Interneuron diversity and normalization specificity in a visual system

H. Sebastian Seung
Neuroscience Institute and Computer Science Dept.
Princeton University, Princeton, NJ 08544

Apr 3, 2024

Abstract. Normalization is a fundamental operation in image processing. Convolutional nets have evolved over the past decade to include a large number of normalizations, and this architectural shift has proved essential for robust visual artificial intelligence. Here I argue that normalization is the function of a large fraction of neuronal cell types in the fly optic lobe. A dozen types of local interneuron in the distal medulla (Dm), in particular, are argued to mediate lateral inhibition, which was classically assumed to depend only on spatial separation. Each Dm type turns out to connect with specific source and target cell types. The source and target are either the same type, or are directly connected, from which it follows that Dm-mediated lateral inhibition has a normalizing function. Therefore the diversity of Dm interneurons is analogous to the ubiquity of normalizations in contemporary convolutional nets. A notable difference is that Dm types normalize over specific spatial scales. A final Dm type is an outlier in its perfect tiling of the medulla, and likely does not mediate lateral inhibition due to a predicted electrical compartmentalization of its arbor. Further candidate normalizers are identified in all other interneuron families, generalizing the Dm findings to the entire optic lobe.

Normalization is often the first step of image processing. Normalization can be subtractive, to center the distribution of image pixel values. And normalization can be divisive (Carandini and Heeger 2011), to control the width of the distribution. In convolutional nets, currently the dominant approach to visual artificial intelligence, normalization is not only the first step, but occurs repeatedly during image processing: every feature map in every layer of a convolutional net is accompanied by a normalization (Ioffe and Szegedy 2015; Ulyanov, Vedaldi, and Lempitsky 2016).

Normalization might seem prosaic relative to convolution, the feature detecting operation that gives convolutional nets their name. But empirical evidence has accumulated over the past decade that normalizations are essential for making convolutional nets robust (He et al. 2015; Bjorck et al. 2018; Santurkar, Tsipras, and Ilyas 2018). Like supporting actors, normalizations are essential but often underappreciated.

A neuronal wiring diagram of a biological visual system, that of the fruit fly *Drosophila melanogaster*, has recently been completed (Zheng et al. 2018; Dorkenwald et al. 2023; Schlegel et al. 2023; Matsliah et al. 2023). A cell type in the fly optic lobe consists of many cells performing the same computation at different locations in the visual field, and is analogous to a feature map in a convolutional net (Lappalainen et al. 2023). The connections from one cell type to another can be averaged over location to estimate a convolutional kernel (Lappalainen et al. 2023), and this approach was used to identify cell types that are predicted to detect features like bars and corners (Seung 2023).

Given that the optic lobe implements convolutions, one might wonder whether normalization is also ubiquitous. Of the 200+ cell types intrinsic to the optic lobe (Matsliah et al. 2023), well over half are spatially confined to a single neuropil of the optic lobe (Fig. 1a). These are analogous to the local

interneurons found in mammalian brains, and are likewise typically inhibitory. My claim is that many of these interneuron types are devoted to normalization.

Since general claims in biology are best argued through specific examples, I will focus on the Dm interneuron family, which consists of 23 types found in the distal medulla (Fig. 1b, c). The functions of 13 Dm types have never been investigated (Fig. 1c). I will argue that almost all of this baker's dozen of Dm types support normalization by mediating lateral inhibition, a kind of connectivity first hypothesized in the 19th century to explain visual psychophysics (Mach 1868; Ratliff 1965), and first observed directly in the 20th century in the *Limulus* retina (Hartline, Wagner, and Ratliff 1956).

Lateral inhibition was classically assumed to depend only on spatial separation, so that cells at one location inhibit cells at neighboring locations. Lateral inhibition in the fly distal medulla turns out to have an extraordinary degree of specificity, because the relevant Dm types are connected with particular source and target columnar types.

Lateral inhibition by itself is not necessarily interpretable in terms of normalization. Evidence for that interpretation will be provided by showing that the source type of a Dm type is directly connected to the target type, or source and target are the same type. The implication is that the input to a target neuron is the activity of a source neuron normalized relative to activities of neighboring source neurons.

In contemporary convolutional nets, feature maps are normalized over the entire visual field (Ioffe and Szegedy 2015; Ulyanov, Vedaldi, and Lempitsky 2016). Dm-mediated inhibition normalizes over a neighborhood, which is expected to improve robustness to variations in illumination across space. The size of the neighborhood varies across Dm types, so normalization in the distal medulla is multiscale, another potential source of robustness.

The thirteenth Dm type (Dm4) in the baker's dozen is the exception that proves the rule. Its perfect tiling of the medulla and looping self-fasciculating neurites make it an outlier in the baker's dozen. It is predicted to not mediate lateral inhibition. Rather, each column of a single Dm4 arbor is predicted to function as an independent unit, much as in the CT1 cell (Meier and Borst 2019).

In addition to the detailed study of Dm interneurons, I also search for candidate normalizers in all other interneuron families in the optic lobe. The long list of search results strongly suggests that normalization is a ubiquitous computation throughout the optic lobe, with specificity and diversity similar to that observed in the Dm interneuron family.

More broadly, the present work addresses a question that also arises in studies of mammalian neuronal diversity: why are there so many types of inhibitory interneuron? Transcriptomic studies have identified 60 interneuron types in mouse visual cortex (Tasic et al. 2018), while functional studies differentiate between only a handful of broad inhibitory classes (Kepecs and Fishell 2014). In the fly optic lobe, the reason for a large number of interneuron types may be simple. If interneurons perform the computation of normalization, but each type supports specific source and target types and a specific spatial scale, there must inevitably be many interneuron types.

Results

Local interneurons of the distal medulla (Dm)

Dm types have been traditionally defined by their stratification in layers of the medulla (Fig. 1b), their tangential arbor sizes (Fig. 1b, g), and their molecular properties (Fischbach and Dittrich 1989; Takemura et al. 2013; Nern, Pfeiffer, and Rubin 2015). For reliable discrimination, connectivity should also be used (Matsliah et al. 2023). The cells of each Dm type cover the medulla with overlap (Fig. 1g).

Predictions of neurotransmitter identity based on the electron microscopic images are provided by FlyWire (Eckstein et al. 2020). Predictions based on transcriptomic measurements are also available (Davis et al. 2020). Most Dm types are predicted to secrete glutamate or GABA, neurotransmitters that are usually inhibitory.

Of the twenty-three Dm types (Matsliah et al. 2023), seven (yDm8, pDm8, Dm9, DmDRA1, DmDRA2, Dm11, and Dm2) are known to receive direct input from photoreceptors R7 and R8, so they are presumably related to color and polarization vision (Schnaitmann, Pagni, and Reiff 2020). In a companion paper, I have provided evidence that Dm3 types (Dm3v, Dm3p, and Dm3q) are likely related to form vision (Seung 2023).

The 13 remaining Dm types (Dm1, Dm4, Dm6, Dm10, and Dm12 through Dm20) with completely unknown functions are the topic of the present work. Of this Dm baker's dozen, ten types contain less than 100 cells each, and three types (Dm12, Dm15, and Dm10) each contain between 100 and 300 cells (Fig. 1c). In contrast, most of the Dm types with existing functional annotations contain larger numbers of cells (Fig. 1c).

Columnar types in the distal medulla

The medulla is divided into 800 columns, which are perpendicular to its layers, and are in one-to-one correspondence with the ommatidia of the compound eye. The main synaptic partners of the Dm baker's dozen are columnar neurons (Fig. 1d). These 12 partner types (with one exception) are confined to a single column, and contain about 800 cells (Fig. 1e). The activity of such a "unicolumnar" neuron mainly encodes information about the visual stimulus at or near a single ommatidium. The exception Lawf1 consists of <200 cells, each of which extends over many columns.

Lamina monopolar cells (L2, L3, L4, and L5) project from the lamina to the distal medulla. L2 and L3 receive direct input from photoreceptors in the lamina. (L1 has little or no connectivity with Dm types.) Tm1, Tm2, and Tm3 project from distal medulla to lobula. Mi4 and Mi9 project from distal to proximal medulla. T1 and Lawf1 project from distal medulla to lamina. T2 projects from distal and proximal medulla to lobula. All types are predicted cholinergic by FlyWire, except that Mi4 is predicted GABAergic, and Mi9 is predicted glutamatergic.

Lateral inhibitory fields

The shapes and orientations of the Dm arbors already hint at their functions. The arbors are typically "wide" enough to connect columnar neurons corresponding to well-separated ommatidia (Fig. 1b). Columnar types usually send information forward (from the photoreceptors towards the central brain) or

sometimes backward (Lawf1 in Fig. 1d). Dm arbors look capable of sending information “sideways,” and this lateral spread of information enables the activity of a columnar cell to be normalized relative to neighboring cells. The following examples will show that the normalization mediated by a Dm type can be highly specific to particular columnar types.

Dm10 and Dm15 are the most numerous of the types in the Dm baker’s dozen (Fig. 1c). Clues to their functions can be gleaned from their presynaptic and partner types. L5 is the strongest input to Dm10 (Figs. 2a, S1), and Mi4 is the strongest output of Dm10 (Figs. 2a, S2). Strengths are quantified in Fig. 2a following the key in Fig. 2b.

The indirect pathway L5-Dm10-Mi4 is paralleled by a direct L5-Mi4 connection (Fig. 2a). L5 is cholinergic and presumed excitatory, while Dm10 is predicted GABAergic and presumed inhibitory. It follows that the direct pathway is excitatory, while the indirect pathway is net inhibitory. Since both direct and indirect pathways converge onto the same target, the Dm10 circuit motif (Fig. 2a) will be called “convergent feedforward inhibition,” or “feedforward inhibition” for short.

For Dm10 the strongest input and output were different columnar types. For Dm15, it turns out that the strongest input and output are the same columnar type, L2 (Fig. 2c, S1, S2). Dm15 is predicted glutamatergic and presumed inhibitory. The motif of Fig. 2b will be called “recurrent inhibition,” meaning that L2 inhibits itself through Dm15.

It might seem counterproductive for L5 to exert opposite effects on Mi4 through the direct and indirect pathways in Fig. 2a. But this diagram ignores spatial organization. Through the direct pathway, Mi4 is excited by L5 in the same column, with negligible input from L5 in neighboring columns (Fig. 2d). Tracing all indirect pathways from L5 cells to Dm10 cells to an Mi4 cell yields an estimate of the L5-Dm10-Mi4 “inhibitory field” (Fig. 2e). This is an estimate of how much the activity of an Mi4 cell is inhibited by L5 activity in neighboring columns.

Figs. 2d and 2e mean that an Mi4 cell compares its direct input from the L5 cell in the same column with Dm10-mediated indirect input from L5 cells in neighboring columns. The lateral inhibition has a “normalizing” function, so that Mi4 activity is driven not by absolute L5 activity, but L5 activity relative to its neighbors.

The above structural analysis can be compared with physiological measurements of the Mi4 receptive field. According to a white noise analysis, the Mi4 receptive field has an ON center that is roughly one column wide (Arenz et al. 2017). The direct pathway L5-Mi4 (Fig. 2d) is likely a major contributor to the ON center, as L5 is the strongest input of Mi4 (Fig. 2a, S1), L5 is ON excitatory, and L5-Mi4 connectivity is uni-columnar.

White noise analysis also reveals an antagonistic OFF surround of the Mi4 receptive field (Arenz et al. 2017). The L5-Dm10-Mi4 inhibitory field extends to the nearest neighbor columns and more weakly to the next nearest neighbor columns (Fig. 2f, bottom), which roughly matches the Mi4 receptive field surround from white noise analysis. Therefore the L5-Dm10-Mi4 pathway (Fig. 2e) likely contributes to the Mi4 receptive field surround.

The reader should be cautioned that Mi4 has other input types besides the ones shown in Fig. 2a (Fig. S3). The Mi4 receptive field center and surround may involve a number of overlapping mechanisms. What is interesting and novel is that at least one of these mechanisms (L5-Dm10-Mi4) can be identified with much more specificity than the generic term “lateral inhibition.”

A similar analysis predicts the L2-Dm15-L2 inhibitory field (Fig. 2g). In this case, the lateral inhibition is also said to be normalizing, because it enables the L2 cell to compare its own activity with that of its neighbors.

The inhibitory field is anisotropic, with a horizontal orientation. This is reminiscent of the old observation of a horizontally oriented inhibitory field in the *Limulus* eye (Barlow 1969). The L2-Dm15-L2 inhibitory field is expected to give rise to a horizontally oriented surround for the L2 receptive field. Such an anisotropic surround has not been reported in experiments with white noise stimuli (Drews et al. 2020). The effect might be subtle, as L2 receives input from many sources (Fig. S3) and the L2 receptive field is likely created by multiple overlapping mechanisms.

The role of GABAergic inhibition in creating the surround of the L2 receptive field was previously investigated (Freifeld et al. 2013). Knockdown of GABA receptors in L2 cells alone had no effect on the receptive field. This is not inconsistent with the idea that Dm15 contributes to the L2 surround, because Dm15 is predicted glutamatergic.

Dm cells are dominated by one or two input types

The above analysis raises two questions. First, the Dm10 and Dm15 circuit motifs (Fig. 2a, c) depicted only the strongest input and output to these Dm types. Is it really a good approximation to neglect the others? Second, does the analysis extend from Dm10 and Dm15 to other Dm types?

These questions can be addressed by examining the bar graphs of Fig. S1, which show the top four input fractions for all Dm types. The strongest input to a Dm type (the leftmost bar in the graphs) will be called the “winner.” For each Dm type, the winner is always some type of lamina monopolar cell (L2 through L5, Fig. S1), except that Tm2 is the winner for Dm16. The second strongest input to a Dm type will be called the “runner-up.” For most Dm types, the winner is much larger than the runner-up and all other inputs. The exceptions are Dm13, Dm4, and Dm20, for which the runner-up is nearly tied with the winner. Therefore the following rule holds:

Rule 1. All Dm types in the baker’s dozen are dominated by one or two input types.

This notion of dominance can be formalized by ordering the inputs (as in Fig. S1), and then quantifying the fractional decrease of each successive input. Find the smallest i such that input i+1 is less than γ times input i. Then we say that inputs 1 through i are dominant. Setting γ=0.6, we see that either one or two inputs are dominant for each type. These are the inputs above the red line in Fig. S1, the height of which is set by γ times input i. The choice of threshold parameter γ is somewhat arbitrary, but modest changes do not qualitatively change the overall picture of high input specificity.

