Patrick, C., Paterson, T., Phillips, E. E., Phillips, E. M., Rampally, N., Ribeiro, C., Robertson, M. K., Rymer, J. T., Ryan, S. M., Sammons, M., Scott, A. K., Scott, A. L., Shinomiya, A., Smith, C., Smith, K., Smith, N. L., Sobeski, M. A., Suleiman, A., Swift, J., Takemura, S., Talebi, I., Tarnogorska, D., Tenshaw, E., Tokhi, T., Walsh, J. J., Yang, T., Horne, J. A., Li, F., Parekh, R., Rivlin, P. K., Jayaraman, V., Ito, K., Saalfeld, S., George, R., Meinertzhagen, I., Rubin, G. M., Hess, H. F., Scheffer, L. K., Jain, V., and Plaza, S. M. (2020). A connectome of the adult drosophila central brain. *bioRxiv*.
- Zheng, Z., Lauritzen, J. S., Perlman, E., Robinson, C. G., Nichols, M., Milkie, D., Torrens, O., Price, J., Fisher, C. B., Sharifi, N., et al. (2018). A complete electron microscopy volume of the brain of adult drosophila melanogaster. *Cell*, 174(3):730–743.

APPENDIX

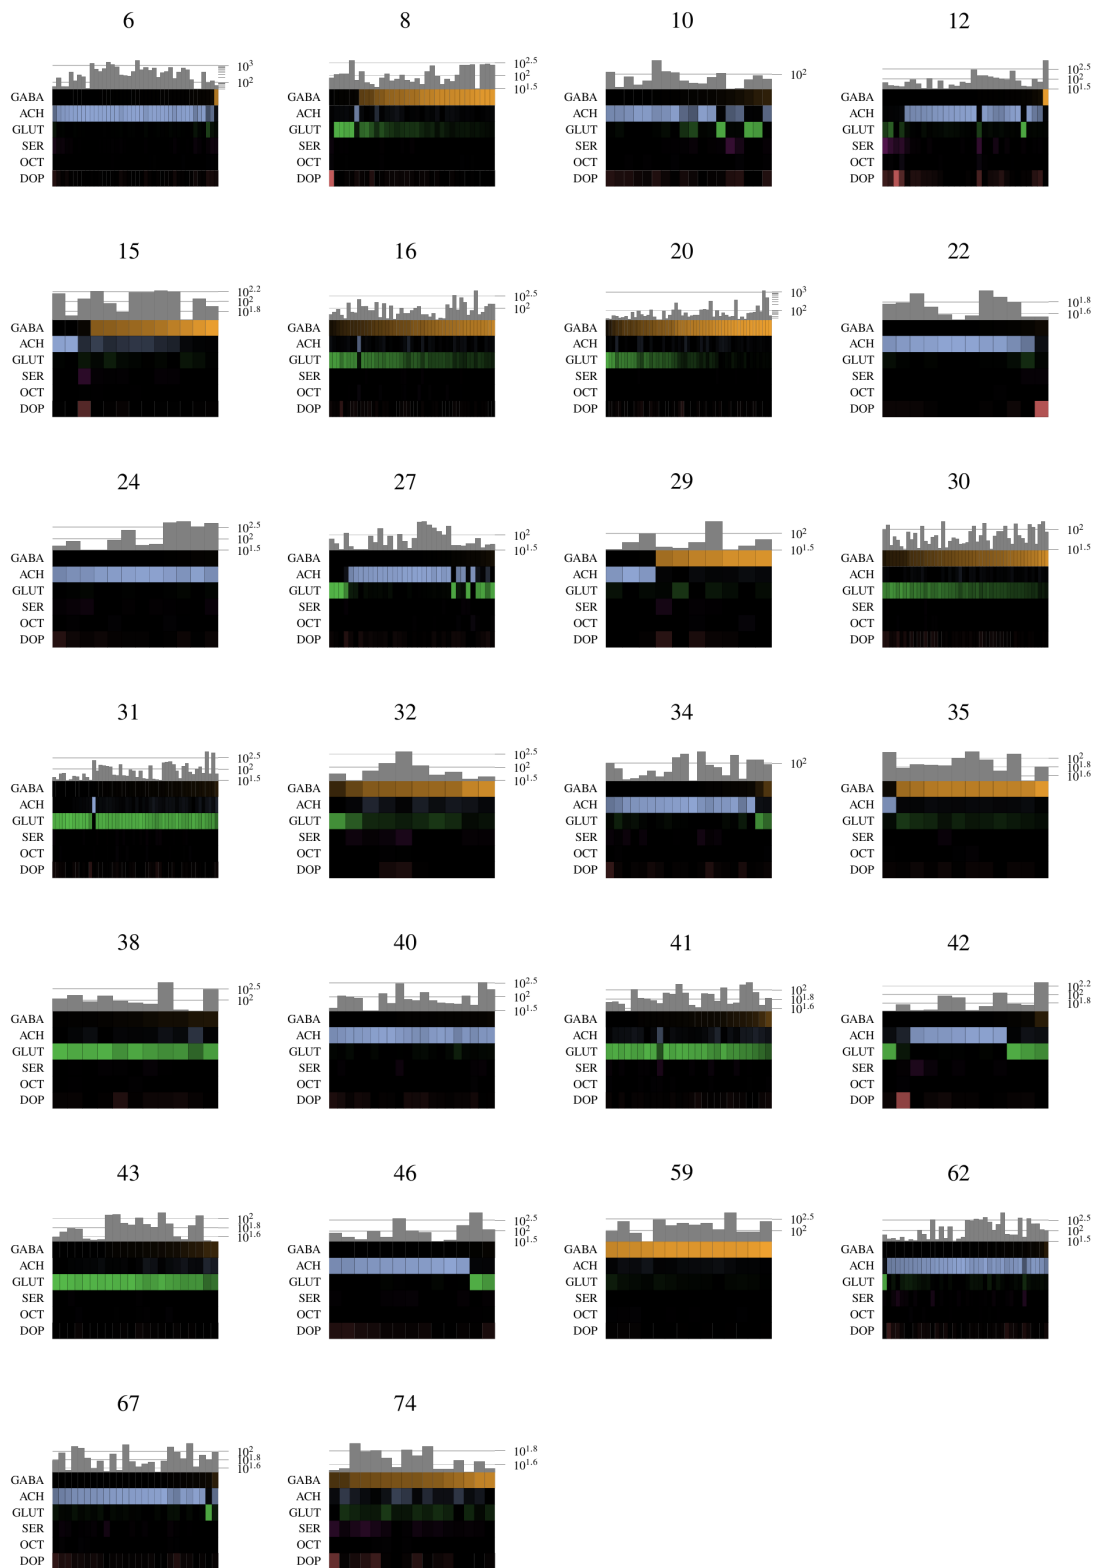


Figure 8: Neurotransmitter barcode plots of our predictions for all hemilineages that have more than 10 neurons with more than 30 synapses each. Each column represents one neuron. The relative number of synapses predicted as neurotransmitter $\hat{y} \in Y = \{GABA, ACh, GLUT, SER, OCT, DOP\}$ is represented by the color intensity of the respective row. Corresponding renderings of neurons and predicted neurotransmitters are shown in Fig. 9. For a mapping of hemilineage ID to hemilineage name see Table 2

