
Connectome-Based Modelling Reveals Orientation Maps in the *Drosophila* Optic Lobe

Jia-Nuo Liew^{1,2,*}, Shenghan Lin^{3,*}, Bowen Chen^{2,4,6,7}, Wei Zhang^{1,2},
Xiaowei Zhu⁴, Wei Zhang^{2,4,6,7,†}, Xiaolin Hu^{1,2,8,†}

¹Department of Computer Science and Technology, BNRist, Tsinghua University, Beijing 100084, China

²IDG/McGovern Institute of Brain Research, Tsinghua University, Beijing 100084, China

³Zhili College, Tsinghua University, Beijing 100084, China

⁴School of Life Sciences, Tsinghua University, Beijing 100084, China

⁵Ant Group, China

⁶State Key Laboratory of Membrane Biology, Tsinghua University, Beijing 100084, China

⁷Tsinghua-Peking Center for Life Sciences, Tsinghua University, Beijing 100084, China

⁸Chinese Institute for Brain Research (CIBR), Beijing 100010, China

{liujn24, linsh24, cbw21, zhangw23}@mails.tsinghua.edu.cn

robert.zxw@antgroup.com

{wei_zhang, xlhu}@mail.tsinghua.edu.cn

Abstract

The ability to extract oriented edges from visual input is a core computation across animal vision systems. Orientation maps, long associated with the layered architecture of the mammalian visual cortex, systematically organise neurons by their preferred edge orientation. Despite lacking cortical structures, the *Drosophila melanogaster* brain contains feature-selective neurons and exhibits complex visual detection capacity, raising the question of whether map-like vision representations can emerge without cortical infrastructure. We integrate a complete fruit fly brain connectome with biologically grounded spiking neuron models to simulate neuroprocessing in the fly visual system. By driving the network with oriented stimuli and analysing downstream responses, we show that coherent orientation maps can emerge from purely connectome-constrained dynamics. These results suggest that species of independent origin could evolve similar visual structures.

1 Introduction

The ability to extract oriented edges from visual input is a core computation across animal vision systems [23]. Orientation selectivity is canonically attributed to the primary visual cortex (V1), where neurons respond selectively to specific edge orientations. Hubel and Wiesel [22, 23] first demonstrated this modular structure in cats. The findings were later substantiated by electrophysiology, optical imaging, and detailed anatomical mapping [3, 41]. These orientation-selective responses are embedded within columnar and laminar architecture and are thought to arise from a combination of spatially organised feedforward inputs and recurrent cortical dynamics [2, 12, 15, 19, 48].

Orientation selectivity has been directly observed in *Drosophila melanogaster* visual system. In particular, T4 and T5 neurons - traditionally known for direction selectivity - also exhibit robust orientation tuning, which sharpens motion detection [16]. These findings provide physiological

*Equal contribution.

†Corresponding author.

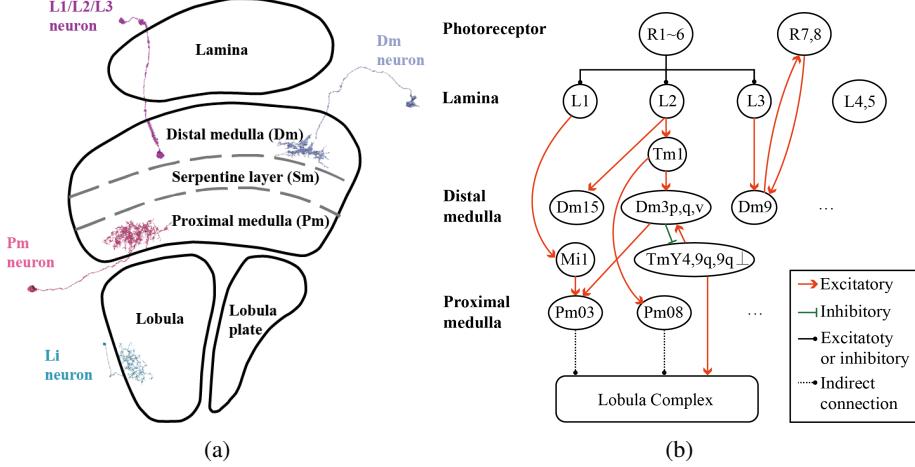


Figure 1: **Anatomical organisation and connectivity in the fly visual system.** (a) Schematic of the fly optic lobe showing the layered structure of the medulla, including the distal (Dm), serpentine (Sm), and proximal (Pm) sublayers. (b) Simplified connectivity diagram highlighting the visual pathway from photoreceptors (R1-6) through lamina neurons (L1-L3), to the orientation-selective neurons, to the lobula complex.

evidence that orientation selectivity exists at the neuron level. Building on this, recent studies have hypothesised that neurons in early visual stages, particularly in the medulla, could support orientation selectivity, based solely on synaptic connectivity patterns from structural data [45]. Though the question of whether such orientation selectivity is organised into coherent maps across the visual system remains unknown. Given the compact, non-layered architecture of the *Drosophila* brain and the absence of large-scale recurrent loops seen in vertebrate cortex, it is unclear whether this system can support the emergence of global tuning structures such as orientation maps [7, 10, 34].

To investigate this, we simulated the visual responses of early-stage neurons (L1-L3) using bar-like stimuli and propagated their activity through a downstream population in the optic lobe modelled with leaky integrate-and-fire (LIF) dynamics, constrained by known synaptic connections [13]. We then analysed the resulting neural population activity, revealing spatially organised orientation tuning reminiscent of cortical orientation maps. Figure 1a provides an overview of the anatomical structure of the *Drosophila* optic lobe, while Figure 1b illustrates the visual connectivity pathway examined in this study.

To summarise, our main contributions are as follows:

- We computationally demonstrated, for the first time, spatially coherent orientation maps in the medulla region of an invertebrate visual system.
- We identified topological singularities and inter-layer columnar alignment in orientation preference within the distal medulla (Dm) and proximal medulla (Pm) regions.

Together, these findings suggest that canonical orientation maps can arise from shared computational principles across species, even in the absence of cortical lamination.

2 Background and related works

Orientation maps in mammals. Orientation selectivity has been extensively studied in the visual cortices of many mammals. At the anatomical level, evidence from mammals indicates that orientation tuning can emerge in neural substrates with vast differences [25, 4, 38]. In rodents, maps often display salt-and-pepper rather than columnar architecture, as observed in the visual cortex, while the retina displays a continuous topographic map of orientation tuning [48]. In contrast, cats and wallabies display a structured pinwheel-like feature that has been documented in the primary visual cortex [24, 25, 37], despite differences in cortical evolutionary divergence across mammalian species [41]. These findings suggest that orientation maps can arise across diverse mammalian species with varying

cortical architectures, raising the possibility that such maps are not an exclusive feature of the neocortex but reflect convergent computational motifs.

Orientation tuning in non-mammals. Studies across non-mammalian species suggest that structured visual coding arises under minimal anatomical constraints. In pigeons, sharply tuned orientation-selective neurons emerge through feedforward circuits alone, without a laminated cortex [32]. In turtles, despite the absence of fine retinotopy, population activity in the dorsal cortex accurately encodes spatial information [17], indicating that a global analysis of the visual scene can arise from distributed representations. Similarly, in zebrafish, orientation-selective neurons have been observed in the optic tectum, and some studies report their laminar segregation and retinotopic organisation [40, 43]. However, while all three species exhibit orientation-selective neurons, neither has demonstrated the presence of continuous, spatially organised orientation maps comparable to the ones found in the mammalian cortex.

Orientation coding in *Drosophila*. In *Drosophila melanogaster*, orientation tuning has been hypothesised based on anatomical structure alone. In particular, Seung [45] predicted that functionally distinct neuron types such as Dm3 and TmY may exhibit orientation selectivity, proposing that such responses could arise purely based on wiring alone. To date, the only direct physiological evidence of orientation selectivity comes from earlier studies on T4 and T5 neurons, which exhibit tuning to oriented edges in addition to their well-known direction selectivity [16]. Beyond these, no other neuron types have been experimentally or computationally confirmed to show orientation selectivity. Moreover, no prior study has computationally demonstrated the emergence of orientation maps in this non-cortical system. Our work fills in this gap by building on these findings to show that coherent, spatially organised orientation maps can emerge in such a system, thus suggesting that species of vastly different evolutionary origin may share common circuit-level principles for encoding visual features.

3 Methods

LIF model. We implemented a leaky integrate-and-fire (LIF) framework to simulate the spiking dynamics of neurons in the *Drosophila* visual system (see Appendix B.1) [46]. Using the full adult fly brain (FAFB) connectome, we constructed a network in which neurons are labelled and connected according to their anatomical synapses [8, 13, 20, 50]. The model encompasses the complete adult *Drosophila* connectome, including 138,639 neurons and 1,508,983 synapses. The FAFB connectome provides a complete and cell-resolved reconstruction of the *Drosophila* brain [50], and has become a foundational resource for structural annotation [44], functional inference [49, 11], and whole-brain spiking simulations validated against behaviour [46]. In our model, visual stimuli were simulated as Poisson spike trains injected into lamina inputs (L1-L3) in the right eye of the *Drosophila*, and the resulting activity was propagated through the whole brain connectome using the embedded LIF framework, allowing us to monitor orientation tuning in downstream visual neurons (Figure 2a).