Predicted contrast polarities of Dm visual responses

If dominated by a single winning input, a Dm type is expected to inherit some of the winner’s visual response properties. In particular, most Dm types are predicted to be OFF cells, because they are dominated by the OFF excitatory types L2, L3, L4, or Tm2 (Fig. S1). Exceptions are Dm1 and Dm10, which are predicted to be ON cells since they are dominated by the ON excitatory type L5.

The subtle cases are Dm13, Dm4, and Dm20, for which the runner-up is comparable to the winner. Dm20 is predicted to be OFF, as its top two inputs are OFF excitatory (L3) and ON inhibitory/GABAergic (Mi4).

Dm13 and Dm4 are predicted to be ON-OFF, because Dm13 receives both ON excitatory (L5) and OFF excitatory (Tm2) input, and Dm4 receives OFF excitatory (L3) and OFF inhibitory/glutamatergic (Mi9) input.

Dominant outputs of Dm types

The output fractions of each Dm type are shown in Fig. S2. Seven types are dominated by a single output, and three types are dominated by two outputs. A few types are dominated by large output sets (Dm1 and Dm6 by four outputs, and Dm19 by five outputs). The top outputs of the Dm cells are more diverse than the top inputs (Fig. S2). They include all the columnar types that are top input of some Dm cell (L2, L3, L5, Tm2) except for L4. In addition, the top outputs include Tm1, Mi1, Mi4, T2, and Lawf1. All of the top outputs are confined to single columns, except for Lawf1 which extends over multiple columns. For five Dm types, the top output is the same as the top input.

Dm types mediate lateral inhibition that is normalizing

Fig. 3 depicts the connectivity of each Dm type with its dominant inputs and outputs. Also drawn are direct connections between inputs and outputs. For clarity, only connections with input fraction or output fraction greater than 10% are included; weaker edges are suppressed. For the Dm types that are dominated by a single input, the following simple rule almost always holds:

> Rule 2. The output target of a Dm type is either the same as its dominant input source (recurrent inhibition), or cell types that are directly downstream of its dominant input source (feedforward inhibition).

This promotes the specific example motifs in Figs. 2a and 2b to be a general rule for the Dm baker's dozen. Rule 2 means that the activity of a cell in one column is compared with the activities of cells of the same type in neighboring columns. In other words, the lateral inhibition has a normalizing function.

Dm12 (Takemura et al. 2015; Xu et al. 2018) and Dm18 are reciprocally connected with L3 and L5, respectively (Fig. 3a), and therefore conform to recurrent inhibition, similar to Fig. 2b.

Dm14 conforms to feedforward inhibition (Fig. 3b), just like the Dm10 motif that was in Fig. 2a and is repeated in Fig. 3b for convenience. The Dm19 motif looks more complex (Fig. 3c), but it is simply four feedforward inhibition motifs that overlap.

Dm15 and Dm17 combine recurrent and feedforward inhibition (Fig. 3d). For example, one target of Dm15 (Tm1) is feedforward inhibition, while the other target (L2) is the recurrent inhibition that was shown in Fig. 2b.

Overall, each Dm type tends to devote large fractions (>10%) of its input and output synapses to connections that conform to either recurrent or feedforward inhibition. In this sense, it seems apt to say that the Dm type "serves" or "supports" its dominant input. On the other hand, the dominant input usually devotes a small fraction (<10%) of its output synapses to the Dm type, and the targets of the

dominant input usually receive a small fraction (<10%) of its input synapses from the Dm type. In this respect, the relationship between the Dm type and the columnar types is asymmetrical.

Dm20 and Dm4 do not fit Rule 2 at all (Figs. 3g, h). These outliers will be discussed in detail later on.

ON-OFF opponency as a secondary function

As combinations of recurrent and feedforward inhibition, the Dm16, Dm6, and Dm13 diagrams (Fig. 3e) largely conform to Rule 2, but also contain “dangling” outputs Mi9 and Mi4 and input L5, which do not conform. It turns out that the Dm16 diagram does not truly lack a connection from Tm2 to Mi9; there is a weak connection that does not show up because it is below the chosen threshold for being displayed. So Dm16 actually does conform to Rule 2.

The other two diagrams of Fig. 3e make functional sense provided that response polarity is taken into account. Because L2 is OFF while Mi4 is ON, the L2-Dm6-Mi4 pathway can be interpreted as mediating an opponent interaction between ON and OFF. Similarly, L5 is ON while Tm2 is OFF, so the L5-Dm13-Tm2 pathway can be interpreted as mediating ON-OFF opponency. According to this idea, the “dangling” types in Fig. 3e violate Rule 2 because their function is not normalization, but opponency.

Based on its rarity, ON-OFF opponency seems only a secondary function of the Dm types shown in Fig. 3e. Normalization appears to be the primary function.

Normalization is multiscale

Figure 3 depicted connectivity from the viewpoint of each Dm type. It is also helpful to visualize connectivity from the viewpoints of columnar types. Three Dm types provide recurrent inhibition to L2 (Fig. 4a), and two Dm types provide recurrent inhibition to L5 (Fig. 4b). This might seem redundant, but visual examination of the types reveals that their properties are quite different. For the Dm partners of L2, the number of cells per type varies from 200 to 30 to 4 (Dm15, Dm6, Dm17, Fig. 1c). The arbors of these types likewise range from small to medium to large (Fig. 4g). Dm6 is predicted GABAergic, Dm15 is predicted glutamatergic, and Dm17 has no confident prediction.

L2 cells extend from the lamina to the medulla, where they synapse onto Dm15 (Fig. 4c), Dm6 (Fig. 4d), and Dm17 cells. Example maps of L2-Dm connectivity are shown in Fig. 4e. These maps can be regarded as predictions of Dm receptive fields. Example maps of Dm-L2 connectivity are shown in Fig. 4f. Through a Dm cell, an L2 cell on one side of the Dm arbor can indirectly inhibit an L2 cell on the other side of the arbor. The L2-Dm15-L2 inhibitory field (Fig. 2h) is wider than the L2-Dm15 (Fig. 4e, top) and Dm15-L2 (Fig. 4f, top) maps, because the inhibitory field of an L2 cell is mediated by the multiple Dm15 cells that are presynaptic to it. L2-Dm6-L2 and L2-Dm17-L2 inhibitory fields (data not shown) are wider than the L2-Dm15-L2 inhibitory field (Fig. 2h), because the range of lateral inhibition is set by the spatial extent of the Dm arbor.

To summarize, lateral inhibition between L2 neurons is mediated by three Dm types that likely operate at different length scales. The example cells (Fig. 4e, f) illustrate the range of variation across types. Similarly, lateral inhibition between L5 neurons is mediated by two Dm types that likely operate at different length scales (Fig. 4b). A more complete picture of intra-type as well as inter-type variation is provided by Fig. 4g. Dm1 and Dm18, which are both coupled to L5, have distinct size distributions,

again consistent with the idea that a single columnar type can be simultaneously normalized by multiple Dm types at different length scales.

Dm4 is an outlier that is unlikely to mediate lateral inhibition

The above analysis implicitly assumed Dm cells are electrotonically compact, meaning that synaptic input causes an electrical effect that spreads across the entire cell with little or no attenuation. This seems plausible for the small Dm10 and Dm15 cells of Fig. 2, but the assumption could be wrong for large Dm cells. If so, then lateral inhibition would not spread as far as the width of the Dm arbor. This possibility should be examined more closely.

Each Dm type covers the medulla, in the sense that every column is covered by at least one cell of that type. In vertebrate retina, the coverage factor of a neuron type has been defined as the average number of neurons of that type covering a single point in the retina, and has been estimated by physiological and anatomical methods (Sterling 1983). Here coverage factor can be quantified for the first time using anatomical connectivity. For a given Dm type, choose a columnar type that is a synaptic partner. For all cells of the columnar type, count the number of partner cells of the Dm type, and then average over columnar cells. This is done in Fig. 5a using the strongest input type for each Dm type. Most of the Dm baker's dozen have coverage factors close to two or three.

The two outliers are Dm4 and Dm20. Dm4 has the lowest coverage factor, which is close to one (Fig. 5a). This is consistent with visual inspection, which reveals that Dm4 cells overlap very little, fitting together beautifully, almost like floor tiles (Fig. S4). Dm20 has the highest coverage factor, which exceeds five (Fig. 5a). Interestingly, Dm4 and Dm20 are also the types that violate Rule 2, and are outliers in the motifs of Fig. 3.

If lateral inhibition were the function of Dm4, then the almost perfect tiling of Dm4 would have marked effects near the borders of Dm4 arbors. This is illustrated by three adjacent Mi4 cells in Fig. 5b. The middle Mi4 cell receives input from two Dm4 cells, while the flanking Mi4 cells each receive input from a single Dm4 cell. Fig. 5c shows the consequences for L3-Dm4-Mi4 connectivity maps. The middle Mi4 cell receives disynaptic input from L3 cells on its left and right. The left Mi4 cell receives input from L3 cells on its left, and the right Mi4 cell receives input from L3 cells on its right. Such marked differences in the inhibitory fields of adjacent Mi4 cells would be observable through visual physiology, and would contradict the concept of translation invariance, that cells of the same type perform the same computation at different locations in the visual field.

However, there is an alternative scenario. Perhaps Dm4 does not mediate lateral inhibition after all. Zooming into the arbor of a Dm4 cell reveals that its neurites take amazingly circuitous paths to travel from one column to the next (Fig. 5d). The neurite path length between two neighboring columns turns out to exceed 200 μm, even though the columns are separated by a Euclidean distance of just 7 μm. This is because of the looping and auto-fasciculation of the Dm4 arbor. This is not the case for the arbors of other Dm types, in which neurites take straighter paths.

I predict that the long looping neurites serve to electrically isolate columns from each other. Each column of the Dm4 arbor is predicted to be an independent functional unit. If this is the case, Dm4 mediates unicolumnar interactions, and the connectivity maps of Fig. 5c are not actually lateral inhibitory fields. Such an "extreme compartmentalization" has been observed in the CT1 amacrine cell of the optic lobe (Meier and Borst 2019).

Dm20 mediates feedback inhibition

The Dm20 graph (Fig. 3g) contains a 3-cycle L3→Dm20→Lawf1→L3. Lawf1 is predicted cholinergic, consistent with transcriptomic evidence (Davis et al. 2020), though immunolabeling suggests that Lawf1 is both cholinergic and glutamatergic (Yuan et al. 2021). Assuming that Lawf1 excites L3, and Dm20 inhibits Lawf1, the 3-cycle is net inhibitory. Three cells in such a 3-cycle are shown in Fig. 5e.

Lawf1 is a columnar neuron, by the traditional definition (Fischbach and Dittrich 1989), because its axon is oriented parallel to the main axis of the columns. However, it has two unusual properties. First, Lawf1 dendritic and axonal arbors cover many columns. The other columnar partners of Dm cells are mostly confined to a single column. This difference is related to the fact that Lawf1 is much less numerous than the other columnar types (Fig. 1e). Second, Lawf1 is a centrifugal neuron, projecting from the medulla to the lamina. In summary, Dm20 likely mediates lateral inhibition between L3 neurons, but the lateral inhibition involves feedback to the lamina.

Dm-Dm disinhibition

Up to now, only connectivity of Dm types with columnar types has been discussed. Dm-Dm connectivity is shown in Fig. S5a. Pixel values indicate the “importance” of the connection, defined as the maximum of input and output fraction (Methods). Some important inter-type Dm connections exist (Fig. S5a), though connectivity is rather sparse. Thresholding importance at 5%, there are only 12 out of 156 possible connections between distinct Dm types. Dm13 and Dm17 have no connections that survive thresholding. The thresholded graph of important connections is acyclic/feedforward (Fig. S5b), the only exception being that Dm10 and Dm20 mutually inhibit each other.

It is not obvious how Dm-Dm connectivity fits into the conceptual framework of normalization. One interesting case is Dm19 inhibition of Dm6 (Fig. S5a). Since both types share L2 as source (Figs. 3c, e), Dm19 can be interpreted as normalizing L2 input to Dm6. In other words, this could be a case of normalizing a normalizer.

All connections above the diagonal in Fig. S5a are from larger to smaller types. Some connections also exist below the diagonal, from smaller to larger types. Intra-type connectivity is weak, with input and output fractions less than 1% for 10 out of 13 types (diagonal pixels, Fig. S5a).

Generalization to other interneuron families

We have seen that many Dm interneurons are likely to act as normalizers. Does this finding generalize to interneuron families other than Dm? Space does not permit a full exploration of this question, but it is straightforward to identify candidate normalizers by searching for strong recurrent loops like Fig. 2b. It turns out that such loops exist in all other interneuron families (Pm, Sm, Li, and LPi) of the optic lobe that were described by (Matsliah et al. 2023).

Candidate Pm normalizers (Fig. 6a) interact with normalizees in ON (Mi1) and OFF (Tm1, Tm2) channels. Notably Pm08 interacts with both Mi1 and Tm1, so it could potentially mediate ON-OFF opponency as well as normalize.

Candidate Sm normalizers (Fig. 6b) interact with targets (Dm2, Dm8a, Dm8b, Tm20, DmDRA1, Mi15) of inner photoreceptor synapses (Kind et al. 2021), as well as with two Sm types. In several cases, the

normalizer is excitatory while the normalizee is inhibitory. This motif can be viewed as lateral inhibition, since the recurrent loop is net inhibitory.

Candidate Li normalizers (Fig. 6c) interact with types related to object (T2, T2a, T3, Tm21, Y3), motion (T5a), and color (Tm5f) vision, as well as types in ON (Tm3) and OFF (Tm2, Tm4) channels.

Candidate LPi normalizers (Fig. 6d) interact with T4/T5 types. These motifs suggest that normalization can be a function of LPi types. Previously, the only proposed function of LPi types was motion opponency (Mauss et al. 2015; Ammer et al. 2023).

Some columnar types are connected to multiple candidate normalizers in the same interneuron family. As in Fig. 4, this could implement normalization at multiple length scales. Furthermore, a single columnar type can be normalized by interneurons distributed across multiple neuropils (Figs. 3, 6). For example, Tm2 is likely normalized by Dm, Pm, and Li types, a rich and complex array of overlapping mechanisms.