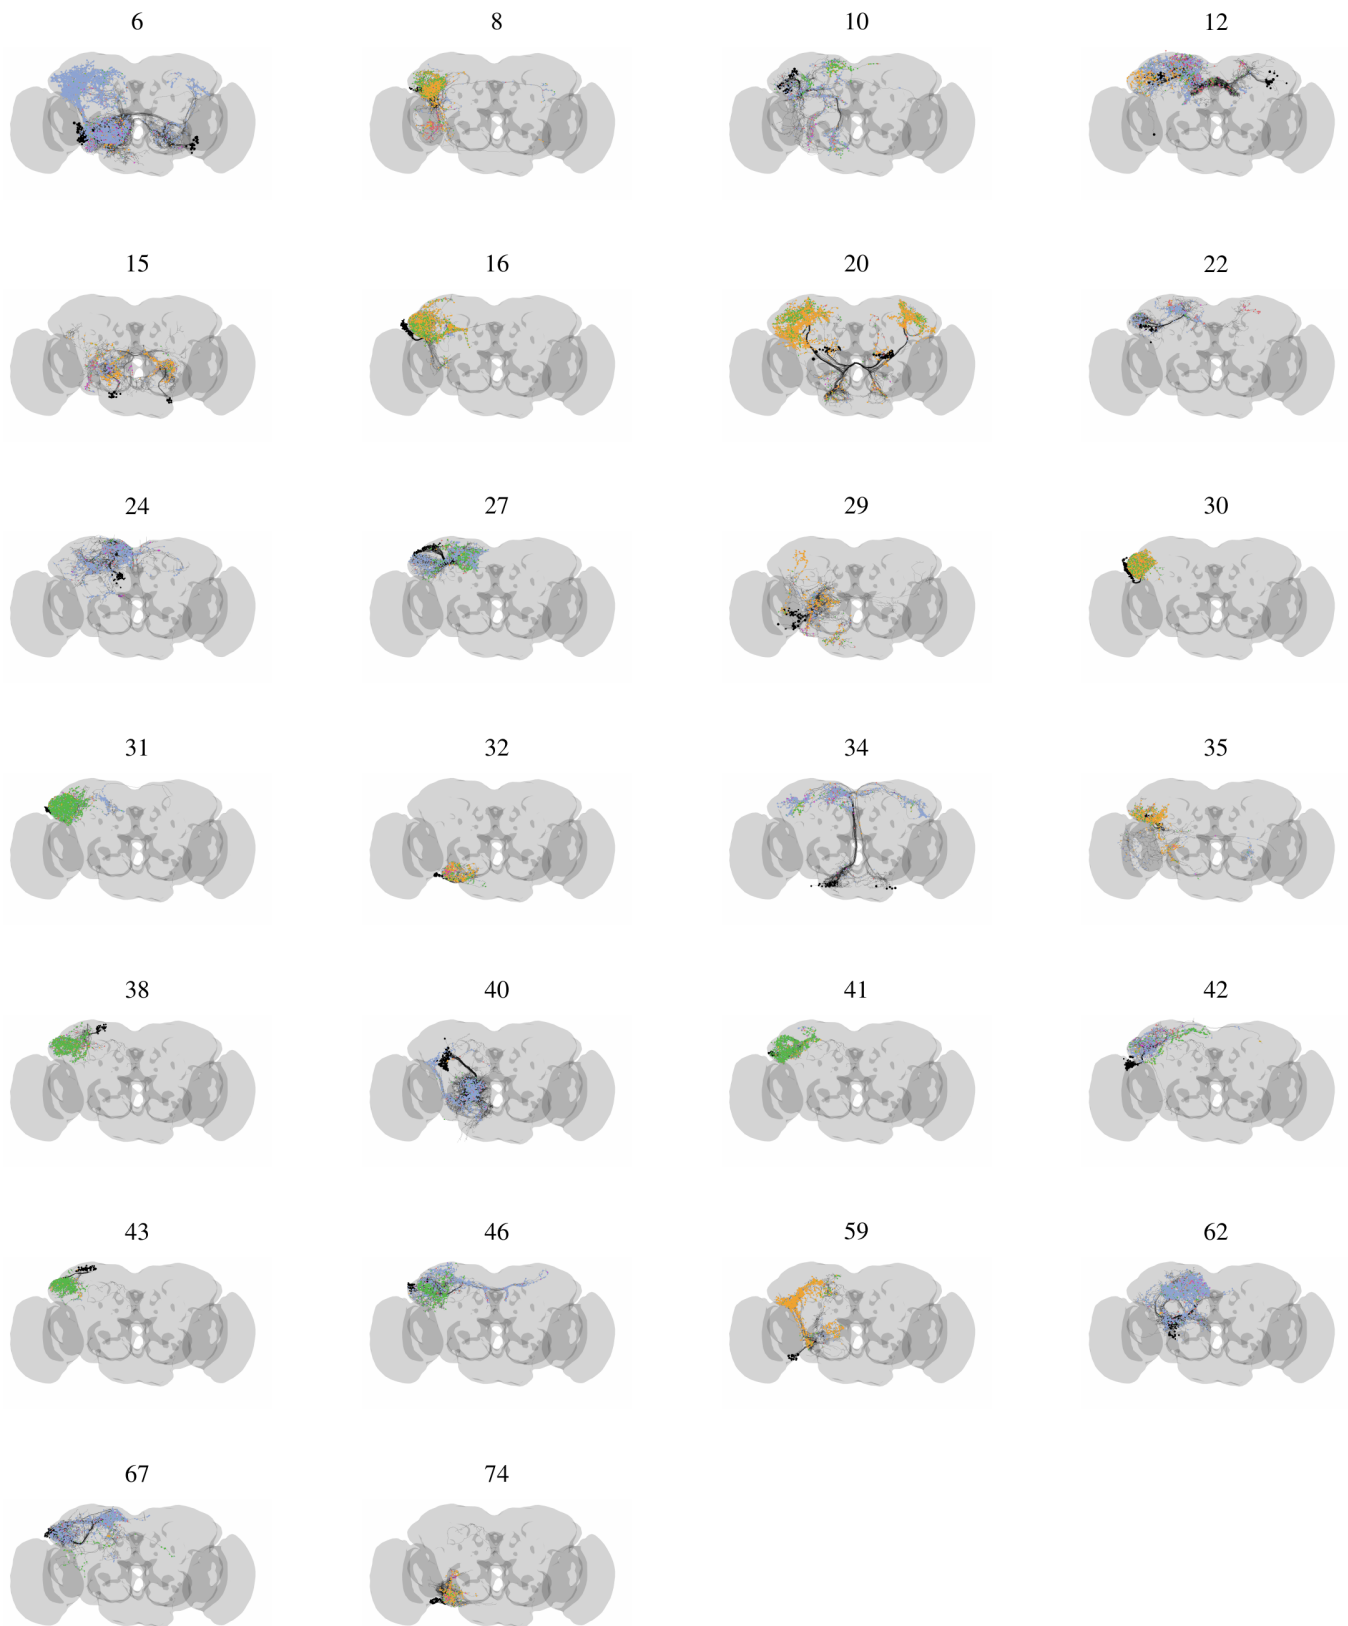


Figure 9: Renderings of neurotransmitter predictions for all neurons within hemilineages that have more than 10 neurons with more than 30 synapses each. Corresponding neurotransmitter barcode plots are shown in Fig. 8. For a mapping of hemilineage ID to hemilineage name see Table 2