Stimuli. We focused on the lamina neurons L1-L3, which are visual columnar neurons - retinotopically organised neurons that are associated with an individual ommatidium in the compound eye, such that each visual column contains a dedicated copy of the neuron [39]. Unlike other columnar neurons, L1-L3 receive direct input from photoreceptors R1-R6, which capture brightness change across the visual field. The ommatidia themselves are arranged in a hexagonal lattice, forming a precise sampling grid over visual space. For this study, we applied input directly to L1-L3 rather than R1-R6, as the connectomic dataset provides visual columnar mappings for L1-L3 but not for photoreceptors [8, 13, 14, 20, 33, 35, 36, 44, 50, 39]. Naively, one can let an L1-L3 neuron receive a fixed strength of Poisson input regardless of the location of the OFF-bar on the receptive field (RF) of the corresponding ommatidium. However, this ignored spatial gradients across the RF. To better reflect spatial structure, the input stimuli were modelled by computing the distance d from each ommatidium to the bar and making the input strength depend on d based on the calcium imaging data [29] (Appendix B.2). Figure 2b illustrates how the distance of the OFF-bar modulates input strength. In the top row, each panel shows a single ommatidium, colour-coded by its corresponding L1-L3 firing rate as a function of d from the OFF-bar (black line). Firing rates increased as the bar approached the centre of the ommatidia, peaking when aligned and tapering off with distance. The

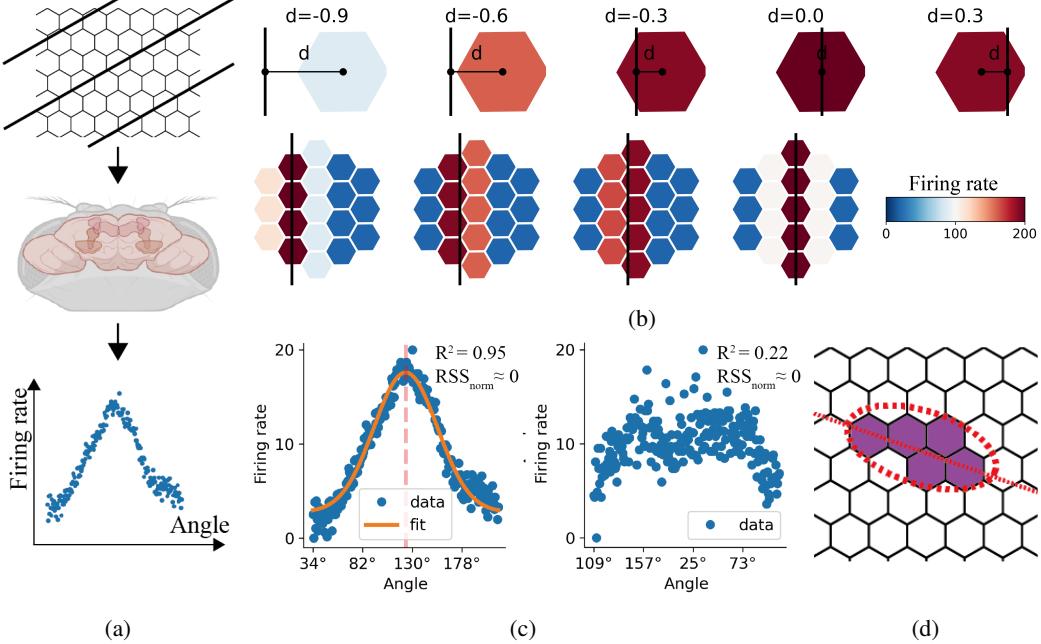


Figure 2: Overview of the simulation framework and analysis methods. (a) Schematic of the simulation pipeline. Poisson spike trains were used to simulate grating stimuli input to L1-L3 neurons. Each neuron's orientation tuning is quantified by its firing rate across stimulus angles. (b) Structured stimulus input to L1-L3 neurons across different OFF-bar positions (black line). Top: Each panel shows a hexagon representing an ommatidium. Bottom: 2D heatmap of firing activity across the simplified hexagonal lattice of ommatidia. Each hexagon represents one ommatidium, colour-coded by its corresponding L1-L3 firing rate. d represents the distance to the centre of the ommatidium. Please note that, even when the bar stimulus lies outside the geometric boundary of an ommatidium, a measurable response is still observed due to the broader receptive field (5° - 8°) relative to the ommatidial spacing (4.6°), resulting in spatial overlap of sensitivity across neighbouring units[5]. (c) Gaussian turning curves of two neurons: the left panel indicates a well-fit neuron (orientation-selective), and the right panel indicates a neuron classified as not orientation-selective. (d) Structural prediction of preferred orientation angles based on upstream columnar connectivity. The red dashed lines indicate the best-fit ellipse drawn through these spatial locations.

bottom row provides a simplified illustration of the 2D activation pattern across the hexagonal lattice for several OFF-bar positions.

Preferred orientation. To quantify the orientation selectivity of the monitored neurons, we fit a Gaussian function to their firing responses across different stimulus orientations. Unlike a standard Gaussian, which assumes a linear domain, our model accounts for the 180° periodicity of angle space by computing the minimal circular distance between the stimulus orientation and the neuron's preferred orientation (Figure 2c; see Appendix B.3). A fit was considered "good" if it satisfied two criteria: a coefficient of determination $R^2 \geq 0.7$ and a normalised residual sum of squares $RSS_{norm} \leq 0.4$ – indicating both high explanatory power and low relative error (Figure 2c). Representative examples of Gaussian fits that are considered good, poor and near threshold are illustrated in Figure S4a. The distribution of R^2 across neurons and the relationship between R^2 and RSS_{norm} are illustrated in Figure S4b, S4c. Unless otherwise noted, the preferred orientation of a neuron in this paper refers to its orientation preference as determined by this Gaussian fit.

Structural prediction. In this paper, structured prediction refers to the predicted preferred orientation angle of a neuron by its direct upstream neurons, similar to the receptive field approach used by Seung [45]. Specifically, we identify each neuron's upstream partners that are also columnar neurons using known synaptic connectivity. We mapped these upstream neurons to their corresponding columns in the compound eye, and a best-fit ellipse is drawn around their spatial locations (Figure 2d).

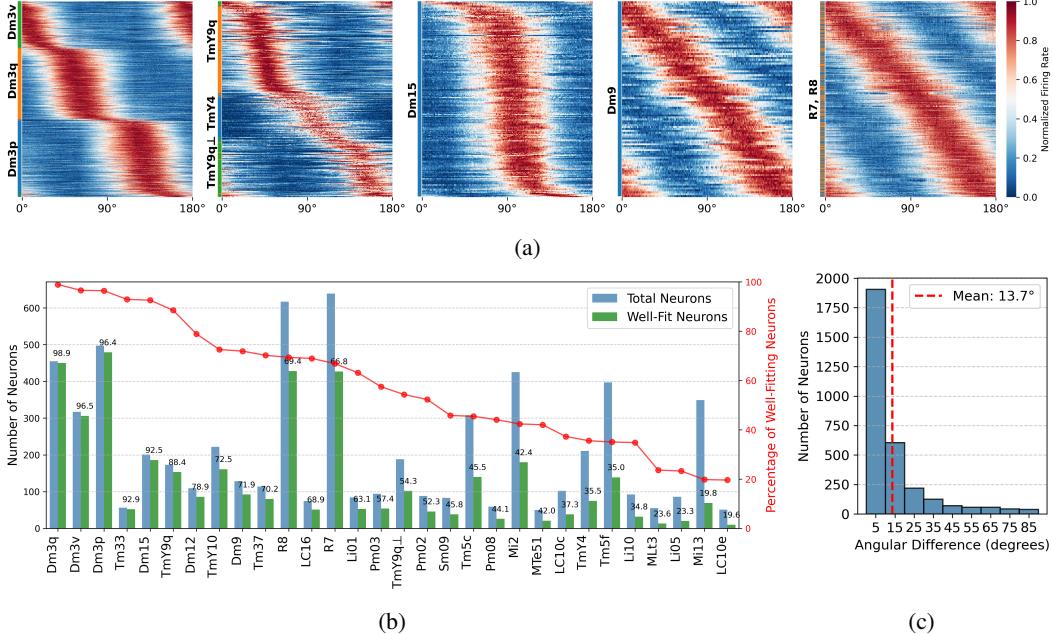


Figure 3: **Orientation tuning and structural prediction accuracy.** (a) Heatmaps showing the normalised firing rates of various neurons across different angular orientations. (b) Proportion of neurons with well-fit orientation tuning across types. Green bars indicate the subset well-fitted by a circular Gaussian model (see methods). Blue bars indicate the total number of neurons in the right optic lobe of the respective type. (c) Distribution of absolute differences between each neuron’s preferred orientations predicted by LIF simulation and by dendritic structure.

The orientation of the ellipse’s major axis, measured relative to the y-axis, is taken as the predicted angle for the downstream neuron.

4 Results

4.1 Orientation selectivity

We first evaluated whether neurons in the *Drosophila* optic lobe exhibit robust orientation selectivity. Using our LIF-based simulation framework, we presented bar-like stimuli across orientations and recorded their firing rates. We then fitted each neuron’s response profile using a circular Gaussian model (see Methods). For neurons that were well-fit, we normalised firing rates to their peak responses and visualised the resulting heatmap of preferred orientations. Neurons were initially sorted via hierarchical clustering, revealing structured patterns in tuning preferences. Upon observing a diagonal-like structure - suggestive of a continuous orientation gradient - we further sorted neurons by the ascending position of their peak response, which highlighted a clearer gradient in orientation tuning (Figure 3a). Our findings demonstrated that *Drosophila* neurons exhibit robust orientation selectivity, with evidence suggesting that orientation maps in this species may exist. The percentage of well-fit neurons relative to the total neurons in the type for a subset of neurons (Dm3v, Dm3p, Dm3q, TmY4, TmY9q, TmY9q[⊥]) aligns well with Seung’s prediction [45], indicating that orientation selectivity arises directly from the spatial arrangement of dendritic inputs (left two panels of Figure 3a).

Moreover, we found that several other neuron types (e.g., Dm15, Tm33, TmY10) also exhibited tuning properties (Figure 3a). To quantify this selectivity, we classified a neuron type as orientation selective if over 40% of its neurons were well fit (Figure 3b; see Methods). These results revealed that most orientation-selective neurons were located within the medulla layer, highlighting it as a key site for early orientation processing in the *Drosophila* visual system. To further investigate the tuning mechanism and verify our hypothesis that orientation tuning could be directly influenced

by its dendritic inputs, we quantified the preferred orientation of each neuron by computing the structural prediction of the neuron (Figure 2a; see Methods). We then compared it to the neuron's preferred orientation using absolute difference. Across the population analysed, the mean absolute difference sits at 13.7° (Figure 3c), indicating a robust correspondence between dendritic geometry and orientation preference. Notably, T4 and T5 neurons, despite being known for their direction and orientation selectivity, did not meet our selection criteria in this analysis. Our results showed that most T4 and T5 neurons had firing rates close to zero. A few neurons that had higher firing rates exhibited orientation selectivity (Figure S5). This was likely due to the stimulus design: the static input may have been insufficient to strongly activate these neurons.

Intriguingly, while analysing the orientation tuning patterns across neuron populations, we incidentally observed that photoreceptor R7 and R8, classically known for detecting colours [36], also exhibit orientation selectivity (Figure 3b). This unexpected finding raises the question of how such tuning arises in primary sensory neurons. We investigated the connectivity of the neurons and found that Dm9, R7 and R8 formed a closed recurrent loop that results in R7 and R8 adopting the same orientation preferences as Dm9 (Figure 1b and right two panels of Figure 3a), which itself receives input from L3. This closed recurrent loop between Dm9, R7 and R8, leading to synchronised orientation preferences, could provide new insights into how feedback within small circuits may amplify or stabilise neural responses. While this was not the primary hypothesis, it raises interesting questions about the role of recurrent circuits in orientation selectivity. Further exploration of this feedback mechanism in the future could contribute to our understanding of how local orientation selectivity is maintained or enhanced within neural networks.