All told, Fig. 6 contains 32 interneuron types that are candidate normalizers. The length of the list depends on a threshold, which was set at 10% of the input and output fraction of the normalizer (Fig. 6 caption). The list would become longer if the threshold were lowered, or if convergent feedforward inhibition like Fig. 2a were included. The list could shorten if types turn out to be electrically compartmentalized, as suspected for Dm4. Therefore Fig. 6 is not intended to be exhaustive; it is merely meant to support the overall claim that a significant fraction of interneuron types throughout the optic lobe are likely to be normalizers.

Discussion

Lateral inhibition was classically assumed to depend only on spatial separation (Hartline and Ratliff 1958; Ratliff 1965; Sperling 1970). In the fly distal medulla, lateral inhibition turns out to exhibit a high degree of cell type specificity. According to Rule 1, Dm type mediates inhibition from a specific (usually one) source type at one location to specific target types at neighboring locations (Fig. 3). According to Rule 2, the target type is either the same as the source type (recurrent inhibition), or the target type is directly downstream from the source type (convergent feedforward inhibition). A Dm interneuron type can mediate either recurrent or feedforward inhibition, or a combination of both (Fig. 3).

Lateral inhibition was classically assumed to have one characteristic length scale, but multiple scales turn out to exist in the fly distal medulla. For example, the notion of “lateral inhibition between L2 cells,” even if specific to cell types, is still coarser than reality. Recurrent inhibition of L2 is mediated by three Dm types of different sizes (Fig. 4a). Each Dm type mediates inhibition over spatial separations that are limited by its arbor size.

Analogy with convolutional nets

The motifs of Fig. 3 are analogous to normalization layers in a convolutional net, but there are some important differences. The first is that Dm interneurons average the activity of a source type over a spatial neighborhood, so that a Dm type can be approximated as computing a convolution of source activity. Batch (Ioffe and Szegedy 2015) and instance (Ulyanov, Vedaldi, and Lempitsky 2016) normalization, in contrast, sum over an entire feature map (and a batch of examples in the former case as well) when computing global statistics (mean and variance). Summing over an entire feature map

looks completely different from convolution, at first glance, but actually global summation can be regarded as convolution by a kernel of the same size as the feature map. One can imagine reducing the size of this kernel, and thereby modifying batch or instance normalization to compute neighborhood statistics like Dm cells. Indeed, the biologically inspired forerunners of convolutional nets did normalize over neighborhoods rather than the whole image (Fukushima 1975, 1980).

Since batch (Ioffe and Szegedy 2015) and instance (Ulyanov, Vedaldi, and Lempitsky 2016) normalization work on individual feature maps, they are analogous to the Dm types that have a single dominant source input. Group (Wu and He 2018) and layer (Ba, Kiros, and Hinton 2016) normalization are quite different because they mix feature maps together. They are more analogous to Dm13 and Dm20, which mediate lateral inhibition but have more than one dominant source input (Fig. S1).

The recurrent inhibition motif (Figs. 2b, 3a) is the closest analog to normalizing a feature map in a convolutional net, because it normalizes a single cell type, and its effects are seen by all downstream cell types. Convergent feedforward inhibition (Fig. 2a, 3b) is potentially more specific. In this motif, a target type compares source activity relative to some local average of source activity. A different target of the same source could compute a different normalization of the source or even none at all. In a convolutional net, this would be a different kind of normalization associated not with a feature map but with a connection between feature maps.

Dm4 is predicted to be electrically compartmentalized

The attenuation of electrical signals along a passive neurite has been studied theoretically (Koch 2004). Such theories cannot generate accurate quantitative predictions about Dm cells because of uncertainty in parameters like the specific transmembrane resistance and the axoplasmic resistivity. These parameters may vary by a factor of two or more, even after constraining them with electrophysiological measurements (Gouwens and Wilson 2009).

Since it is difficult to predict the absolute degree of Dm cell compartmentalization, I have resorted to comparing Dm cell morphologies to predict the relative degree of compartmentalization. Unique among the Dm baker's dozen, the Dm4 cell has long looping neurites that make the path length between neighboring columns many times greater than the Euclidean distance. This outlier morphology suggests that Dm4 is much more electrically compartmentalized than other Dm types.

Similar long looping neurites have been observed in the CT1 cell, an example of a large amacrine cell in the optic lobe that exhibits "extreme compartmentalization."(Meier and Borst 2019) Because of these long neurite paths, CT1 terminals in the medulla are electrically isolated from each other to a first approximation, even those in neighboring columns. Calcium imaging experiments show that CT1 terminals act as independent functional units; each CT1 terminal responds to visual input in its own column (Meier and Borst 2019).

Cable theory predicts that attenuation depends more strongly on path length, and rather weakly on diameter and membrane properties (due to a square root in the formula for length constant) (Koch 2004). Therefore varying path length across Dm types should be an effective way of varying degree of compartmentalization.

Translation invariance

Convolutional nets are often said to be translation invariant, meaning that the neurons of a feature map perform the same computation at different locations in the visual field. However, it is well-known that translation invariance is broken if feature maps are downsampled. For example, some convolutional nets contain an operation known as strided convolution that subsamples the result of regular convolution (Dumoulin and Visin 2016). A stride of two means that every other pixel is retained, resulting in a smaller output feature map with a spatial resolution that is halved relative to the input feature map.

The stride of a convolution controls coverage. If the stride is equal to the kernel size, then the coverage factor is exactly one, similar to Dm4. If the stride is equal to one, the coverage factor is equal to the kernel size, and translation invariance holds. A stride greater than one means that translation invariance is broken, with more severe effects for larger stride.

Each Dm type is analogous to a feature map that is constructed via strided convolution. Most Dm types have higher coverage factors (Fig. 5c), so that breaking of translation invariance is mild. But Dm4 has a low coverage factor (Fig. 5a, S4), so translation invariance of its target Mi4 would be broken quite severely (Fig. 5b, c) if lateral inhibition were the function of Dm4. If, on the other hand, Dm4 terminals function as independent units, then perfect tiling of Dm4 would not break translation invariance. Dm4 would be much like an assembly of unicolumnar cells, implementing convolution with stride one.

The higher coverage factors of most Dm types (Fig. 5a) are also typical of most ganglion cell types in the mouse retina (Bae et al. 2018). A higher coverage factor means that translation invariance is approximately preserved, which may be a design principle in visual systems.

Dm4 may still act as a functional unit for top-down inputs

If the columns of a Dm4 arbor are electrically independent, could there be any advantage of uniting them into a single cell? One intriguing possibility is that the columns are not electrically independent when the cell receives synapses from centrifugal pathways originating in the central brain. Such synapses tend to lie on thicker dendrites that are closer to the cell body than the terminals (data not shown). In a passive dendrite, electrical attenuation along the thick-to-thin direction is less than along the thin-to-thick direction (Rall and Rinzel 1973; Roth and Häusser 2001). Therefore, synaptic input on the thicker portions of dendrites may propagate to the Dm4 terminals with relatively little attenuation, so that top-down input influences all terminals as one functional unit.

Is lateral inhibition subtractive or divisive?

Network models of normalization by lateral inhibition have been classified as subtractive vs. divisive (Furman 1965). In the subtractive model, the currents from excitatory and inhibitory synapses are approximated as combining linearly as in Kirchhoff's First Law. The divisive or shunting model arises from the ratio of current and conductance in Ohm's Law, because synapses affect voltage by causing conductance changes. The center-surround receptive field in the *Limulus* eye was modeled with subtractive inhibition by (Hartline and Ratliff 1958). The model of human spatial vision due to (Sperling 1970) is divisive.

Normalization in convolutional nets combines subtractive and divisive steps (Ioffe and Szegedy 2015; Ulyanov, Vedaldi, and Lempitsky 2016). The biologically inspired forerunners of convolutional nets made use of divisive normalization (Fukushima 1975, 1980), as did influential models of mammalian neurophysiology (Carandini and Heeger 1994, 2011).

Visual responses of fly medulla neurons are divisively normalized by background image contrast (Drews et al. 2020). There is also direct electrophysiological evidence that inhibition is divisive in fly T4 visual neurons (Groschner et al. 2022). Physiological studies of inhibition from Dm interneurons are clearly warranted.

Whether inhibition is divisive or subtractive has been controversial for mammalian neurons (Holt and Koch 1997; Silver 2010). Many theoretical studies have attempted to relate this question to synapse location, motivated by the fact that mammalian interneuron types often prefer to synapse on particular compartments of the postsynaptic neuron. In the fly medulla, Dm neurons synapse onto the axons of L cells and onto the dendrites of other columnar types, and these facts might be relevant for understanding the effects of inhibition.

Luminance vs. contrast encoding

Normalization by Dm-mediated lateral inhibition can be probed by observing how visual responses of a columnar neuron are affected by neurons in other columns (Drews et al. 2020; Ketkar et al. 2020, 2022, 2023). This is a complex subject, and here I venture only some simple remarks.

Since L2 receives three types of Dm-mediated lateral inhibition (Fig. 4a), its output is expected to be highly normalized. In contrast, L1 has little or no connectivity with Dm types (Fig. S1, S2), so its output is expected to be relatively unnormalized. L3 receives lateral inhibition in the lamina (Fig. 3g), and its lateral inhibition from Dm12 (Fig. 3a) is rather short-range (Fig. 4g). Therefore L3 is expected to be less normalized than L2 but more than L1.

L2 is insensitive to luminance (Ketkar et al. 2020), consistent with L2 being highly normalized. L1 (Ketkar et al. 2022) encodes luminance rather linearly, consistent with the idea of no normalizing mechanisms in the medulla. L3 (Ketkar et al. 2020) encodes luminance but rather nonlinearly, consistent with some degree of normalization.

These arguments are merely hand-waving preludes to a future computational model of neural activity dynamics in optic lobe circuitry. They serve to illustrate the general idea that the connectivity of Dm types should have strong implications for visual response properties of columnar neurons.

Functional annotation of interneuron types

In addition to analyzing Dm circuits in detail, I also searched for candidate normalizers in other interneuron families throughout the optic lobe (Fig. 6). Because the relevant motifs are so widespread, a significant fraction of optic lobe cell types seem likely to be normalizers. This shows the power of connectomics to provide clues to the functions of cell types before any visual physiology experiments have been performed.

The present work can be regarded as a contribution to the emerging field of connectome annotation (Scheffer et al. 2020; Dorkenwald et al. 2023; Schlegel et al. 2023). Connectomic analysis provides not

only cell types but also rules of their connectivity. These rules can be used to deduce functional annotations for the cell types, and the present work shows by example how to arrive at such deductions. The traditional method of investigating the function of a visual neuron has been to record its responses to visual stimuli, but such experiments are time-consuming and laborious. I expect that such physiological studies will be accelerated by functional annotations of the connectome like the ones provided here.

Methods

All analyses are based on v783 of the FlyWire neuronal wiring diagram of a female adult fly brain (Zheng et al. 2018; Dorkenwald et al. 2023; Schlegel et al. 2023), and cell types assigned in the right optic lobe (Matsliah et al. 2023). Cells were assigned to a hexagonal lattice of columns as described in a companion paper (Seung 2023).

Input and output fractions

Let $N_{\mathrm{A}\to\mathrm{B}}$ be the number of synapses from type A neurons to type B neurons. This number is normalized to yield the input fraction,

$$f^{\mathrm{in}}_{\mathrm{A}\to\mathrm{B}} = \frac{N_{\mathrm{A}\to\mathrm{B}}}{N_{\mathrm{A}\to}}$$

Here $N_{\mathrm{A}\to}$ is the number of synapses from A neurons to any neuron, and $N_{\to\mathrm{B}}$ is the number of synapses from any neuron to B neurons. The output fraction is analogously defined as

$$f^{\mathrm{out}}_{\mathrm{A}\to\mathrm{B}} = \frac{N_{\mathrm{A}\to\mathrm{B}}}{N_{\to\mathrm{B}}}$$

Acknowledgements

I thank Ben Silverman for the Dm4 skeleton tracing of Fig. 5d, Ryan Willie for illustrations in Fig. 1, and the FlyWire team for the connectome and cell types on which this analysis is based. I am grateful to Michael Hausser for advice concerning dendritic biophysics.

Figures

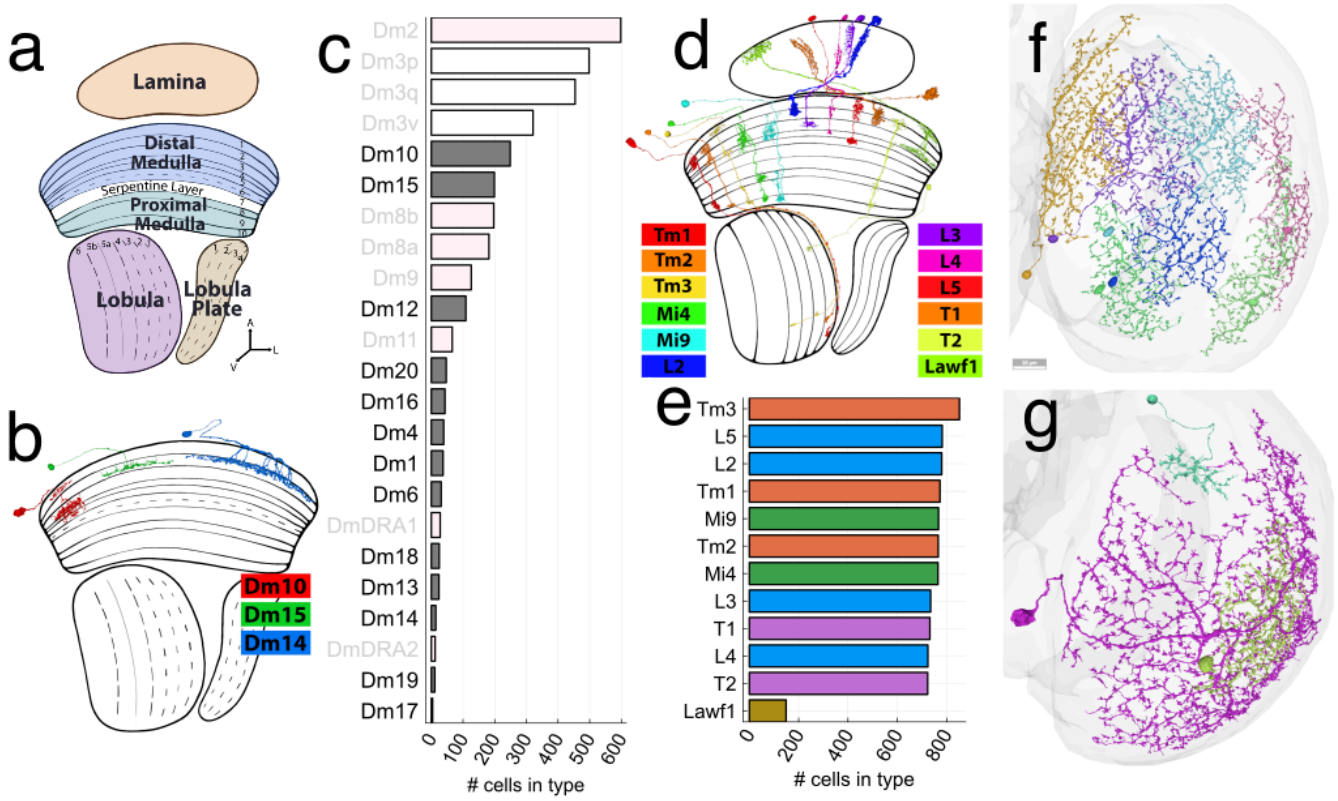

Figure 1. Dm interneurons and columnar neurons

(a) Main neuropils of the *Drosophila* optic lobe. The distal medulla is defined as the first six of ten medulla layers.