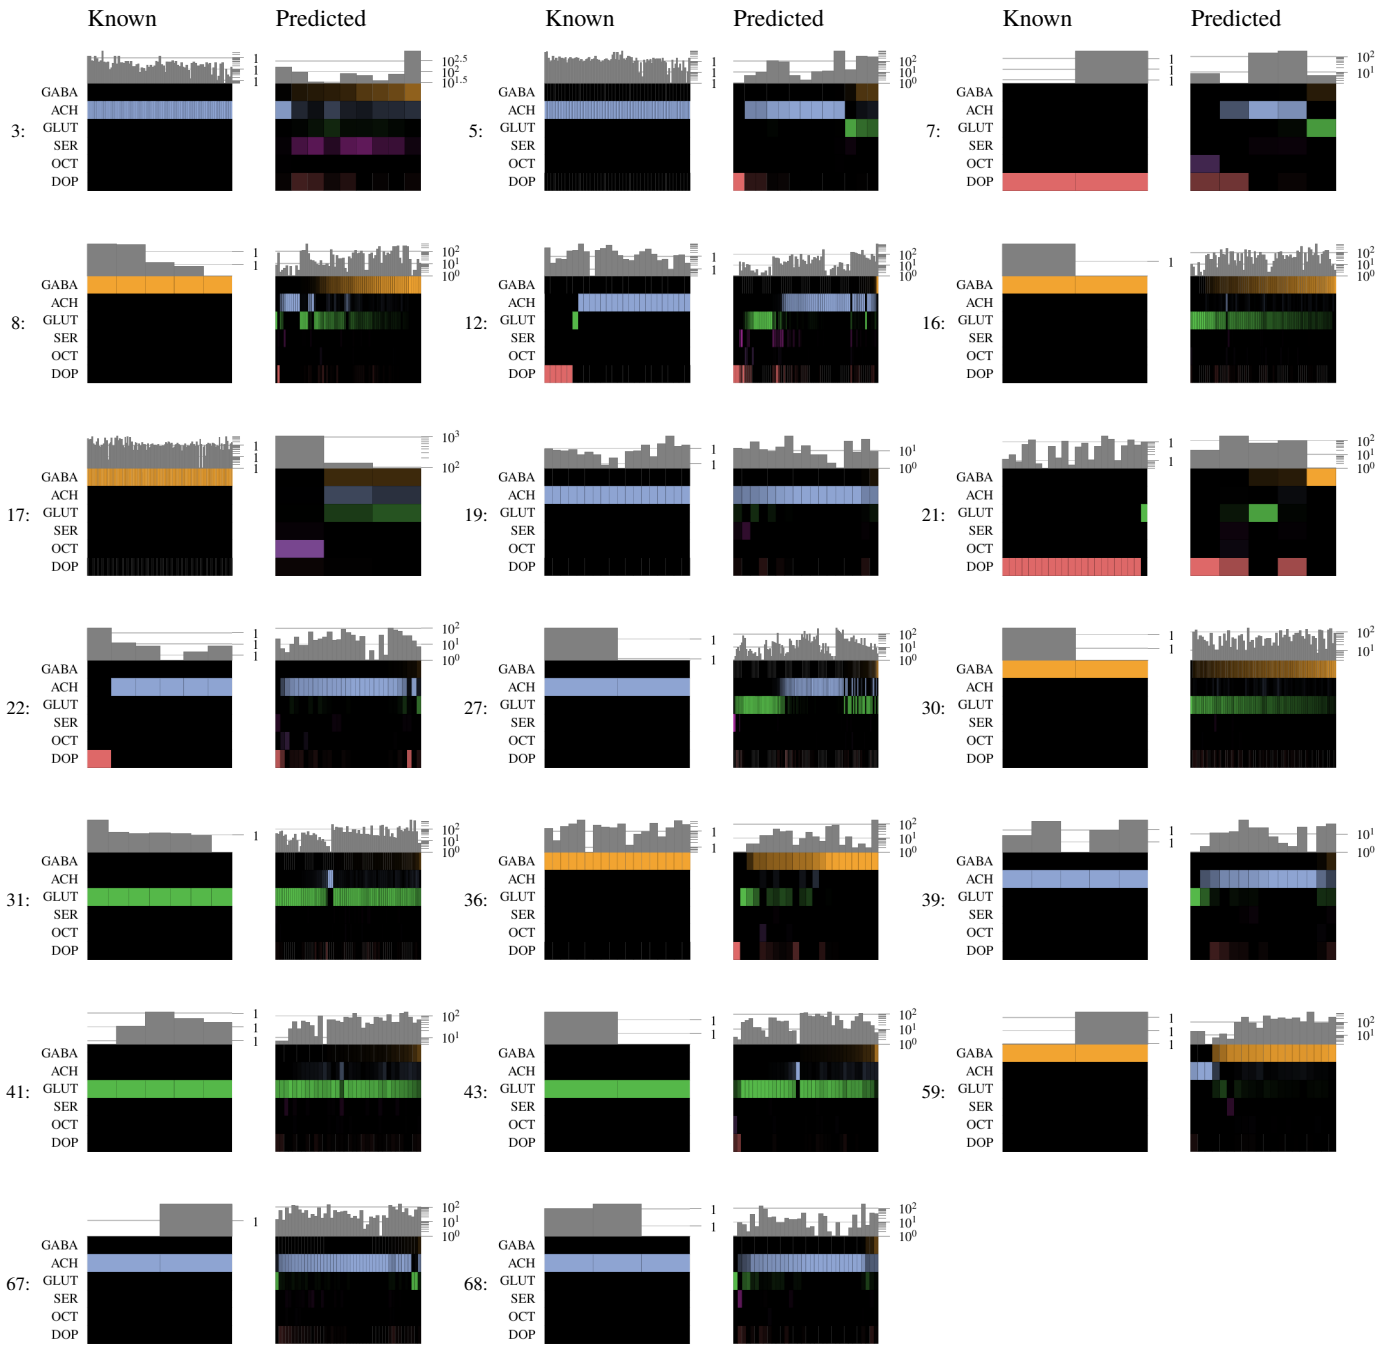


Figure 10: Neurotransmitter barcode plots of all hemilineages, for which a subset of neurons N_{gt}^h have genetically determined, known neurotransmitters (Known) and our predictions (Predicted) on the remaining neurons N_{pred}^h in the same hemilineage. Each column in the neurotransmitter barcode represents one neuron. The relative number of synapses with neurotransmitter y , $\hat{y} \in Y = \{GABA, ACh, GLUT, SER, OCT, DOP\}$ is represented by the color intensity of the respective row. Note that $N_{gt}^h \cap N_{pred}^h = \emptyset$. For a mapping of hemilineage ID to hemilineage name see Table 2