4.2 Orientation maps

To evaluate whether orientation selectivity is spatially organised in the fly visual system, we focused on the medulla, a key visual processing region that receives direct input from L1-L3 neurons and contains the highest density of neurons with strong orientation selectivity (Figure 3b). Across the medulla sublayers - which are anatomically divided into distal (Dm), serpentine (Sm) and proximal (Pm) layers, shown in Figure 1a, we identified 1710 neurons in the Dm layer, of which 1245 belonged to the Dm3v, Dm3q and Dm3p subtypes. Additionally, we identified 143 well-fit neurons in the Pm layer and 63 in the Sm layer, based on the Gaussian tuning criteria (see Methods). It was found that neurons in all three sublayers exhibited clustered patterns - neurons with similar preferred orientations tended to cluster together (Figure S6). We then focused on the Dm and Pm layers for detailed spatial analysis, as these layers contained a sufficient number of neurons spanning the full spatial extent of the medulla (Figure S6).

We first visualised the spatial organisation of preferred orientation by flattening the 3D morphology of the neurons in the Dm and Pm layers onto layer-specific 2D coordinate systems. To achieve this, we computed a best-fit plane through the full set of 3D skeleton coordinates from each layer, providing a reference frame for spatial alignment and visualisation. This approach preserved the columnar layout while minimising distortions introduced by the *Drosophila*'s optic lobe curvature. To enable cross-layer comparisons, a scaling transformation was applied to normalise for differences in physical size and curvature between the Dm and Pm regions. This projection preserves relative topographic relationships while enabling 2D spatial smoothing and angular comparisons. We colour-coded the neurons into their preferred orientation across Dm and Pm layers using unsmoothed maps to examine the spatial distribution of preferred orientations (Figure S6). These maps revealed a clear spatial clustering of similar orientation preferences, in contrast to a salt-and-pepper pattern in the rodent V1 area.

To better visualise the global structure, we applied local spatial smoothing by computing the circular mean of preferred orientations within a circular neighbourhood of fixed radius r around each neuron. Specifically, for each neuron in the 2D projected plane, we identified all neurons located within a disk of radius $r \approx 5 \times 10^3$ nm and computed the circular mean of their orientation preferences. This smoothing preserved mesoscale spatial patterns while reducing local variability. The resulting heatmap produced orientation maps shown in Figure 4a. These smooth maps reveal pinwheel-like singularities - points around which orientation preference rotates continuously (the circled points in Figure 4a) - indicating a degree of spatial coherence previously unreported in invertebrate visual systems.

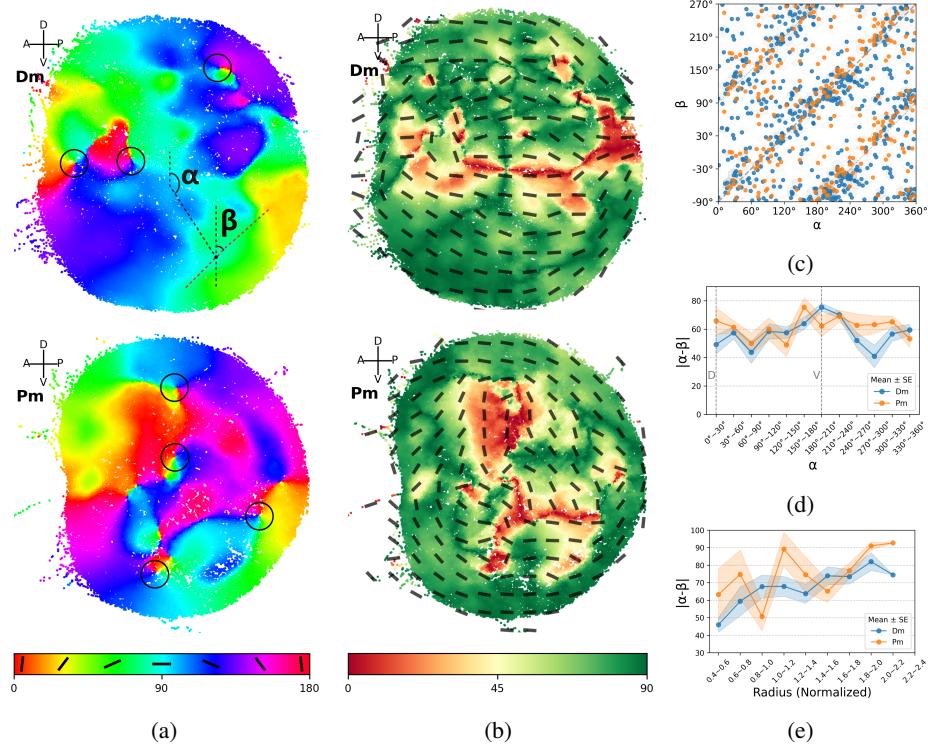


Figure 4: Spatial and orientation map in the optic lobe. (a) The smoothed heatmap of the optic lobe coloured by its preferred orientations at each layer of the optic lobe. The circles highlight the observed pinwheel structures. α denotes the position angle where dorsal-ventral represents 0° and β denotes the preferred orientation of the neurons (see Methods). (b) Vector field plots of the preferred orientation in the Dm (top) and Pm (bottom) layers. Each line represents the preferred orientation of a neuron, and the background heatmap encodes the absolute angular difference between the neuron's α and β . Green regions indicate strong tangential alignment ($\Delta \approx 90^\circ$), red regions indicate radial alignment ($\Delta \approx 0^\circ$). (c) Scatter plot (position angle, preferred orientations) pairs are displayed with coordinate axes shifted for visualisation (y-axis: -90° to 270° , x-axis: 0° to 360°). (d) Relationship of $|\alpha - \beta|$ and α . D: dorsal area; V: ventral area. (e) Relationship between $|\alpha - \beta|$ and radial distance from the centre.

Initial observations of Figure 4a revealed that both layers exhibited a roughly centrosymmetric distribution of preferred orientations. To characterise this structure more precisely, we defined a polar coordinate system for each layer, with the origin set at the centroid of the neuron population, computed by the mean of the spatial coordinates, and using the V-D (bottom-top) direction as the axis. This allowed us to compute a spatial position angle α for each neuron and directly compare it to its preferred orientation β , derived from the smoothed map in Figure 4a. Building on this smoothed representation, we treated each neuron's orientation as representative of its surrounding neighbourhood, based on the same local averaging kernel described earlier. To quantify alignment, we computed the angular difference $\Delta = |\alpha - \beta|$ for each neuron and projected it onto a 2D heatmap (Figure 4b), where $\Delta = 90^\circ$ indicates tangential alignment and $\Delta = 0^\circ$ indicates radial alignment. This visualisation revealed a clear trend: neurons in peripheral regions tended to prefer tangential orientations ($\beta \approx \alpha + 90^\circ$). These findings suggest that orientation tuning in the medulla is not randomly distributed, but spatially organised relative to each neuron's anatomical position.

We then further examined the relationship between α and β for all neurons in Dm and Pm layers. Here, each neuron's position angle was computed as the circular mean of the angular positions of all its voxels in the 2D projection. The resulting scatter plot showed a clear diagonal band, consistent with a systematic offset where $\beta \approx \alpha + 90^\circ$, supporting the presence of tangential alignment (Figure 4c). To quantify this relationship, we computed Δ and examined how it varied across spatial dimensions. Across angular bins, the Δ remains relatively stable with minor regional differences

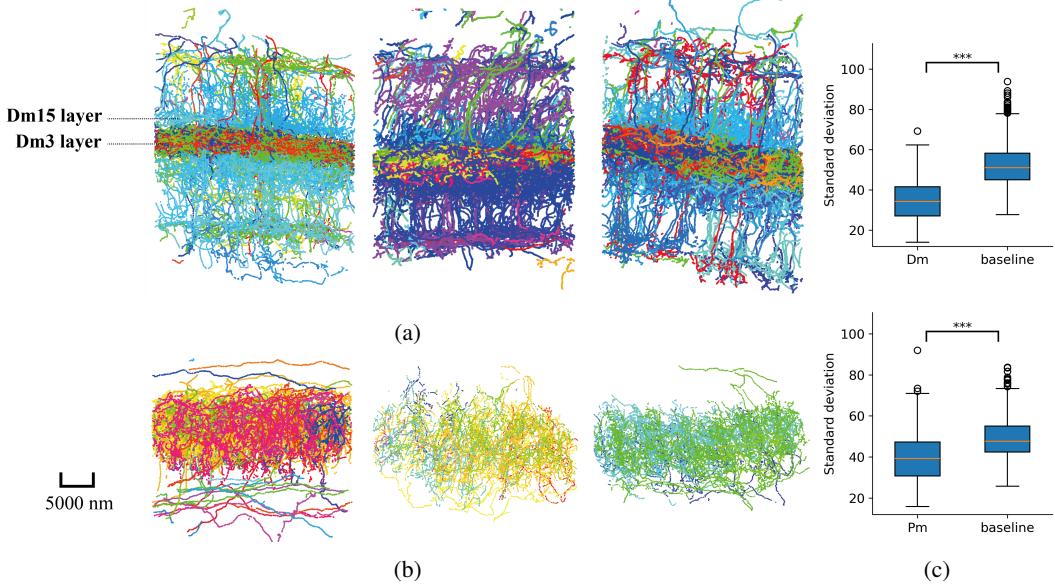


Figure 5: Orientation column structure in the optic lobe. (a) Three examples of orientation columns in the Dm layer, colour-coded by preferred orientation. (b) Three examples of orientation columns in the Pm layer, colour-coded by preferred orientation. (c) Quantification of tuning consistency within columns. Top: Dm layer. Bottom: Pm layer. *** indicates significance assessed via Mann-Whitney U test comparing standard deviations to baseline ($p < 10^{-20}$).

(Figure 4d). In contrast, alignment improved monotonically with radial distance from the centre (Figure 4e), indicating stronger alignment at the periphery. Although perfect tangential alignment would yield $\Delta \approx 90^\circ$, the observed means approaching 90° suggest a systematic but approximate offset. We speculate that radial and tangential recognition neurons organise an overall pattern in the medulla corresponding to the horizontal direction. This pattern is especially evident in the Dm layer, suggesting a non-uniform organisation of orientation preferences, which may play a key role in the previously reported recognition of horizons and horizontal objects by the biological flight control system [47].