(b) Examples of Dm interneuron types. Dm10 (red) is the smallest and stratifies in multiple medulla layers. Dm15 (green) and Dm14 (blue) have slightly different sizes and stratification depths, but very different synaptic partners (Figs. S1, S2).

(c) Twenty-three Dm types ordered by the number of cells contained in each type. Dm types in dark letters are the “baker’s dozen” for which no function is known. Other Dm types (gray letters) are related to color vision (pink) or form vision (white) (Seung 2023).

(d) Columnar synaptic partners of Dm baker’s dozen, left to right in medulla: Tm1, Tm2, Tm3, Mi4, Mi9, L2, L3, L4, L5, T1, T2, Lawf1.

(e) Number of reconstructed cells for each columnar type in (d), colored by neuropil family. Almost all types are roughly as numerous as columns. The exception is Lawf1, which has a wider arbor in (d).

(f) Seven Dm6 cells. All 31 Dm6 cells would cover the entire medulla with overlap.

(g) Dm15 (cyan), Dm6 (green), and Dm17 (purple) cells vary greatly in size. Cell ids 720575940609136476, 720575940634638562, 720575940636785380

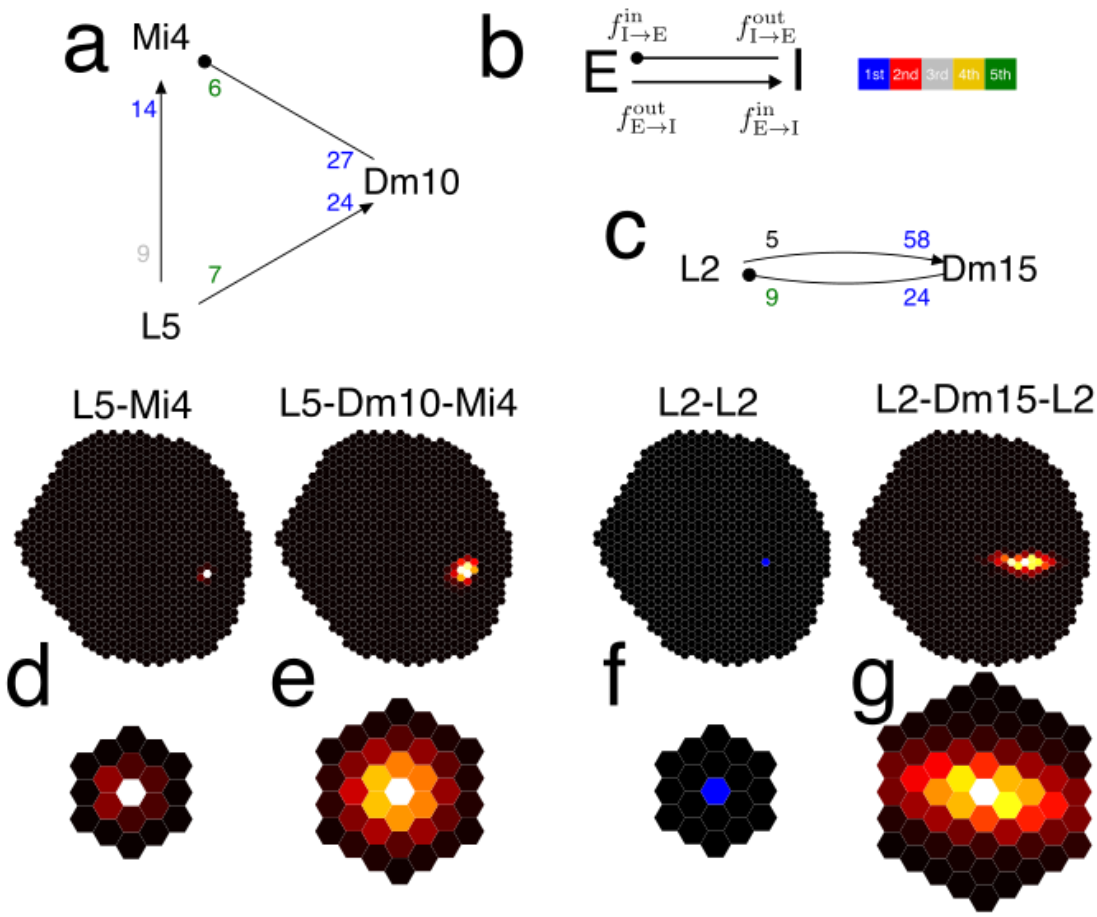

Figure 2. Dm-mediated lateral inhibition between columnar types

(a) Dm10 mediates convergent feedforward inhibition from L5 to Mi4.

(b) Key to arrowheads, numbers, and colors in motif diagrams. Arrowheads indicate excitatory (triangular) and inhibitory (circular) connections. The numbers at the head and tail of an arrow are the input and output fractions (×100) of that connection. The color of the number is an ordinal ranking of strength. The input rank of a connection is obtained by sorting all connections that share the same output type in order of decreasing input fraction. The output rank of a connection is defined by sorting all connections that share the same input type in order of decreasing output fraction.

(c) Dm15 mediates recurrent inhibition of L2.

(d) Example and average L5-Mi4 connectivity maps. Top, all L5 cells presynaptic to an example Mi4 cell. Hexel color indicates the number of synapses from the L5 cell at that location to the Mi4 cell. Bottom, aligned and averaged for all Mi4 cells

(e) Example and average L5-Dm10-Mi4 connectivity maps. Top, all L5 cells presynaptic to Dm10 cells presynaptic to the same Mi4 cell as in (d, top). Hexel color indicates the products of the L5-Dm10 and Dm10-Mi4 synapse numbers summed over all disynaptic paths from the L5 at that hexel to the example Mi4 cell. Bottom, aligned and averaged for all Mi4 cells.

(f) Example and average L2-L2 connectivity maps. These maps are everywhere zero or negligible; the blue hexel just guides the eye to the location of the target L2 cell.

(g) Example and average L2-Dm15-L2 connectivity maps.

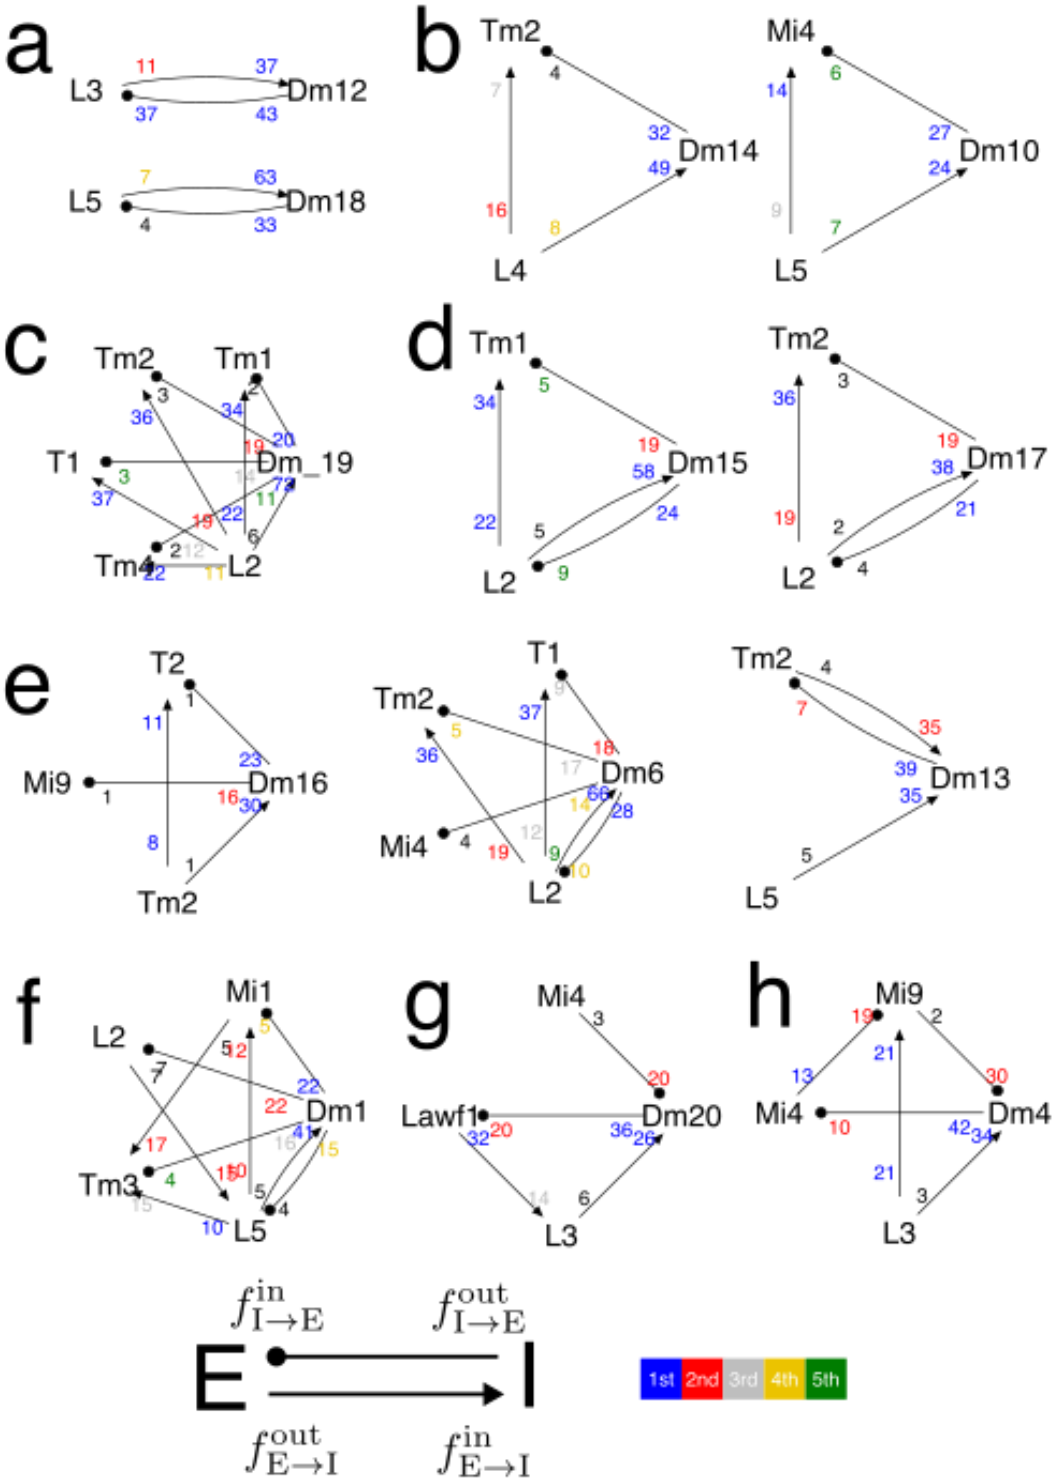

Figure 3. Lateral inhibition is normalizing

Dm types with their dominant inputs and outputs (defined in Fig. S1, S2). Only connections with importance (Methods) exceeding 10% are shown. Most motifs obey Rule 2, so lateral inhibition can be interpreted as normalizing.

(a) Dm12 and Dm18 mediate recurrent inhibition of L3 and L5, respectively.
(b) Dm14 and Dm10 mediate feedforward inhibition from L4 and L5, respectively.
(c) Dm19 mediates feedforward inhibition from L2 to multiple targets.
(d) Dm15 and Dm17 mediate a combination of feedforward and recurrent inhibition.
(e) Dm16, Dm6, and Dm13 mostly conform to Rule 2. The “dangling” connections to Mi4 and from L5 can be interpreted as mediating ON-OFF opponency.
(f) Dm1 mostly conforms to Rule 2, but contains a 3-cycle L5→Dm1→L2→L5, and an “extra” connection from Mi1 to Tm3.
(g) Dm20 mediates 3-cycle inhibition, which breaks Rule 2.
(h) Dm4 is another outlier that does not conform to Rule 2.