Hemilineage ID	ItoLee Name	Hartenstein Name	Hemilineage ID	ItoLee Name	Hartenstein Name
1	ALad1	BAmv3	45	LHa3	BLVa2
2	AOTUv1	DALcm2	46	LHp2 medial	DPLp1 medial
3	ALi1 dorsal	BAlc dorsal	47	VLPi2 dorsal	BLAv2 dorsal
4	SLPav1 lateral	BLA1 lateral	48	SLPad1 posterior	DPLi3 posterior
5	ALlv1	BAIp4	49	VLPi2 ventral	BLAv2 ventral
6	ALi1 ventral	BAlc ventral	50	SLPav3	BLVa2a
7	unnamed	unnamed	51	LHi2 lateral	DPLal2 lateral
8	VPNp and v1 posterior	BLP1 posterior	52	VESa1	BAla3
9	LHi4 posterior	BLD1 posterior	53	SMPpd1	DPLc1
10	VLPd1	DPLam	54	SLPal3 dorsal	BLAd3 dorsal
11	FLAa3	BAmas2	55	VLPi1	BLVa3/4
12	DL1 dorsal	CP2 dorsal	56	VLPi and p1 posterior	BLVp2 posterior
13	unnamed	BLP3 ventral	57	SIPa1 ventral	BLAd2 ventral
14	DM1	DPMm1	58	LHa2 ventral	BLVa1 ventral
15	LALv1	BAmv1	59	WEDA2	BAIp3
16	LHi1 lateral	BLD4 lateral	60	SIPa1 dorsal	BLAd2 dorsal
17	ALv1	BAla1	61	SLPav2 dorsal	BLD2 dorsal
18	LHd1	DPLd	62	VLPi and p2 posterior	BLVp1 posterior
19	SLPal2 ventral	DPLal3 ventral	63	SLPp1	DPLi1
20	CREa1 ventral	BAmD1 ventral	64	LHp2 lateral	DPLp1 lateral
21	CREa2	DALcm1	65	SLPp3 lateral	unnamed
22	DL2 dorsal	CP3 dorsal	66	VLPi and d1 lateral	BLAv1 lateral
23	SMPpv1	DPMp1	67	LHi2 medial	DPLal2 medial
24	SIPp1	DPMp2	68	VLPi and p1 anterior	BLVp2 anterior
25	DL1 ventral	CP2 ventral	69	VLPi and d1 dorsal	BLAv1 dorsal
26	SLPal1	DPLal1	70	AOTUv2	DALi1
27	VLPd and p1 posterior	DPLi2 posterior	71	AOTUv3 dorsal	DALcl1 dorsal
28	SMPpv2	CP1	72	VLPp and i1 ventral	DPLpv ventral
29	unnamed	unnamed	73	VPNd1	BLD6
30	LHi4 lateral	BLD1 lateral	74	WEDA1 dorsal	BAIv
31	LHp1	BLP4	75	VLPd and p1 anterior	DPLi2 anterior
32	WEDA1 ventral	BAIv	76	unnamed	unnamed
33	SLPad1 anterior	DPLi3 anterior	77	VLPp1	BLP2
34	FLAa2	BAmas1	78	PSa1	BAIp1
35	VLPp and i1 dorsal	DPLpv dorsal	79	Primary	Primary
36	EBa1	DALv2	80	VPNp1 posterior	BLD5 posterior
37	SMPad2	DAMd2/3	81	unnamed	unnamed
38	SLPpm1	DPLm1	82	DL2 ventral	CP3 ventral
39	SLPal2 dorsal	DPLal3 dorsal	83	LHa2 dorsal	BLVa1 dorsal
40	WEDd1	DALd	84	SLPp3 posterior	unnamed
41	SLPp and v1 posterior	DPLp2 posterior	85	PSp3	DPMi1/2
42	LHa1 medial	BLAd1 medial	86	SLPa and i1 lateral	BLA1 lateral
43	LHd2 dorsal	DPLm2 dorsal	87	SLPa and i1 anterior	BLAvm1 anterior
44	VPNi and d1 dorsal	BLAvm2 dorsal	88	AOTUv3 ventral	DALcl1 ventral
			89	CLp1	DPLc2/4

Table 2: Mapping of hemilineage ids to ItoLee and Hartenstein hemilineage names.

Operation	Size	Feature Maps
Conv (1)	(3,3,3)	8
BatchNorm		8
ReLU		8
Conv (2)	(3,3,3)	8
BatchNorm		8
MaxPool	(1,2,2)	8
Conv (3)	(3,3,3)	16
BatchNorm		16
ReLU		16
Conv (4)	(3,3,3)	16
BatchNorm		16
MaxPool	(1,2,2)	16
Conv (5)	(3,3,3)	32
BatchNorm		32
ReLU		32
Conv (6)	(3,3,3)	32
BatchNorm		32
MaxPool	(1,2,2)	32
Conv (7)	(3,3,3)	64
BatchNorm		64
ReLU		64
Conv (8)	(3,3,3)	64
BatchNorm		64
MaxPool	(1,2,2)	64
Linear	4096	1
ReLU		1
Dropout		1
Linear	4096	1
ReLU		1
Dropout		1
Linear	6	1

Table 3: Best performing 3D-VGG-type architecture used.

Parameter	Value	Augmentation	Parameter	Value
Input Shape	(16, 160, 160)	Elastic	control point spacing	(4,40,40)
Loss	CrossEntropy		jitter sigma	(0, 2, 2)
Optimizer	Adam		subsample	8
Learning Rate	1E-04	Rotation	axis	z
β_1	0.95		angle	in $[0, \frac{\pi}{4}]$
β_2	0.999	Section Defects	slip probability	0.05
Iterations	315,000		shift probability	0.05
			max misalign	10
		Mirror	n/a	
		Transpose	axes	x, y
		Intensity	scale	in [0.9, 1.1]
			shift	in [-0.1, 0.1]

Table 4: Training parameters for best performing model. Augmentations from <http://funkey.science/gunpowder>.