This finding suggests that orientation tuning in *Drosophila*'s optic lobe is not randomly distributed but exhibits a structured, position-dependent organisation. To access the contribution of upstream visual pathways to orientation selectivity, we performed a targeted ablation analysis (Appendix C). Silencing Mi neurons abolished tuning across both Dm and Pm layers, while silencing Tm neurons produced selective loss in Pm but minimal effect in Dm, suggesting that Mi neurons provide essential input for orientation tuning throughout the optic lobe. Consistent with this observation, an analysis of synaptic connectivity by orientation preference (Appendix D) revealed that excitatory neurons with similar tuning are more likely to form recurrent connections, supporting topographic organisation and collinear facilitation as proposed by Seung [45]. In contrast, inhibitory connections displayed distinct off-diagonal structure, suggesting selective cross-orientation motifs that may sharpen tuning and maintain balance within the network. Peripheral neurons preferentially align their orientation selectivity tangentially relative to the centre of the optic lobe, reminiscent of contour-aligned representations seen in higher animals.

4.3 Orientation columns

We next asked whether such tuning is preserved across the depth of the medulla. We examined the existence of column structures - clusters of neurons aligned vertically across layers that share similar orientation preferences. Putative orientation columns were defined by anchoring each analysis region around a Tm3 neuron, chosen due to their well-defined columnar morphology. For each anchor neuron, we extracted the surrounding neurons within a cylindrical region of 1.5×10^4 nm radius for both the Dm and Pm layers according to their skeleton coordinates. Each such grouping was considered as a single column. A total of 858 columns were extracted, each containing on average

approximately 158.29 neurons in the Dm layers, of which 124.14 belong to Dm3 and Dm15 neuron types and 29.7 neurons in the Pm layers.

We visualised the neuron morphology by projecting a rectangular slice passing through the centre of the cylinder onto a 2D plane. This plane was defined by the two orthogonal axes perpendicular to the principal axis - the longest direction along which the neuron's morphology extends - of the Tm3 skeleton, ensuring that the projection aligned with the local column geometry. Specifically, neuron positions were transformed into this local coordinate frame by projecting each skeleton point onto the two orthogonal basis vectors to the Tm3 axis. This allowed us to flatten the 3D column into a 2D view while preserving vertical alignment across layers. We randomly selected three distinct columns each from Dm (Figure 5a) and Pm (Figure 5b) for visualisation. The Dm layer did exhibit orientation column organisations, with vertically clustered neurons sharing similar orientation preferences. However, Dm3 (Dm3v, Dm3p, Dm3q) and Dm15 neurons seemed to form a distinct horizontal layer that bisects each column into upper and lower segments, creating a clear visual stratification within individual columns (Figure 5a). This pattern was observed consistently across columns, indicating a widespread structural feature in the medulla. Notably, these neuron types exhibited only a single orientation preference (Dm3v: $\approx 0^\circ$; Dm3p: $\approx 60^\circ$; Dm3q: $\approx 120^\circ$; Dm15: $\approx 100^\circ$), in contrast to other types that exhibited broader tuning curves or gradual changes in preferred orientation across space (Figure 2a). On the other hand, the Pm layer also exhibited orientation column structure (Figure 5b), though there are a few columns that seem less pronounced in the structure (left two panels of Figure 5b).

To quantify this alignment, we measured the circular standard deviation (see Appendix B.4) of preferred orientations within each column in the Dm layer. The circular standard deviation was computed over the preferred orientations of all voxels within the column region, providing a measure of how tightly aligned the local tuning preferences are. Lower values indicated stronger orientation coherence within the column. To further ensure that the observed orientation columns were not an artefact of the simulation framework or analysis pipeline, we generated a null baseline by randomising the preferred orientations across neurons while preserving their spatial locations and column definitions. The resulting distribution revealed that most Dm columns, excluding Dm3 and Dm15, exhibited sharp tuning, with a peak at low deviation values and a minority displaying broader orientation spread (top panel of Figure 5c). We repeated the same analysis for columns in the Pm layer. In contrast to Dm, the Pm layer showed weaker columnar correspondence, with a distribution of standard deviations centred closer to the shuffled baseline (bottom panel of Figure 5c). Compared to the Pm layer, the Dm layer exhibited a more pronounced separation between the observed and baseline distributions, indicating stronger tuning coherence.

Together, these results demonstrate that *Drosophila*'s optic lobe contains not only intra-layer orientation maps, but also orientation columns across different layers of Dm, reminiscent of columnar architectures in vertebrate visual systems.

5 Conclusion and Discussion

In this study, we investigated the orientation preferences in the *Drosophila* visual system and explored the spatial organisation of these preferences, revealing a structured pattern. While these maps are less discretely defined than in mammals, likely due to the lower neuronal density, the presence of spatially organised orientation selectivity and columnar structure suggests that *Drosophila* possesses a rudimentary form of orientation map. Biologically, the organisation shows a population-level bias toward tangential tuning, where neurons prefer orientations aligned with the local contour of the visual field and may reflect an adaptation for encoding object boundaries or motion, enhancing contrast sensitivity across the curved surface of the compound eye [6].

The results align with studies in other non-mammalian species, such as birds, turtles and zebrafishes, where sharply tuned orientation-selective neurons are found in non-laminated visual pathways as mentioned earlier [32, 17, 40, 43]. However, as of now, *Drosophila* is currently the only invertebrate where spatially structured orientation preferences, resembling a primitive orientation map, have been computationally demonstrated. As such, we hypothesise that structured visual maps may reflect convergent evolution, arising from shared computational demands rather than shared anatomy. The *Drosophila* ommatidium \rightarrow Dm \rightarrow lobula circuit may serve a functionally analogous role to the mammalian retina \rightarrow LGN \rightarrow V1 pathway [1, 19, 27], suggesting that efficient visual processing

follows common architectural motifs across species. Mechanistically, dendritic integration, closed-loop motifs, and lateral inhibition likely contribute to orientation tuning in *Drosophila*, as proposed in prior works [2, 25, 45]. Consistent with this interpretation, analysis of synaptic connectivity by orientation preference (Appendix D) revealed that neurons with similar tuning are more likely to form recurrent connections. The study by Klapoetke et al. [30] reveals functionally clustered representations of complex visual features in the lobular columnar neurons projecting into the central brain. In contrast, our work focuses on orientation selectivity in the earlier stages of the optic lobe, based on connectome-driven circuit structure. Together, these complementary findings highlight how structured visual maps may extend from early visual encoding to higher-order feature integration within the *Drosophila* brain.

The orientation maps at the early stages of visual processing provide critical advantages for efficient visual detection for the flies. The neurons that selectively respond to similar orientations are spatially adjacent to each other, enabling them to fire maximally when a line or edge appears at their preferred orientation, which in mammals is achieved by long-distance connections among columns with similar orientation preference. The columns are arranged in a primitive pinwheel-like map, enabling more stable and precise edge perception. This organisation supports higher-level vision by feeding processed orientation data to the central brain for shape recognition and motion detection. *Drosophila* serves as a powerful minimal model for exploring core principles of vision, with potential applications in bio-inspired machine vision and low-power autonomous systems.

Limitations. This study relies on biologically grounded simulations and lacks in vivo validation, warranting experimental follow-up to confirm our predictions. Similar to prior work based on anatomical reconstructions [45], our approach emphasised predictive modelling to inform future experiments. The use of a simplified LIF model omits nonlinear firing dynamics and neurotransmitter effects, which may affect network behaviour. However, similar simplified models have been effectively used in large-scale cortical simulations to gain insights into emergent network dynamics [46, 18].

Acknowledgments and Disclosure of Funding

This work was supported by the National Key Research and Development Program of China under Grant 2021ZD0200301, the National Natural Science Foundation of China under Grant U2341228, the Fundamental and Interdisciplinary Disciplines Breakthrough Plan of the Ministry of Education of China under Grant JYB2025XDXM504, and Ant Group under Grant 20252000099.

References

- [1] Jose-Manuel Alonso, W. Martin Usrey, and R. Clay Reid. Rules of connectivity between geniculate cells and simple cells in cat primary visual cortex. *Journal of Neuroscience*, 21:4002–4015, 2001. doi: <https://doi.org/10.1523/JNEUROSCI.21-11-04002.2001>.
- [2] Alessandra Angelucci and Paul C. Bressloff. Contribution of feedforward, lateral and feedback connections to the classical receptive field center and extra-classical receptive field surround of primate v1 neurons. *Progress in Brain Research*, 154:93–120, 2006. doi: [https://doi.org/10.1016/S0079-6123\(06\)54005-1](https://doi.org/10.1016/S0079-6123(06)54005-1).
- [3] Gary Blasdel, Klaus Obermayer, and Lynne Kiorpis. Organization of ocular dominance and orientation columns in the striate cortex of neonatal macaque monkeys. *Visual Neuroscience*, 12:589–603, 1995. doi: <https://doi.org/10.1017/S0952523800008476>.
- [4] Gary G. Blasdel and Guy Salama. Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. *Nature*, 321:579–585, 1986. ISSN 0028-0836. doi: <https://doi.org/10.1038/321579a0>.
- [5] Alexander Borst. *Drosophila*'s view on insect vision. *Current Biology*, 19(1):R36–R47, 2009. ISSN 0960-9822. doi: <https://doi.org/10.1016/j.cub.2008.11.001>. URL <https://www.sciencedirect.com/science/article/pii/S0960982208014322>.
- [6] Alexander Borst and Thomas Euler. Seeing things in motion: models, circuits, and mechanisms. *Neuron*, 71:974–994, 2011. doi: <https://doi.org/10.1016/j.neuron.2011.08.031>.