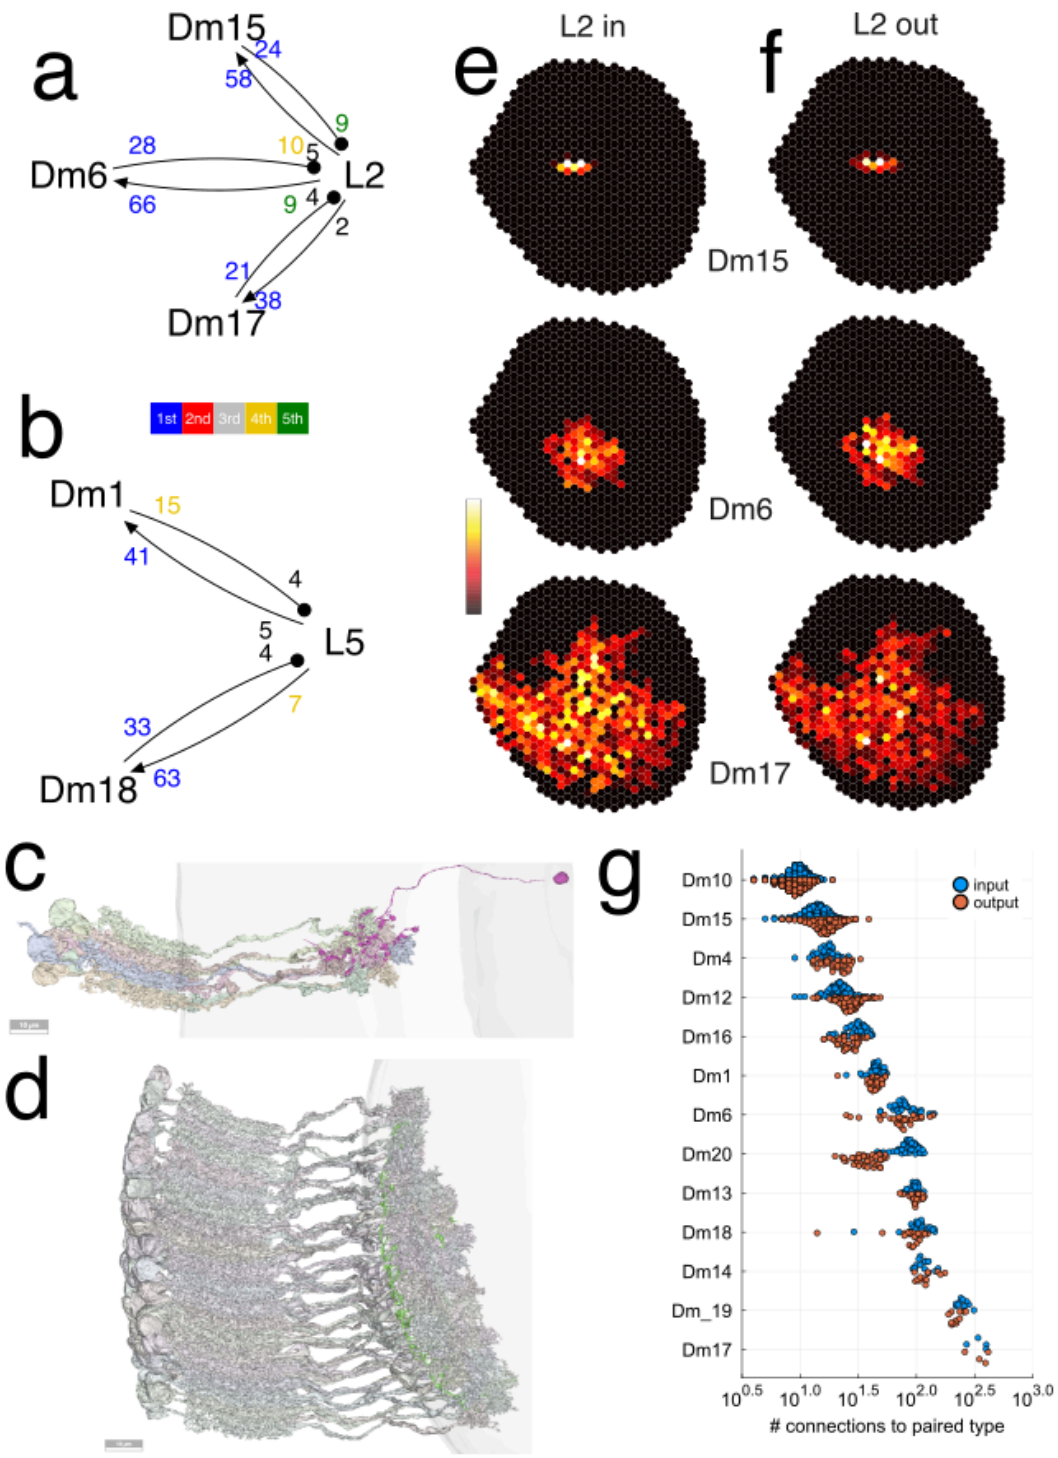

Figure 4. Lateral inhibition is multiscale

(a) Recurrent inhibition of L2 is mediated by three Dm types.
(b) Recurrent inhibition of L5 is mediated by two Dm types.
(c) Dm15 cell (red, 720575940644990062) and 9 presynaptic L2 cells (>4 synapses)
(d) Dm6 cell (green, 720575940609136476) and 74 presynaptic L2 cells (>4 synapses)

(e) L2-Dm connectivity maps for example Dm15, Dm6, and Dm17 cells.
(f) Dm-L2 connectivity maps for example Dm15, Dm6, and Dm17 cells. The inputs and outputs to a Dm cell have the same spatial organization, because they correspond to the columns covered by the Dm arbor. Colorbar maximum 12, 37, 15 synapses for L2-Dm and 16, 18, 17 synapses for Dm-L2. Cell IDs are 720575940609136476, 720575940634638562, 720575940644990062.
(g) For cells of each Dm type, number of connections with strongest input and output types. For each Dm type, the distributions are similar for input and output. Dm20 is an exception because its strongest output is Lawf1, which is multicolumnar.

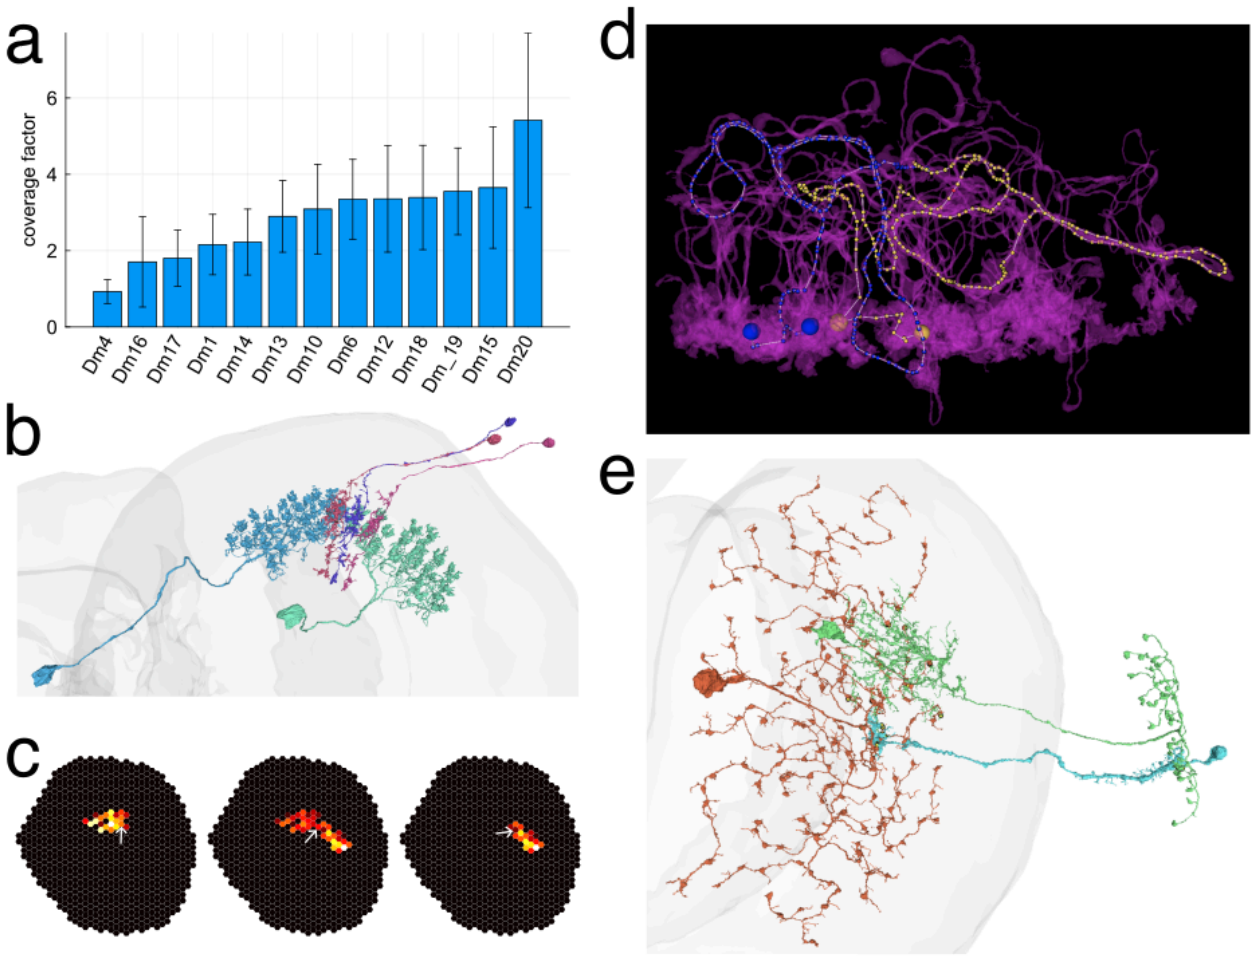

Figure 5. Dm4 and Dm20 are exceptions

(a) Coverage factors of Dm types. Dm4 is the only type with coverage factor close to one, meaning almost perfect tiling of the medulla (Fig. S4). Dm20 has an unusually large coverage factor.

(b) Mi4 receiving input from two Dm4 cells, flanked by two Mi4 receiving input from single Dm4 cell. Cell ids are 720575940613047794 720575940613604581,720575940643591368 720575940620113969 720575940614309416.

(c) Example L3-Dm4-Mi4 connectivity maps for the Mi4 cells of (b). The white arrow in each map points to the location of the Mi4 cell. The dramatic differences between these maps constitute a severe breaking of translation invariance.

(d) Dm4 may not be electrotonically compact, because of long path lengths between columns. Two neighboring columns (large blue and large yellow balls indicate two pairs of terminals in the two columns) are separated by a Euclidean distance of roughly 7 μm, but the neurite path length (blue and yellow skeletons) between the columns exceeds 200 μm. The blue path from the blue terminals and the yellow path from the yellow terminals meet at a branch point. Each path exceeds 100 μm. Cell id is 720575940627092016 in v630.

(e) Example of L3 (cyan) to Dm20 (orange) to Lawf1 (green) mediating 3-cycle recurrent inhibition. Lawf1 is a centrifugal cell that projects from the medulla to the lamina (left to right), where it synapses onto L3. Cell ids are 720575940607554434 720575940608960009 720575940633352100.

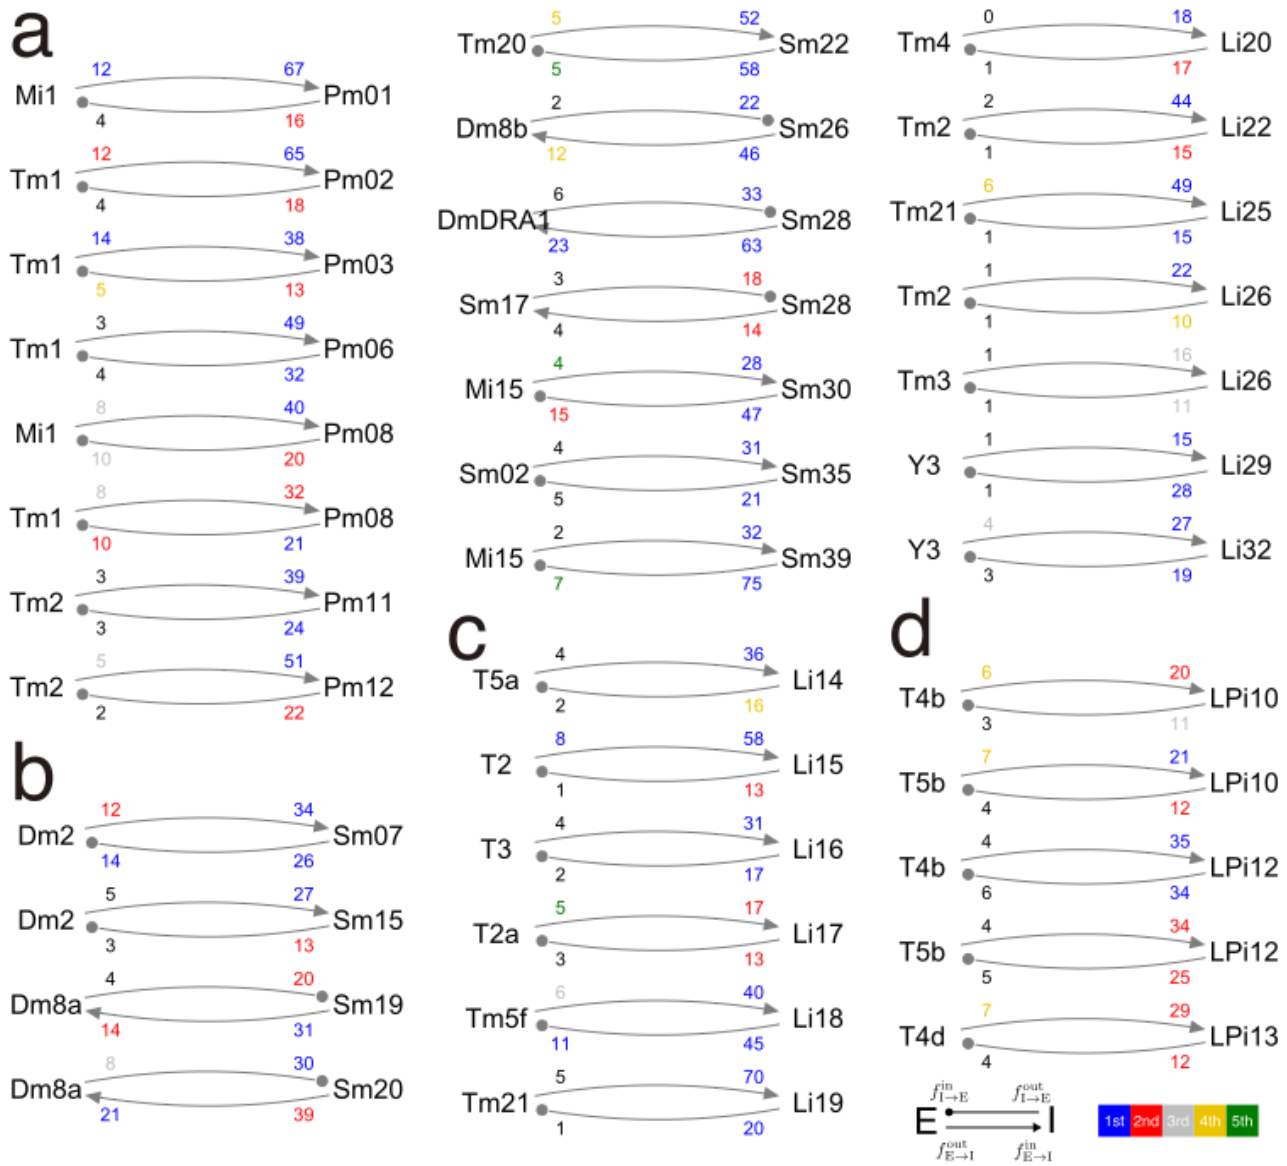

Figure 6. Candidate normalizee-normalizer cell type pairs

In all motifs, left type is normalizee and right type is normalizer. Normalizee is source of more than 10% of input synapses and target of more than 10% of output synapses of normalizer (percentages are small numbers near the normalizer). In almost all cases, normalizee is the strongest or second strongest input and output of normalizer (almost all percentages are blue or red). Normalizers include interneuron types in the Pm (a), Sm (b), Li (c), and LPi (d) families. Some normalizees are connected to multiple normalizers, similar to Fig. 4, and some normalizers are connected to multiple normalizees.