- [7] Alexander Borst and Moritz Helmstaedter. Common circuit design in fly and mammalian motion vision. *Nature Neuroscience*, 18:1067–1076, 2015. doi: <https://doi.org/10.1038/nn.4050>.
- [8] Julia Buhmann, Arlo Sheridan, Caroline Malin-Mayor, Philipp Schlegel, Stephan Gerhard, Tom Kazimiers, Renate Krause, Tri M. Nguyen, Larissa Heinrich, Wei-Chung Allen Lee, Rachel Wilson, Stephan Saalfeld, Gregory S.X.E. Jefferis, Davi D. Bock, Srinivas C. Turaga, Matthew Cook, and Jan Funke. Automatic detection of synaptic partners in a whole-brain drosophila electron microscopy dataset. *Nature Methods*, 18:771–774, 2021. ISSN 1548-7091. doi: <https://doi.org/10.1038/s41592-021-01183-7>.
- [9] Matthew A. Churgin, Danylo O. Lavrentovich, Matthew A-Y Smith, Ruixuan Gao, Edward S Boyden, and Benjamin L de Bivort. A neural correlate of individual order preference in drosophila. *eLife*, 12:RP90511, 2023. doi: <https://doi.org/10.7554/eLife.90511>.
- [10] Damon A. Clark, Limor Bursztyn, Mark A. Horowitz, Mark J. Schnitzer, and Thomas R. Clandinin. Defining the computational structure of the motion detector in drosophila. *Neuron*, 70:1165–1177, 2011. doi: <https://doi.org/10.1016/j.neuron.2011.05.023>.
- [11] Benjamin R. Cowley, Adam J. Calhoun, Nivedita Rangaranjan, Elise Ireland, Maxwell H. Turner, Jonathan W. Pillow, and Mala Murthy. Mapping model units to visualise neurons reveals population code for social behaviour. *Nature*, 629:1100–1108, 2024. doi: <https://doi.org/10.1038/s41586-024-07451-8>.
- [12] Michael C. Crair, Deda C. Gillespie, and Michael P. Stryker. The role of visual experience in the development of columns in cat visual cortex. *Science*, 279:566–570, 1997. ISSN 0036-8075. doi: <https://doi.org/10.1126/science.279.5350.566>.
- [13] Sven Dorkenwald, Arie Matsliah, Amy R. Sterling, Philipp Schlegel, Szi-chieh Yu, Claire E. Mckellar, Albert Lin, Marta Costa, Katharina Eichler, Yijie Yin, Will Silversmith, Casey Schneider-Mizel, Chris S. Jordan, Derrick Brittain, Akhilesh Halageri, Kai Kuehner, Oluwaseun Ogedengbe, Ryan Morey, Jay Gager, Krzystof Kruck, Eric Perlman, Runzhe Yang, David Deutsch, Doug Bland, and The FlyWire Consortium. Neuronal wiring diagram of an adult brain. *Nature*, 634:124–138, 2024. ISSN 0039-0836. doi: <https://doi.org/10.1038/s41586-024-07558-y>.
- [14] Nils Eckstein, Alexander Shakeel Bates, Andrew Champion, Michelle Du, Yijie Yin, Philipp Schlegel, Alicia Kun-Yang Lu, Thomson Rymer, Samantha Finley-May, Tyler Paterson, Ruchi Parekh, Sven Dorkenwald, Arie Matsliah, Szi-Chieh Yu, Claire Mckellar, Amy Sterling, Katharina Eichler, Marta Costa, Sebastian Seung, Mala Murthy, Volker Hartenstein, Gregory S.X.E. Jefferis, and Jan Funke. Neurotransmitter classification from electron microscopy images at synaptic sites in drosophila melanogaster. *Cell*, 187:2574–2594, 2024. doi: <https://doi.org/10.1016/j.cell.2024.03.016>.
- [15] David Ferster and Kenneth D. Miller. Neural mechanisms of orientation selectivity in the visual cortex. *Annual Review of Neuroscience*, 23:441–471, 2000. doi: <https://doi.org/10.1146/annurev.neuro.23.1.441>.
- [16] Yvette E. Fisher, Marion Silies, and Thomas R. Clandinin. Orientation selectivity sharpens motion detection in drosophila. *Neuron*, 88:390–402, 2015. doi: <https://doi.org/10.1016/j.neuron.2015.09.033>.
- [17] Julien Fournier, Christian M. Müller, Ingmar Schneider, and Gilles Laurent. Spatial information in a non-retinotopic visual cortex. *Neuron*, 97:164–180e.7, 2018. doi: <https://doi.org/10.1016/j.neuron.2017.11.017>.
- [18] Jake Gavenas, Ueli Rutishauser, and Uri Maoz. Slow ramping emerges from spontaneous fluctuations in spiking neural networks. *Nature Communications*, 15:7285, 2024. doi: <https://doi.org/10.1038/s41467-024-51401-x>.
- [19] David Hansel and Carl van Vreeswijk. The mechanism of orientation selectivity in primary visual cortex without a functional map. *Journal of Neuroscience*, 12:4049–4064, 2012. doi: <https://doi.org/10.1523/JNEUROSCI.6284-11.2012>.

[20] Larissa Heinrich, Jan Funke, Constantin Pape, Juan Nunez-Iglesias, and Stephan Saalfeld. Synaptic cleft segmentation in non-isotropic volume electron microscopy of the complete drosophila brain. In *Medical Image Computing and Computer Assisted Intervention - MICCAI 2018. Lecture Notes in Computer Science*, 2018. ISBN 978-3-030-00933-5. doi: https://doi.org/10.1007/978-3-030-00934-2_36.

[21] Yu-Chi Huang, Cheng-Te Wang, Ta-Shun Su, Kuo-Wei Kao, Yen-Jen Lin, Chao-Chun Chuang, Ann-Shyn Chiang, and Chung-Chuan Lo. A single-cell level and connectome-derived computational model of the drosophila brain. *Frontiers in Neuroinformatics*, 12, 2018. doi: <https://doi.org/10.3389/fninf.2018.00099>.

[22] David H. Hubel and Torsten N. Wiesel. Receptive fields of single neurones in the cat's striate cortex. *The Journal of Physiology*, 148:574–591, 1959. doi: <https://doi.org/10.1113/jphysiol.1959.sp006308>.

[23] David H. Hubel and Torsten N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, 160:106–154.2, 1962. doi: <https://doi.org/10.1113/jphysiol.1962.sp006837>.

[24] Young Jun Jung, Ali Almasi, Shi H. Sun, Molis Yunzab, Shaun L. Cloherty, Sebastien H. Bauquier, Marilyn Renfree, Hamish Meffin, and Michael R. Ibbotson. Orientation pinwheels in primary visual cortex of a highly visual marsupial. *Science Advances*, 8:eabn0954, 2022. doi: <https://doi.org/10.1126/sciadv.abn0954>.

[25] Young Jun Jung, Ali Almasi, Shi Sun, Molis Yunzab, Sebastien H. Baquier, Marilyn Renfree, Hamish Meffin, and Michael R. Ibbotson. Feature selectivity and invariance in marsupial primary visual cortex. *The Journal of Physiology*, 603:423–445, 2024. ISSN 0022-3751. doi: <https://doi.org/10.1113/JP285757>.

[26] Anna-Maria Jürgensen, Afshin Khalili, Elisabetta Chicca, Giacomo Indiveri, and Martin Paul Nawrot. A neuromorphic model of olfactory processing and sparse coding in the drosophila larva brain. *Neuromorphic Computing and Engineering*, 1, 2021. doi: <https://doi.org/10.1088/2634-4386/ac3ba6>.

[27] Hulusi Kafaligonul, Bruno G. Breitmeyer, and Haluk Öğmen. Feedforward and feedback processes in vision. *Frontiers in Psychology*, 6, 2015. doi: <https://doi.org/10.3389/fpsyg.2015.00279>.

[28] Kyobi S. Kakaria and Benjamin L. de Bivort. Ring attractor dynamics emerge from a spiking model of the entire protocerebral bridge. *Frontiers in Behavioral Neuroscience*, 11, 2017. doi: <https://doi.org/10.3389/fnbeh.2017.00008>.

[29] Madhura D. Ketkar, Burak Gür, Sebastian Molina-Obando, Maria Ioannidou, Carlotta Martelli, and Marion Silies. First-order visual interneurons distribute distinct contrast and luminance information across on and off pathways to achieve stable behavior. *eLife*, 11:e74937, 2022. doi: <https://doi.org/10.7554/eLife.74937>.

[30] Nathan C. Klapoetke, Aljoscha Nern, Edward M. Rogers, Gerald M. Rubin, Michael B. Reiser, and Gwyneth M. Card. A functionally ordered visual feature map in the drosophila brain. *Neuron*, 110:1700–1711, 2022. doi: <https://doi.org/10.1016/j.neuron.2022.02.013>.

[31] Aurel A. Lazar, Tingkai Liu, Mehmet Kerem Turkcan, and Zhoum Yiyin. Accelerating the flybrainlab the discovery of the functional logic of the drosophila brain in the connectomic and synaptomic era. *eLife*, 10:e62362, 2021. doi: <https://doi.org/10.7554/eLife.62362>.

[32] Da-Peng Li, Qian Xiao, and Shu-Rong Wang. Feedforward construction of the receptive field and orientation selectivity of visual neurons in the pigeon. *Cerebral Cortex*, 17:885–893, 2007. doi: <https://doi.org/10.1093/cercor/bhk043>.

[33] Albert Lin, Runzhe Yang, Sven Dorkenwald, Arie Matsliah, Amy R. Sterling, Philipp Schlegel, Szi-chieh Yu, Claire E. McKellar, Marta Costa, Katharina Eichler, Alexander Shakeel Bates, Nils Eckstein, Jan Funke, Gregory S.X.E. Jefferis, and Mala Murthy. Network statistics of the whole-brain connectome of drosophila. *Nature*, 634:153–165, 2024. ISSN 0028-0836. doi: <https://doi.org/10.1038/s41586-024-07968-y>.

[34] Matthew S. Maisak, Juergen Haag, Georg Ammer, Etienne Serbe, Matthias Meier, Aljoscha Leonhardt, Tabea Schilling, Armin Bahl, Gerald M. Rubin, Aljoscha Nern, Barry J. Dickson, Dierk F. Reiff, Elisabeth Hopp, and Alexander Borst. A directional tuning map of drosophila elementary motion detectors. *Nature*, 500:212–216, 2013. doi: <https://doi.org/10.1038/nature12320>.

[35] Arie Matsliah, Amy R. Sterling, Sven Dorkenwald, Kai Kuehner, Ryan Morey, H. Sebastian Seung, and Mala Murthy. Codex: connectome data explorer. 2023. doi: <https://doi.org/10.13140/RG.2.2.35928.67844>.

[36] Arie Matsliah, Yu Szi-chieh, Krzysztof Kruk, Doug Bland, Austin T. Burke, Jay Gager, James Hebditch, Ben Silverman, Kyle Patrick Willie, Ryan Willie, Marissa Sorek, Amy R. Sterling, Emil Kind, Dustin Garner, Gizem Sancer, Mathias F. Wernet, Sung Soo Kim, Mala Murthy, H. Sebastian Seung, and The FlyWire Consortium. Neuronal parts list and wiring diagram for a visual system. *Nature*, 634:166–180, 2024. doi: <https://doi.org/10.1038/s41586-024-07981-1>.

[37] N. Merkulyeva, V. Lyakhovetskii, and A. Mikhalkin. Anisotropy of the orientation selectivity in the visual cortex area 18 of cats reared under normal and altered visual experience. *European Journal of Neuroscience*, 61(2):e70004, 2025. ISSN 0953-816X. doi: <https://doi.org/10.1111/ejn.70004>.

[38] Ian Nauhaus, Andrea Benucci, Matteo Carandini, and Dario L. Ringach. Neuronal selectivity and local map structure in visual cortex. *Neuron*, 57:673–679, 2008. doi: <https://doi.org/10.1016/j.neuron.2008.01.020>.

[39] Aljoscha Nern, Frank Loesche, Shin-ya Takemura, Laura E. Burnett, Marisa Dreher, Eyal Gruntman, Judith Hoeller, Gary B. Huang, Michał Januszewski, Nathan C. Klapoetke, Sanna Koskela, Kit D. Longden, Zhiyuan Lu, Stephan Preibisch, Wei Qiu, Edward M. Rogers, Pavithraa Seenivasan, Arthur Zhao, John Bogovic, Brandon S. Canino, Jody Clements, Michael Cook, Samantha Finley-May, Miriam A. Flynn, Imran Hameed, Alexandra M. C. Fragniere, Kenneth J. Hayworth, Gary Patrick Hopkins, Philip M. Hubbard, William T. Katz, Julie Kovalyak, Shirley A. Lauchie, Meghan Leonard, Alanna Lohff, Charli A. Maldonado, Caroline Mooney, Nneoma Okeoma, Donald J. Olbris, Christopher Ordish, Tyler Paterson, Emily M. Phillips, Tobias Pietzsch, Jennifer Rivas Salinas, Patricia K. Rivlin, Philipp Schlegel, Ashley L. Scott, Louis A. Scuderi, Satoko Takemura, Iris Talebi, Alexander Thomson, Eric T. Trautman, Lowell Umayam, Claire Walsh, John J. Walsh, C. Shan Xu, Emily A. Yakal, Tansy Yang, Ting Zhao, Jan Funke, Reed George, Harald F. Hess, Gregory S. X. E. Jefferis, Christopher Knecht, Wyatt Korff, Stephen M. Plaza, Sandro Romani, Stephan Saalfeld, Louis K. Scheffer, Stuart Berg, Gerald M. Rubin, and Michael B. Reiser. Connectome-driven neural inventory of a complete visual system. *Nature*, 2025. doi: <https://doi.org/10.1038/s41586-025-08746-0>.

[40] Linda M Nevin, Estuardo Robles, Herwig Baier, and Ethan K Scott. Focusing on optic tectum circuitry through the lens of genetics. *BMC Biology*, 8, 2010. doi: <https://doi.org/10.1186/1741-7007-8-126>.

[41] Kenichi Ohki, Sooyoung Chung, Prakash Kara, Mark Hübener, Tobias Bonhoeffer, and R. Clay Reid. Highly ordered arrangement of single neurons in orientation pinwheels. *Nature*, 442: 925–928, 2006. ISSN 0028-0836. doi: <https://doi.org/10.1038/nature05019>.

[42] Mila M. Paul, Martin Pauli, Nadine Ehmann, Stefan Hallermann, Markus Sauer, Robert J. Kittel, and Manfred Heckmann. Brunchpilot and synaptotagmin collaborate to drive rapid glutamate release and active zone differentiation. *Frontier in Cellular Neuroscience*, 9, 2015. doi: <https://doi.org/10.3389/fncel.2015.00029>.

[43] Sebastián A. Romano, Thomas Pietri, Adrien Pérez-Schuster, Verónica Jouary, Mathieu Haudrechy, and Germán Sumbre. Spontaneous neuronal network dynamics reveal circuit's functional adaptations for behavior. *Neuron*, 85:1070–1085, 2015. doi: <https://doi.org/10.1016/j.neuron.2015.01.027>.

[44] Philipp Schlegel, Yijie Yin, Alexander S. Bates, Sven Dorkenwald, Katharina Eichler, Paul Brooks, Daniel S. Han, Marina Gkantia, Marcia dos Santos, Eva J. Munnelly, Griffin Badalamente, Laia Serratosa Capdevila, Varun A. Sane, Alexandra M. C. Fragniere, Ladann Kiassat,

Markus W. Pleijzier, Tomke Stürner, Imaan F. M. Tamimi, Christopher R. Dunne, Irene Salgarella, Alexandre Javier, Siqi Fang, Eric Perlman, Tom Kazimiers, Sridhar R. Jagannathan, Arie Matsliah, Amy R. Sterling, Szi-chieh Yu, Claire E. McKellar, The FlyWire Consortium, Marta Costa, H. Sebastian Seung, Mala Murthy, Volker Hartenstein, and Davi D. Bock. Whole-brain annotation and multi-connectome cell typing of *drosophila*. *Nature*, 634:139–152, 2024. ISSN 0028-0836. doi: <https://doi.org/10.1038/s41586-024-07686-5>.

[45] H. Sebastian Seung. Predicting visual function by interpreting a neuronal wiring diagram. *Nature*, 634:113–123, 2024. doi: <https://doi.org/10.1038/s41586-024-07953-5>.

[46] Philip K. Shiu, Gabriella R. Sterne, Nico Spiller, Romain Franconville, Andrea Sandoval, Joie Zhou, Neha Simha, Chan Hyuk Kang, Seongbong Yu, Jinseop S. Kim, Sven Dorkenwald, Arie Matsliah, Philipp Schlegel, Szi-chieh Yu, Claire E. McKellar, Amy Sterling, Marta Costa, Katharina Eichler, Alexander Shakeel Bates, Nils Eckstein, Jan Funke, Gregory S.X.E. Jefferis, Mala Murthy, Salil S. Bidaye, Stefanie Hampel, Andrew M. Seeds, and Kristin Scott. A *drosophila* computational brain model reveals sensorimotor processing. *Nature*, 634:210–219, 2024. ISSN 0028-0836. doi: <https://doi.org/10.1038/s41586-024-07763-9>.

[47] Andrew D Straw, Serin Lee, and Michael H Dickinson. Visual control of altitude in flying *drosophila*. *Current Biology*, 20(17):1550–1556, 2010.

[48] Dominic J. Vita, Fernanda S. Orsi, Nathan G. Stanko, and Natalie A. Clark. Development and organisation of the retinal orientation selectivity map. *Nature Communications*, 15:4829, 2024. doi: <https://doi.org/10.1038/s41467-024-49206-z>.

[49] Sydney R. Walker, Marco Peña-Garcia, and Anita V. Devineni. Connectomic analysis of taste circuits in *drosophila*. *Scientific Reports*, 26:5278, 2025. doi: <https://doi.org/10.1038/s41598-025-89088-9>.

[50] Zhihao Zheng, J. Scott Lauritzen, Eric Perlman, Camenzind G. Robinson, Matthew Nichols, Daniel Milkie, Omar Torrens, John Price, Corey B. Fisher, Nadiya Sharifi, Steven A. Calle-Schuler, Lucia Kmecova, Iqbal J. Ali, Bill Karsh, Eric T. Trautman, John A. Bogovic, Philipp Hanslovsky, Gregory S.X.E. Jefferis, Michael Kazhdan, Khaled Khairy, Stephan Saalfeld, Richard D. Fetter, and Davi D. Bock. A complete electron microscopy volume of the brain of adult *drosophila* melanoaster. *Cell*, 174:730–743, 2018. doi: <https://doi.org/10.1016/j.cell.2018.06.019>.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The paper's abstract and introduction clearly state that the work tests whether structured orientation maps can emerge in *Drosophila* purely from biological connectivity, and the results fully support this claim through connectome-driven simulations and analysis.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The paper includes a dedicated "Limitations" section where we clearly acknowledge that our findings are based on simulations without *vivo* validation, and discuss the modelling simplifications.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The paper does not include formal theoretical results, theorems, or proofs; it is an experimental and simulation-driven study grounded in biological data rather than mathematical derivations.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The paper fully discloses the experimental setup, including model parameters, connectome sources, input stimulus design, simulation details and analysis equations, providing sufficient information for others to reproduce the main results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).

(d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The full code for simulation and analysis is appended in the Github link in the Appendix A, along with sufficient instructions to reproduce the main experimental results.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The paper specifies all necessary experimental details, including the LIF model parameters, input stimulus construction, simulation setup and neuron types studied.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: The paper reports statistical robustness through visual summaries and null-model comparisons. Specifically, distributions of orientation tuning sharpness are shown using histograms, and baseline significance is assessed via shuffled control data. While

traditional error bars are not uniformly presented, the use of population-wide distribution and comparison to randomised baselines appropriately conveys statistical reliability for the claims made.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: The paper specifies the computing resources used by the simulation in Appendix A.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [\[Yes\]](#)

Justification: The research fully conforms to the NeurIPS Code of Ethics, focusing on non-invasive, simulation-based studies using publicly available biological datasets, with no involvement of human subjects, personal data, or potentially harmful applications.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.

- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes] .

Justification: The paper discusses the broader impacts of its findings by suggesting that understanding minimalistic biological circuits could inspire lightweight, sustainable AI systems.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper does not release any pretrained models, datasets, or tools that pose a risk of misuse or dual-use concerns, as it focuses solely on simulated neural network analyses based on publicly available biological data.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: The paper properly cites the FAFB connectome and other publicly available *Drosophila* datasets, referencing the original sources.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not introduce new datasets or models beyond the use of existing publicly available connectome data and simulated networks.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The study does not involve any crowdsourcing, surveys, or human subject research, as all experiments were performed using simulated neural models based on biological datasets.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.

- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve human subjects, crowdsourcing, or participant-based studies, and therefore IRB approvals is not applicable.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The paper's core research methods do not involve large language models (LLMs) as an important, original, or non-standard component.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

Supplementary Material

A Code and data availability

The full code and data can be found at https://github.com/JNLiew/flylif_orientation_maps. The model ran on an Intel(R) CPU at 2.90GHz machine using 20 threads in parallel. Each thread required approximately 72 hours to complete a single round of simulation of different angles in the range 0° - 180° with an interval of 0.4°.