Supplementary Figures

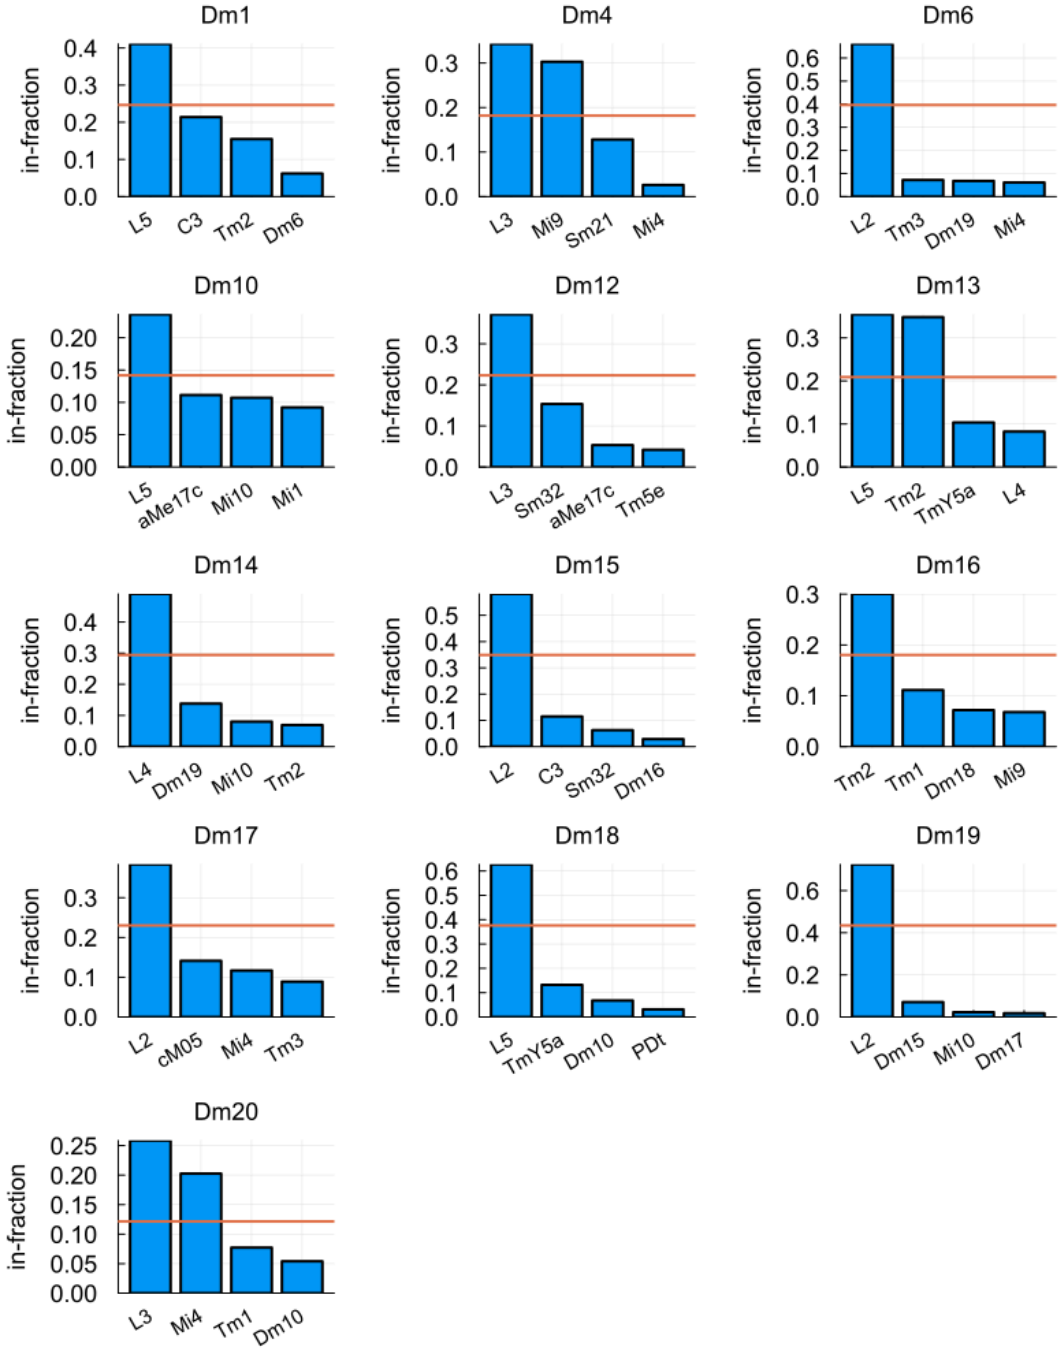

Figure S1. Each Dm type is dominated by one or two input types

Top four input types for each Dm type. The bars poking above the horizontal red line are said to be dominant inputs. The red line is placed at 60% of the height of the shortest bar above the red line. Most Dm types have a single dominant input. Dm13, Dm20, and Dm4 have two dominant inputs.

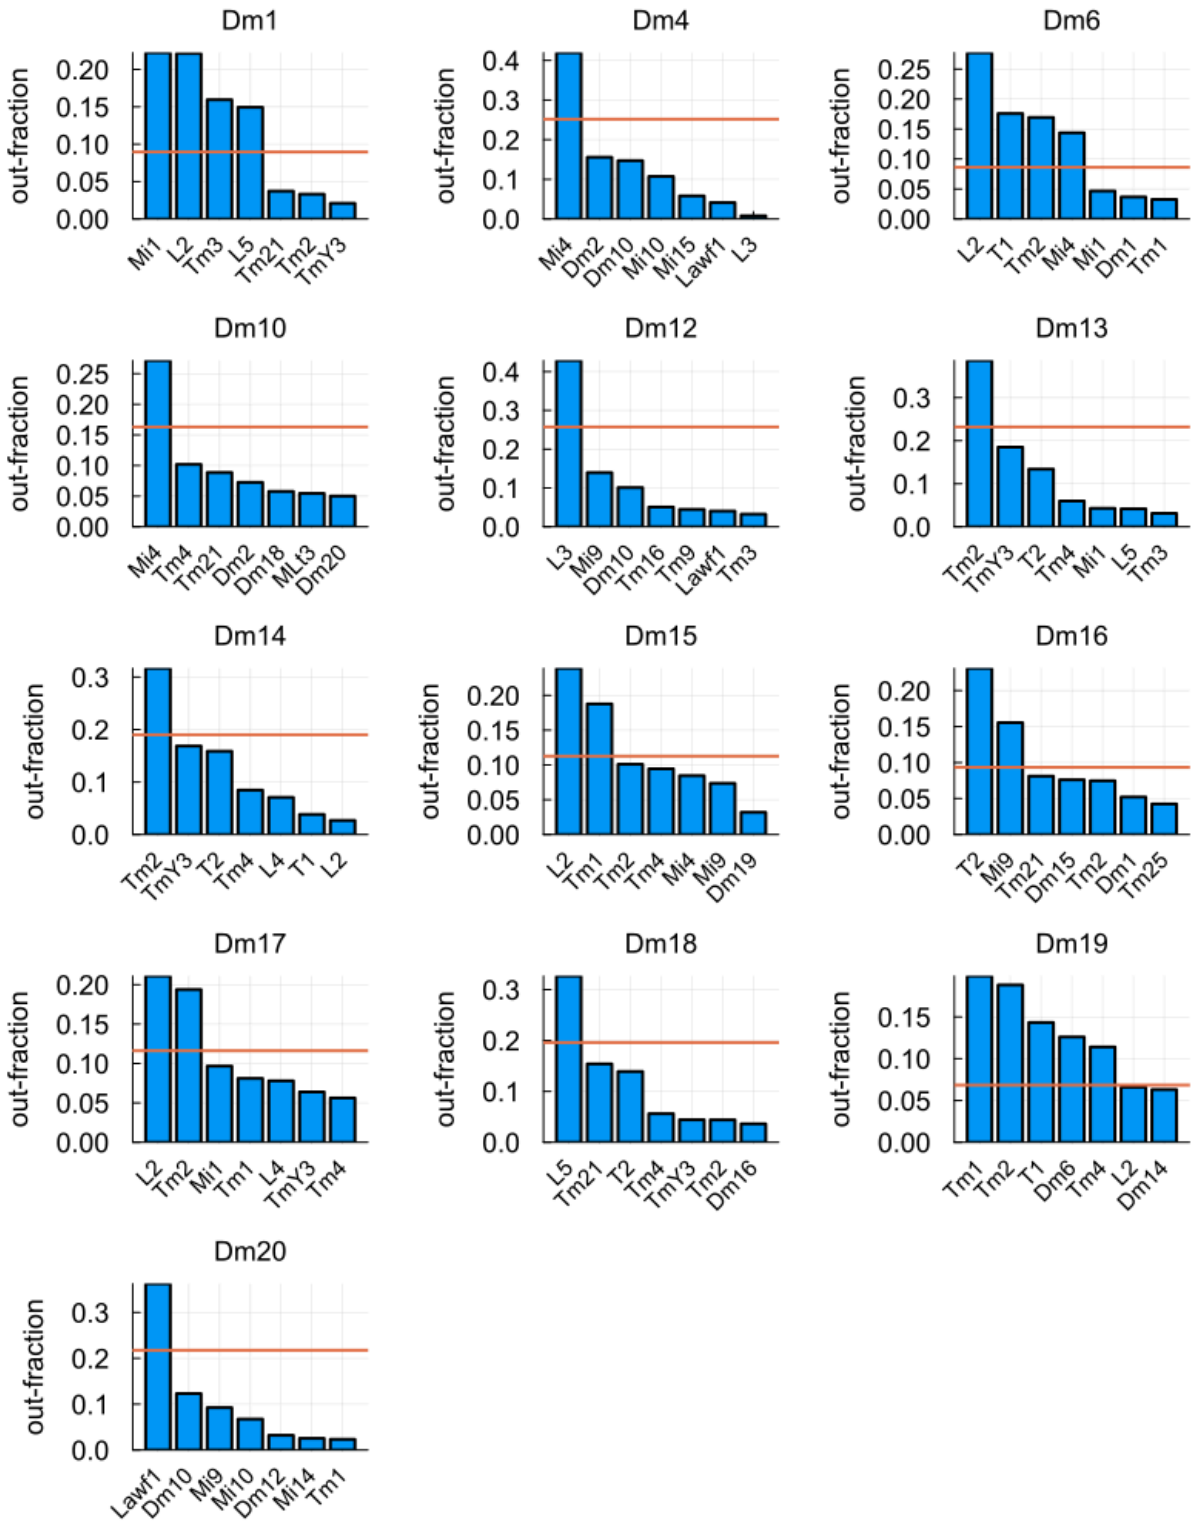

Figure S2. Dominant outputs for Dm interneurons

Top six output types for each Dm type. The bars poking above the horizontal red line are said to be dominant outputs. The red line is placed at 60% of the height of the shortest bar above the red line. Roughly half the Dm types have a single dominant output, but some have four or even five.