B Experimental and analysis definition

B.1 LIF model design

We implemented a conductance-based leaky integrate-and-fire (LIF) model [9, 21, 26, 28, 31, 46] and integrated the FAFB dataset into this LIF model to run our simulations [8, 13, 14, 20, 33, 35, 36, 44, 50]. The membrane potential v_i of neuron i evolves as:

$$\frac{dv_i}{dt} = \frac{g_i - (v_i - V_{\text{resting}})}{T_{\text{mbr}}}, \quad (1)$$

$$\frac{dg_i}{dt} = -\frac{g_i}{\tau}, \quad (2)$$

$$g_i \leftarrow g_i + (w_{j,i} * w_{\text{syn}}) \quad \text{upon spike from neuron } j. \quad (3)$$

A spike is emitted when $v_i \geq V_{\text{threshold}}$, after which the membrane potential is reset to V_{reset} for a refractory period $T_{\text{refractory}}$. Synaptic transmission is delayed by a fixed latency T_{dly} . The model used the following parameter values obtained from *Drosophila* modelling or electrophysiology efforts [9, 26, 28, 42, 46]:

- $V_{\text{resting}} = -52$ mV: resting potential [28, 46]
- $V_{\text{reset}} = -52$ mV: reset potential [28, 46]
- $V_{\text{threshold}} = -45$ mV: spiking threshold [28, 46]
- $R_{\text{mbr}} = 10$ kΩ cm²: membrane resistance [28, 46]
- $C_{\text{mbr}} = 2$ μF cm⁻²: membrane capacitance [28, 46]
- $T_{\text{mbr}} = R_{\text{mbr}} \cdot C_{\text{mbr}}$: membrane time constant [46]
- $T_{\text{refractory}} = 2.2$ ms: refractory period [28, 31, 46]
- $\tau = 5$ ms: synaptic decay time constant [26, 46]
- $T_{\text{dly}} = 1.8$ ms: synaptic transmission delay [46]
- $w_{\text{syn}} = 1.5$ mV: synaptic weight; free parameter
- g_i : the synaptic conductance resulting from the aggregate firing of neurons presynaptic to neuron i .

w_{syn} is a free parameter representing the strength of excitatory and inhibitory postsynaptic potentials. The value was chosen to ensure sufficient firing activity across the medulla for analysis, given the low baseline firing rates of visual neurons. The number of synapses of an upstream neuron j on a downstream neuron i is represented by the connectivity $w_{j,i}$ such that if neuron j fires, the membrane potential of neuron i changes in proportion to its connectivity $w_{j,i} \cdot w_{\text{syn}}$.

To assess sensitivity to parameter perturbations, we performed a variation analysis where all parameters (except V_{rest} and $V_{\text{threshold}}$) were randomly perturbed up to a fixed maximum percentage (Figure S1a). At 10% variation, we observed minimal changes in preferred orientation across neurons (Figure S1b). At 50% orientation selectivity was noticeably degraded (Figure S1c; however, surprisingly, the ODEs (Equations 1 and 2) remained stable and did not diverge. This suggests that the system retains a degree of functional robustness even under large perturbations.

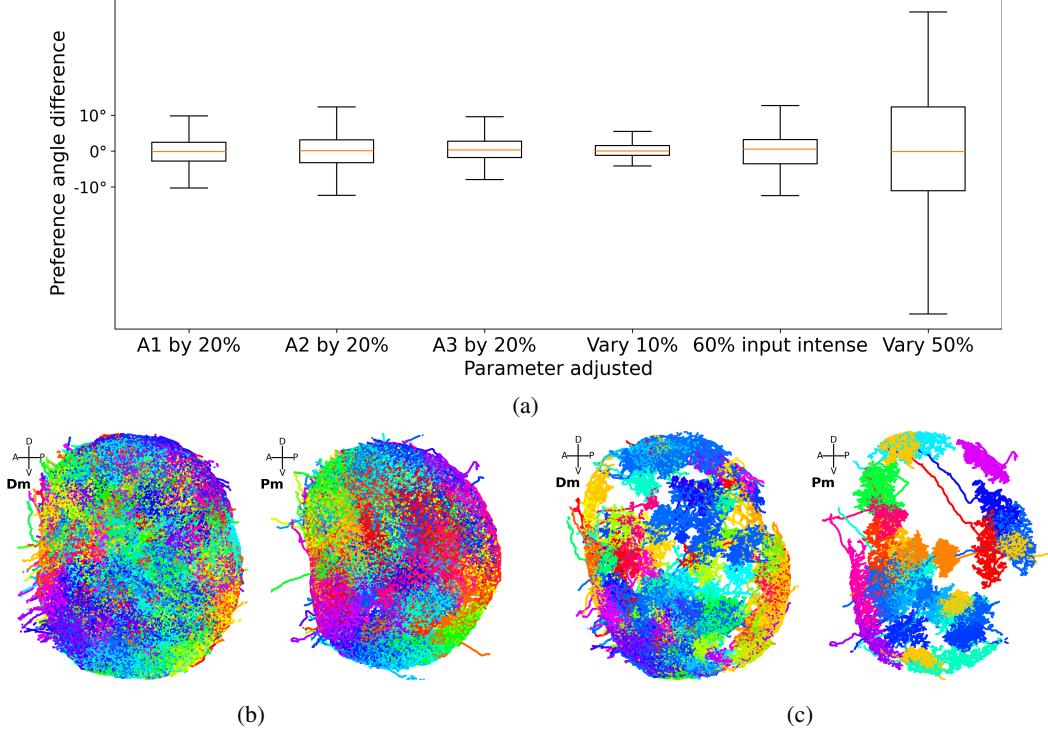


Figure S1: LIF model parameter validation. (a) Shows the difference in neurons' preferred orientation angles when stimulus parameters (A1, A2, A3) are increased by 20% respectively. "Vary 20%" indicates simulations where all LIF parameters (except V_{rest} and $V_{threshold}$) were randomly perturbed by $\pm 10\%$. "60% input intense" refers to simulations where all inputs were reduced to 60%. "Vary 50%" is similar to "Vary 10%" but instead is randomly perturbed by $\pm 50\%$. (b) Orientation maps under 10% LIF parameters variation. (c) Orientation maps under 50% LIF parameters variation.

B.2 Stimulus designs and validation

To simulate early visual input, we designed a synthetic stimulus that mimics the spatial layout and angular sampling properties of *Drosophila* photoreceptors. This allowed us to approximate the input that L1-L3 would receive *in vivo*. Our model design was guided by calcium imaging data reported in Ketkar et al. [29], which showed that L3 neurons exhibit the strongest response variation across changes in light intensity, followed by L1, with L2 showing the weakest modulation. These responses increased nonlinearly with light intensity, starting slowly and then rising more steeply. To capture this, we defined each neuron's firing rate as a function of spatial distance from a bar stimulus, which we refer to as the OFF-bar.

We assumed that spatial distance from the stimulus approximates variations in local light intensity across the visual field. The firing rate of neuron group L1-L3 was then defined as:

$$\text{FiringRate}_{Lx}(d) = \begin{cases} A_x \cdot \sqrt{1 - d^2} \cdot \frac{\max(fr)}{10}, & \text{if } |d| < 1, \\ B_x, & \text{otherwise,} \end{cases} \quad (4)$$

where d denotes the distance from the centre of the ommatidium to the OFF-bar in the visual field, measured in units where $d = 1$ corresponds to the distance between the centres of two neighbouring ommatidia. $\max(fr)$ represents the maximum firing rate, set to 200Hz. The parameters A_x and B_x , layer-specific gain and baseline levels, were estimated by fitting model responses to experimental data [29]:

- L1: $A_1 = 7, B_1 = 20$ Hz
- L2: $A_2 = 5, B_2 = 20$ Hz
- L3: $A_3 = 10, B_3 = 20$ Hz.