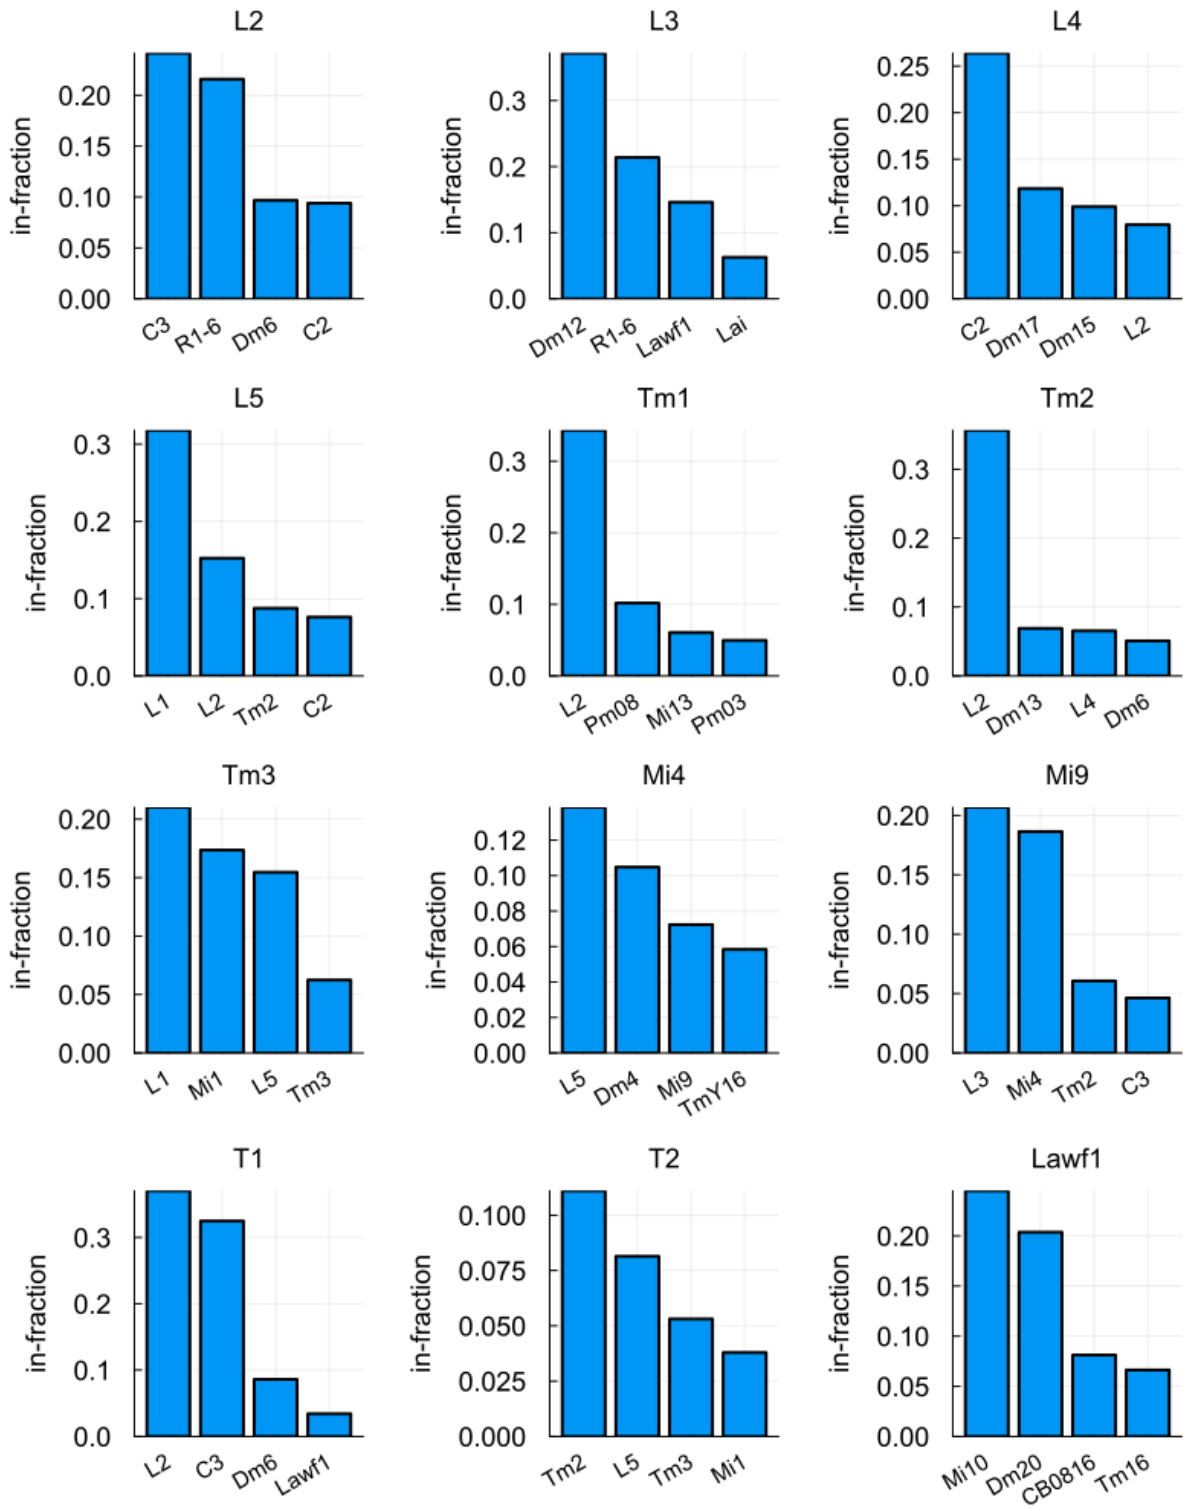

Figure S3. Inputs to columnar types

Top four inputs to the columnar types that are the main partners (dominant inputs or outputs) of Dm types (same columnar types as Fig. 1e). R1-6 input fractions are known to be inaccurate due to underdetection of photoreceptor synapses, and a fix is in progress.

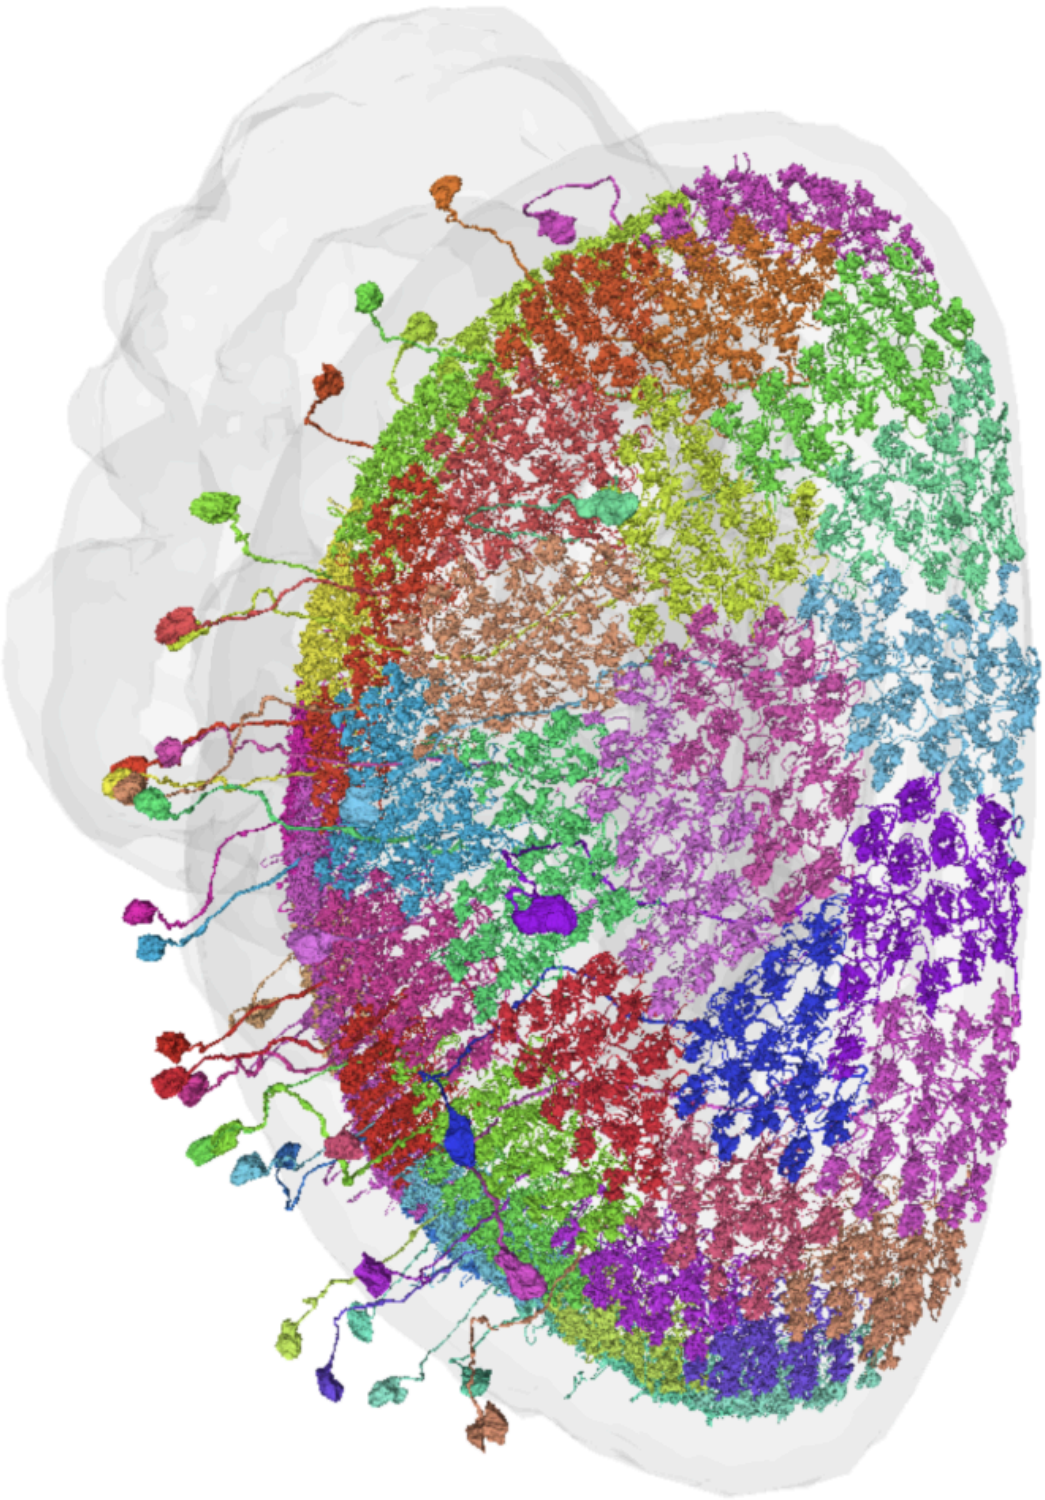

Figure S4. Perfect tiling of the visual field by Dm4 cells

Dm4 is the only type in the Dm4 baker's dozen that perfectly tiles the medulla, leading to a coverage factor that is approximately one (Fig. 5a). Other Dm types contain cells that overlap significantly with each other, with coverage factors that are typically 2 or 3 (Fig. 5a).

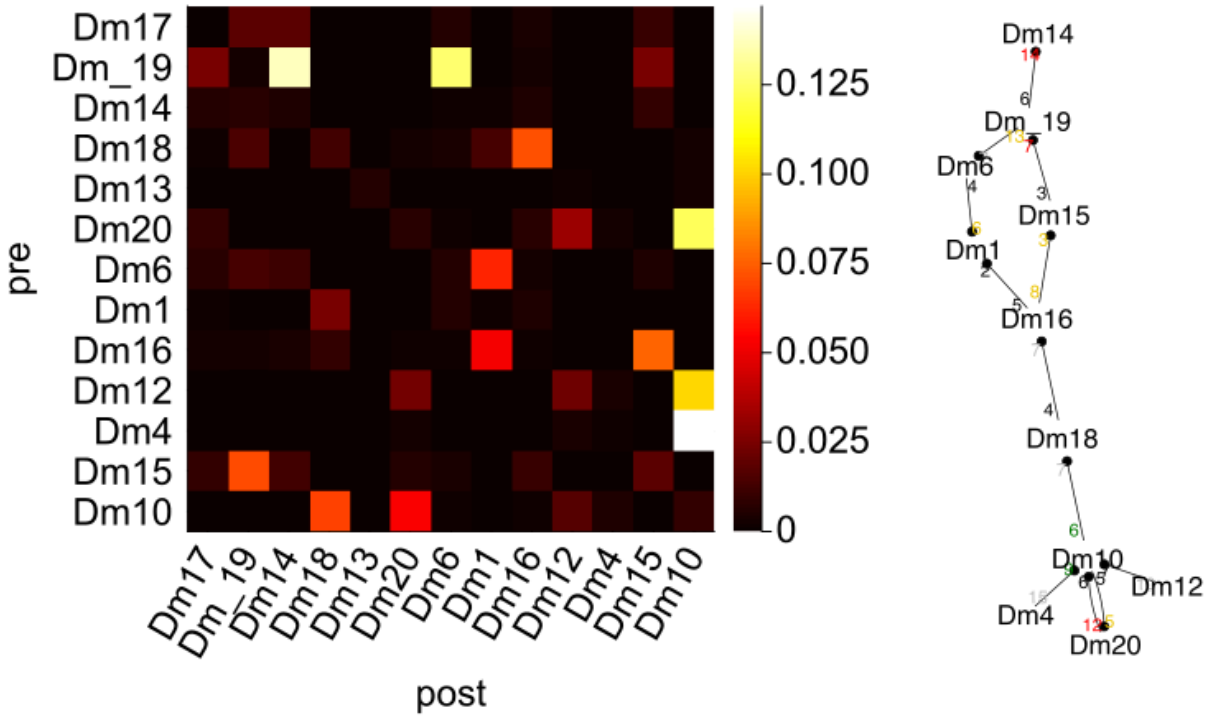

Figure S5. Dm-Dm disinhibition

(a) Connectivity between Dm types, ordered with decreasing arbor size. Each pixel value denotes "importance," defined as the maximum of the input and output fractions for that connection. Pixels above the diagonal correspond to connections from larger to smaller cells, and vice versa for pixels below the diagonal.

(b) Graph showing connections with importance greater than 5%.

References

Ammer, Georg, Etienne Serbe-Kamp, Alex S. Mauss, Florian G. Richter, Sandra Fendl, and Alexander Borst. 2023. "Multilevel Visual Motion Opponency in Drosophila." *Nature Neuroscience* 26 (11): 1894–1905.

Arenz, Alexander, Michael S. Drews, Florian G. Richter, Georg Ammer, and Alexander Borst. 2017. "The Temporal Tuning of the Drosophila Motion Detectors Is Determined by the Dynamics of Their Input Elements." *Current Biology: CB* 27 (7): 929–44.

Bae, J. Alexander, Shang Mu, Jinseop S. Kim, Nicholas L. Turner, Ignacio Tartavull, Nico Kemnitz, Chris S. Jordan, et al. 2018. "Digital Museum of Retinal Ganglion Cells with Dense Anatomy and Physiology." *Cell* 173 (5): 1293–1306.e19.

Ba, Jimmy Lei, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. "Layer Normalization." *arXiv [stat.ML]*. arXiv. http://arxiv.org/abs/1607.06450.

Barlow, R. B., Jr. 1969. "Inhibitory Fields in the Limulus Lateral Eye." *The Journal of General Physiology* 54 (3): 383–96.

Bjorck, Nils, Carla P. Gomes, Bart Selman, and Kilian Q. Weinberger. 2018. "Understanding Batch Normalization." *Advances in Neural Information Processing Systems* 31.

Carandini, M., and David J. Heeger. 2011. "Normalization as a Canonical Neural Computation." *Nature Reviews. Neuroscience* 13 (1): 51–62.

Carandini, M., and D. J. Heeger. 1994. “Summation and Division by Neurons in Primate Visual Cortex.” *Science* 264 (5163): 1333–36.
Davis, Fred P., Aljoscha Nern, Serge Picard, Michael B. Reiser, Gerald M. Rubin, Sean R. Eddy, and Gilbert L. Henry. 2020. “A Genetic, Genomic, and Computational Resource for Exploring Neural Circuit Function.” *eLife* 9 (January). https://doi.org/10.7554/eLife.50901.
Dorkenwald, Sven, Arie Matsliah, Amy R. Sterling, Philipp Schlegel, Szi-Chieh Yu, Claire E. McKellar, Albert Lin, et al. 2023. “Neuronal Wiring Diagram of an Adult Brain.” *bioRxiv : The Preprint Server for Biology*, July. https://doi.org/10.1101/2023.06.27.546656.
Drews, Michael S., Aljoscha Leonhardt, Nadezhda Pirogova, Florian G. Richter, Anna Schuetzenberger, Lukas Braun, Etienne Serbe, and Alexander Borst. 2020. “Dynamic Signal Compression for Robust Motion Vision in Flies.” *Current Biology: CB* 30 (2): 209–21.e8.
Dumoulin, Vincent, and Francesco Visin. 2016. “A Guide to Convolution Arithmetic for Deep Learning.” *arXiv [stat.ML]*. arXiv. http://arxiv.org/abs/1603.07285.
Eckstein, Nils, Alexander S. Bates, M. Du, V. Hartenstein, Gsxe Jefferis, and J. Funke. 2020. “Neurotransmitter Classification from Electron Microscopy Images at Synaptic Sites in Drosophila.” *BioRxiv*. https://www.biorxiv.org/content/10.1101/2020.06.12.148775v2.abstract.
Fahrenbach, W. H. 1985. “Anatomical Circuitry of Lateral Inhibition in the Eye of the Horseshoe Crab, Limulus Polyphemus.” *Proceedings of the Royal Society of London. Series B, Containing Papers of a Biological Character. Royal Society* 225 (1239): 219–49.
Fischbach, K-F, and A. P. M. Dittrich. 1989. “The Optic Lobe of Drosophila Melanogaster. I. A Golgi Analysis of Wild-Type Structure.” *Cell and Tissue Research* 258 (3). https://doi.org/10.1007/bf00218858.
Freifeld, Limor, Damon A. Clark, Mark J. Schnitzer, Mark A. Horowitz, and Thomas R. Clandinin. 2013. “GABAergic Lateral Interactions Tune the Early Stages of Visual Processing in Drosophila.” *Neuron* 78 (6): 1075–89.
Fukushima, K. 1975. “Cognitron: A Self-Organizing Multilayered Neural Network.” *Biological Cybernetics* 20 (3-4): 121–36.
———. 1980. “Neocognitron: A Self Organizing Neural Network Model for a Mechanism of Pattern Recognition Unaffected by Shift in Position.” *Biological Cybernetics* 36 (4): 193–202.
Furman, G. G. 1965. “Comparison of Models for Subtractive and Shunting Lateral-Inhibition in Receptor-Neuron Fields.” *Kybernetik* 2 (6): 257–74.
Gouwens, Nathan W., and Rachel I. Wilson. 2009. “Signal Propagation in Drosophila Central Neurons.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 29 (19): 6239–49.
Groschner, Lukas N., Jonatan G. Malis, Birte Zuidinga, and Alexander Borst. 2022. “A Biophysical Account of Multiplication by a Single Neuron.” *Nature* 603 (7899): 119–23.
Hartline, H. K., and F. Ratliff. 1958. “Spatial Summation of Inhibitory Influences in the Eye of Limulus, and the Mutual Interaction of Receptor Units.” *The Journal of General Physiology* 41 (5): 1049–66.
Hartline, H. K., H. G. Wagner, and F. Ratliff. 1956. “Inhibition in the Eye of Limulus.” *The Journal of General Physiology* 39 (5): 651–73.
Heath, Sarah L., Matthias P. Christenson, Elie Oriol, Maia Saavedra-Weisenhaus, Jessica R. Kohn, and Rudy Behnia. 2020. “Circuit Mechanisms Underlying Chromatic Encoding in Drosophila Photoreceptors.” *Current Biology: CB* 30 (2): 264–75.e8.
Heeger, David J., and Klavdia O. Zemlianova. 2020. “A Recurrent Circuit Implements Normalization, Simulating the Dynamics of V1 Activity.” *Proceedings of the National Academy of Sciences of the United States of America* 117 (36): 22494–505.
He, Kaiming, X. Zhang, Shaoqing Ren, and Jian Sun. 2015. “Deep Residual Learning for Image Recognition.” *Proceedings / CVPR, IEEE Computer Society Conference on Computer Vision and Pattern Recognition. IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, December, 770–78.
Holt, G. R., and C. Koch. 1997. “Shunting Inhibition Does Not Have a Divisive Effect on Firing Rates.” *Neural Computation* 9 (5): 1001–13.
Ioffe, Sergey, and Christian Szegedy. 2015. “Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift.” In *Proceedings of the 32nd International Conference on*

Machine Learning, edited by Francis Bach and David Blei, 37:448–56. Proceedings of Machine Learning Research. Lille, France: PMLR.
Kepecs, Adam, and Gordon Fishell. 2014. “Interneuron Cell Types Are Fit to Function.” *Nature* 505 (7483): 318–26.
Ketkar, Madhura D., Burak Gür, Sebastian Molina-Obando, Maria Ioannidou, Carlotta Martelli, and Marion Silies. 2022. “First-Order Visual Interneurons Distribute Distinct Contrast and Luminance Information across ON and OFF Pathways to Achieve Stable Behavior.” *eLife* 11 (March). https://doi.org/10.7554/eLife.74937.
Ketkar, Madhura D., Shuai Shao, Julijana Gjorgjieva, and Marion Silies. 2023. “Multifaceted Luminance Gain Control beyond Photoreceptors in Drosophila.” *Current Biology: CB* 33 (13): 2632–45.e6.
Ketkar, Madhura D., Katja Sporar, Burak Gür, Giordano Ramos-Traslosheros, Marvin Seifert, and Marion Silies. 2020. “Luminance Information Is Required for the Accurate Estimation of Contrast in Rapidly Changing Visual Contexts.” *Current Biology: CB* 30 (4): 657–69.e4.
Kind, Emil, Kit D. Longden, Aljoscha Nern, Arthur Zhao, Gizem Sancer, Miriam A. Flynn, Connor W. Laughland, et al. 2021. “Synaptic Targets of Photoreceptors Specialized to Detect Color and Skylight Polarization in Drosophila.” *eLife* 10 (December). https://doi.org/10.7554/eLife.71858.
Koch, Christof. 2004. *Biophysics of Computation: Information Processing in Single Neurons*. Oxford University Press.
Lappalainen, Janne K., Fabian D. Tschopp, Sridhama Prakhya, Mason McGill, Aljoscha Nern, Kazunori Shinomiya, Shin-Ya Takemura, Eyal Gruntman, Jakob H. Macke, and Srinivas C. Turaga. 2023. “Connectome-Constrained Deep Mechanistic Networks Predict Neural Responses across the Fly Visual System at Single-Neuron Resolution.” *bioRxiv*.
Mach, E. 1868. “Über Die Physiologische Wirkung Räumlich Vertheilter Lichtreize. 4. Abhandlung.” *Sitzungsberichte Der Kaiserlichen Akademie Der Wissenschaften in Wien. Mathematisch-Naturwissenschaftliche Classe* 57 (2): 11–19.
Mangan, S., and U. Alon. 2003. “Structure and Function of the Feed-Forward Loop Network Motif.” *Proceedings of the National Academy of Sciences of the United States of America* 100 (21): 11980–85.
Matsliah, Arie, Szi-Chieh Yu, Krzysztof Kruk, Doug Bland, Austin Burke, Jay Gager, James Hebditch, et al. 2023. “Neuronal ‘Parts List’ and Wiring Diagram for a Visual System.” *bioRxiv : The Preprint Server for Biology*, October. https://doi.org/10.1101/2023.10.12.562119.
Mauss, Alex S., Katarina Pankova, Alexander Arenz, Aljoscha Nern, Gerald M. Rubin, and Alexander Borst. 2015. “Neural Circuit to Integrate Opposing Motions in the Visual Field.” *Cell* 162 (2): 351–62.
Meier, Matthias, and Alexander Borst. 2019. “Extreme Compartmentalization in a Drosophila Amacrine Cell.” *Current Biology: CB* 29 (9): 1545–50.e2.
Nern, A., B. D. Pfeiffer, and G. M. Rubin. 2015. “Optimized Tools for Multicolor Stochastic Labeling Reveal Diverse Stereotyped Cell Arrangements in the Fly Visual System.” *Proceedings of the*. https://www.pnas.org/content/112/22/E2967.short.
Pouille, Frédéric, Antonia Marin-Burgin, Hillel Adesnik, Bassam V. Atallah, and Massimo Scanziani. 2009. “Input Normalization by Global Feedforward Inhibition Expands Cortical Dynamic Range.” *Nature Neuroscience* 12 (12): 1577–85.
Rall, W., and J. Rinzel. 1973. “Branch Input Resistance and Steady Attenuation for Input to One Branch of a Dendritic Neuron Model.” *Biophysical Journal* 13 (7): 648–87.
Ratliff, Floyd. 1965. “Mach Bands: Quantitative Studies on Neural Networks.” *Retina. San Francisco, CA: Holden-Day*.
Roth, A., and M. Häusser. 2001. “Compartmental Models of Rat Cerebellar Purkinje Cells Based on Simultaneous Somatic and Dendritic Patch-Clamp Recordings.” *The Journal of Physiology* 535 (Pt 2): 445–72.
Santurkar, S., D. Tsipras, and A. Ilyas. 2018. “How Does Batch Normalization Help Optimization?” *Advances in Neural Information Processing Systems*. https://proceedings.neurips.cc/paper/2018/hash/905056c1ac1dad141560467e0a99e1cf-Abstract.html.
Sato, Tatsuo K., Bilal Haider, Michael Häusser, and Matteo Carandini. 2016. “An Excitatory Basis for

Divisive Normalization in Visual Cortex." *Nature Neuroscience* 19 (4): 568–70.
Scheffer, Louis K., C. Shan Xu, Michal Januszewski, Zhiyuan Lu, Shin-Ya Takemura, Kenneth J. Hayworth, Gary B. Huang, et al. 2020. "A Connectome and Analysis of the Adult Drosophila Central Brain." *eLife* 9 (September). https://doi.org/10.7554/eLife.57443.
Schlegel, Philipp, Yijie Yin, Alexander Shakeel Bates, Sven Dorkenwald, Katharina Eichler, Paul Brooks, Daniel S. Han, et al. 2023. "Whole-Brain Annotation and Multi-Connectome Cell Typing Quantifies Circuit Stereotypy in Drosophila." *bioRxiv : The Preprint Server for Biology*, July. https://doi.org/10.1101/2023.06.27.546055.
Schnaitmann, Christopher, Manuel Pagni, and Dierk F. Reiff. 2020. "Color Vision in Insects: Insights from Drosophila." *Journal of Comparative Physiology. A, Neuroethology, Sensory, Neural, and Behavioral Physiology* 206 (2): 183–98.
Seung, H. Sebastian. 2023. "Insights into Vision from Interpretation of a Neuronal Wiring Diagram." *bioRxiv*. https://doi.org/10.1101/2023.11.15.567126.
Silver, R. Angus. 2010. "Neuronal Arithmetic." *Nature Reviews. Neuroscience* 11 (7): 474–89.
Sperling, George. 1970. "Model of Visual Adaptation and Contrast Detection." *Perception & Psychophysics* 8 (3): 143–57.
Sterling, P. 1983. "Microcircuitry of the Cat Retina." *Annual Review of Neuroscience* 6: 149–85.
Takemura, Shin-Ya, Arjun Bharioke, Zhiyuan Lu, Aljoscha Nern, Shiv Vitaladevuni, Patricia K. Rivlin, William T. Katz, et al. 2013. "A Visual Motion Detection Circuit Suggested by Drosophila Connectomics." *Nature* 500 (7461): 175–81.
Takemura, Shin-Ya, C. Shan Xu, Zhiyuan Lu, Patricia K. Rivlin, Toufiq Parag, Donald J. Olbris, Stephen Plaza, et al. 2015. "Synaptic Circuits and Their Variations within Different Columns in the Visual System of Drosophila." *Proceedings of the National Academy of Sciences of the United States of America* 112 (44): 13711–16.
Tasic, Bosiljka, Zizhen Yao, Lucas T. Graybuck, Kimberly A. Smith, Thuc Nghi Nguyen, Darren Bertagnolli, Jeff Goldy, et al. 2018. "Shared and Distinct Transcriptomic Cell Types across Neocortical Areas." *Nature* 563 (7729): 72–78.
Ulyanov, Dmitry, Andrea Vedaldi, and Victor Lempitsky. 2016. "Instance Normalization: The Missing Ingredient for Fast Stylization." *arXiv [cs.CV]*. arXiv. http://arxiv.org/abs/1607.08022.
Wu, Yuxin, and Kaiming He. 2018. "Group Normalization." *International Journal of Computer Vision* 128 (March): 742–55.
Xu, Shuwa, Qi Xiao, Filip Cosmanescu, Alina P. Sergeeva, Juyoun Yoo, Ying Lin, Phinikoula S. Katsamba, et al. 2018. "Interactions between the Ig-Superfamily Proteins DIP-α and Dpr6/10 Regulate Assembly of Neural Circuits." *Neuron* 100 (6): 1369–84.e6.
Yuan, Deliang, Xiaoxiao Ji, Shun Hao, Julia Yvonne Gestrich, Wenlan Duan, Xinwei Wang, Yuanhang Xiang, et al. 2021. "Lamina Feedback Neurons Regulate the Bandpass Property of the Flicker-Induced Orientation Response in Drosophila." *Journal of Neurochemistry* 156 (1): 59–75.
Zheng, Zhihao, J. Scott Lauritzen, Eric Perlman, Camenzind G. Robinson, Matthew Nichols, Daniel Milkie, Omar Torrens, et al. 2018. "A Complete Electron Microscopy Volume of the Brain of Adult Drosophila Melanogaster." *Cell* 174 (3): 730–43.e22.