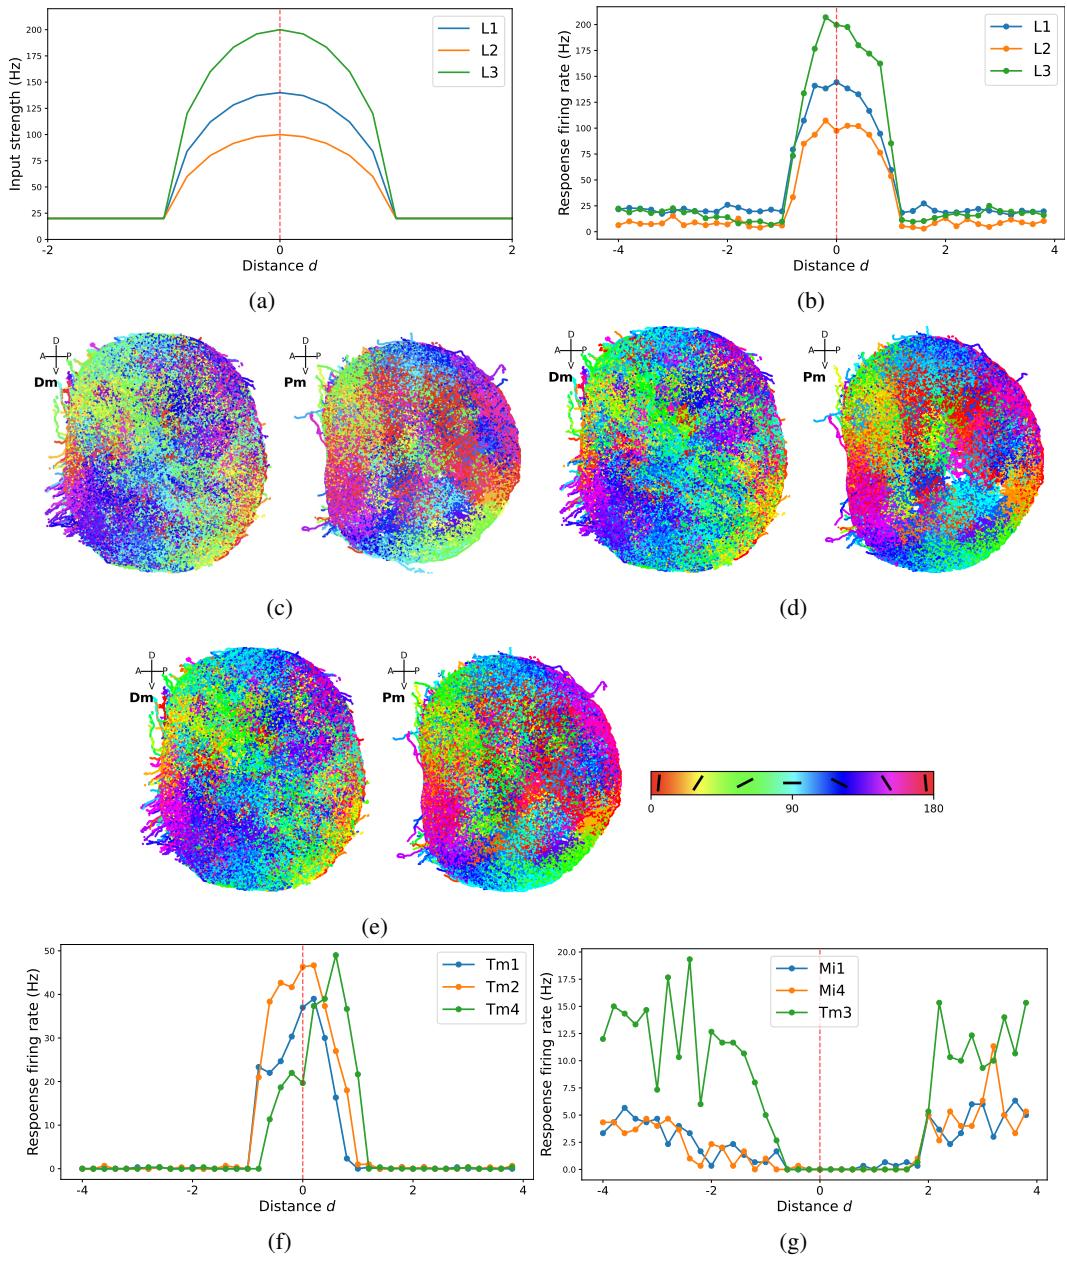


Figure S2: Activation of neuron types to bar-like stimuli. (a) Modelled input strength profiles for lamina neurons L1-L3 as a function of distance d from the OFF-bar stimulus. (b) Firing rate traces of a representative neuron from each of the L1-L3 types against their distance from the OFF-bar. The red line indicates where the stimulus OFF-bar is located. (c) Orientation maps of where A1 parameters are varied by 10%. (d) Orientation maps of where A2 parameters are varied by 10%. (e) Orientation maps of where A3 parameters are varied by 10%. (f) Firing responses of a representative downstream OFF-response neuron from each of the Tm1, Tm2 and Tm4 types against their distance from the OFF-bar. In both (f) and (g), neurons were randomly sampled to verify that OFF-bar stimuli reliably activated these neuron types as expected. (g) Firing responses of a representative downstream ON-response neuron from each of the Mi1, Mi4 and Tm3 types against their distance from the OFF-bar. OFF-bar stimuli reliably suppress firing rates in these neurons within OFF-bar regions.

The configuration ($A_3 > A_1 > A_2$) captured the relative sensitivity of L1-L3, as shown in Figure S2a. The baseline parameter B_x accounted for residual activity even when the calcium signal is low or flat, as $\Delta F/F = 0$ does not necessarily imply zero firing due to background fluorescence and imaging noise. We therefore set $B_x = 20\text{Hz}$ to reflect plausible baseline activity. To assess the robustness of our results to the gain parameters A_x , we conducted a sensitivity analysis by decreasing each of A_1 , A_2 and A_3 by 20%, one at a time. The changes in well-fit downstream neurons' preferred orientation (measured by preference angle difference) remained small across all conditions (Figure S1a), and the qualitative tuning properties were preserved. Similarly, we plotted the orientation maps of where A_1 , A_2 and A_3 parameters were varied. We observed minimal changes to the orientation selectivity (Figure S2c, S2d, S2e).

To ensure that the simulated stimulus was able to elicit a stimulus-consistent pattern among lamina, which is convinced to be a direct downstream of the photoreceptors, we examined the spatial and temporal activation patterns of lamina cells L1-L3 and their downstream neurons. Since our stimuli (OFF-bars) presented a straight bar pattern, we examined the firing rate of lamina cells when the OFF-bars were at different distances from them. We investigated the firing rate of lamina cells in response to the OFF-bars positioned at varying distances from them. Our findings revealed a clear trend: as the distance between the OFF-bars and the cell increased, the firing rate decreased correspondingly. Notably, the lamina cells (L1-L3) exhibited a peak firing rate when the OFF-bars were precisely aligned with the cell's location (Figure S2b). These results confirmed that artificial stimulus-driven lamina neuron activity in our digital *Drosophila* brain was consistent with the L1-L3 response properties known from biological experiments [29].

We then sampled a set of downstream neurons - specifically Tm1, Tm2, Tm4, Mi1, Mi4 and Tm3, which are known from anatomical studies to be postsynaptic targets of L1-L3 neurons [36]. These neuron types were not shown in Figure 1b to maintain visual simplicity, but their connectivity is well established. We randomly selected a single neuron from each of the Tm1, Tm2 and Tm4 subtypes as they are known OFF-cells [36] and evaluated if they exhibited firing patterns that reassemble the expected response profile to OFF-bar stimuli - that is, peaking when the bar aligns with their receptive field and decreasing with distance. As shown in Figure S2f, these OFF-responsive downstream neurons exhibited patterns as expected. In contrast, we also sampled a single neuron from each of the Mi1, Mi4 and Tm3 subtypes, known ON-cells [36], and evaluated whether they showed the opposite trend - stronger inhibition, when the OFF-bar was nearby (Figure S2g). These neurons' spatial activation pattern mirrors the OFF-bar's location on the compound eye, consistent with known visual topology [36].

B.3 Circular Gaussian function

We calculated the preferred orientations of each neuron using a circular Gaussian function defined as:

$$\text{FiringRate}_n(\theta) = C \cdot \exp\left(-\frac{d(\theta, A)^2}{B}\right) + D, \quad (5)$$

where:

- $\text{FiringRate}_n(\theta)$: predicted firing rate of neuron n at the orientation θ
- θ : input orientation (in degrees, $0^\circ \leq \theta < 180^\circ$)
- A : preferred orientation (in degrees, fit parameters)
- B : width parameters (sharpness of tuning)
- C : amplitude of the tuning curve
- D : baseline firing rate
- $d(\theta, A) = \min(|\theta - A|, 180 - |\theta - A|)$: smallest circular distance between input orientation and preferred orientation.

B.4 Circular standard deviation

We quantified the sharpness of orientation tuning within each column using the circular standard deviation (see Figure 5c) in the `scipy` Python library, defined as:

$$s = \sqrt{-2 \ln R} \quad (6)$$

Table S1: Ablation study results.

Z	Tm silenced	Mi silenced
Dm	90.7%	2.4%
Pm	10.4%	3.9%

where:

- s : circular standard deviation (in radians)
- R : mean resultant length, computed as:

$$R = \frac{1}{n} \sqrt{\left(\sum_{i=1}^n \cos \alpha_i\right)^2 + \left(\sum_{i=1}^n \sin \alpha_i\right)^2} \quad (7)$$

- α_i : preferred orientation of neuron i (in radians)
- n : number of neurons in the column

This metric captured how tightly clustered the orientations are around the mean. A value of $s = 0$ indicates perfect alignment, with larger values reflecting broader tuning distributions.

C Targeted ablation analysis

To evaluate the contribution of specific upstream visual pathways to orientation selectivity, we performed targeted silencing of key input neuron populations. Using the equation: $Z\% = X/Y$, where X represents the number of good fits in ablation, Y represents the number of good fits in the original result, and Z is the percentage of neurons remaining that still exhibit orientation selectivity. Silencing Tm neurons resulted in an 89.6% reduction in well-fit orientation-selective neurons in the Pm layer but only a 9.3% reduction in Dm. In contrast, silencing Mi neurons eliminated tuning across both layers, leaving fewer than 5% of well-fit neurons in either Dm or Pm (Table S1). These findings suggest that Mi neurons represent essential upstream sources of orientation selectivity throughout the optic lobe.

D Orientation-based synaptic connectivity structure

To examine whether recurrent connections in the connectome reflect functional similarity, we quantified synaptic connectivity as a function of orientation preference. Each synapse (p_i, q_i) was represented as a point $(\theta_{p_i}, \theta_{q_i})$, where θ denotes the preferred orientation (in degrees) of the pre- and postsynaptic neurons, respectively. In the excitatory subnetwork, connections cluster along the diagonal ($\theta_{p_i} \approx \theta_{q_i}$), indicating that neurons with similar orientation preferences are more likely to be recurrently connected (Figure S3a). This pattern is consistent with topographic organisation and colinear facilitation as proposed by Seung [45], supporting the notion of like-to-like connectivity in local circuits. In contrast, inhibitory connections exhibit off-diagonal structure, suggesting cross-orientation interactions and selective suppression mechanisms (Figure S3b). These results suggest that recurrent circuits in the *Drosophila* optic lobe preferentially link neurons with similar selectivity, providing network-level evidence for structured functional organisation within the connectome.

E Extended figures

Figure S4 shows the Gaussian model fitting used to quantify orientation selectivity. Figure S5 shows the orientation tuning profile of T4 and T5 subtypes. Figure S6 presents additional visualisation of orientation maps across different neuron types and layers within the optic lobe.

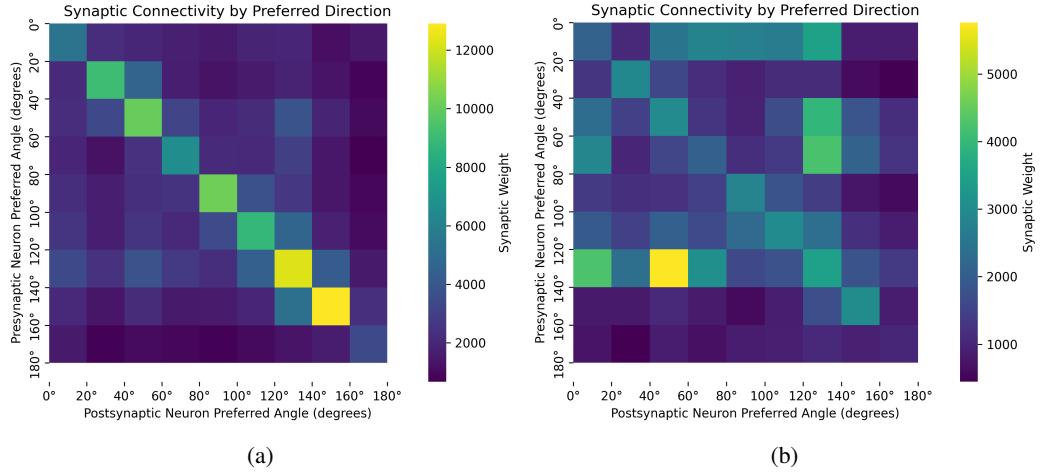


Figure S3: **Orientation-specific synaptic connectivity.** (a) Excitatory network showing strong diagonal structure, indicating neurons with similar orientation preferences preferentially connect. (b) Inhibitory network showing off-diagonal structure, consistent with cross-orientation suppression motifs.

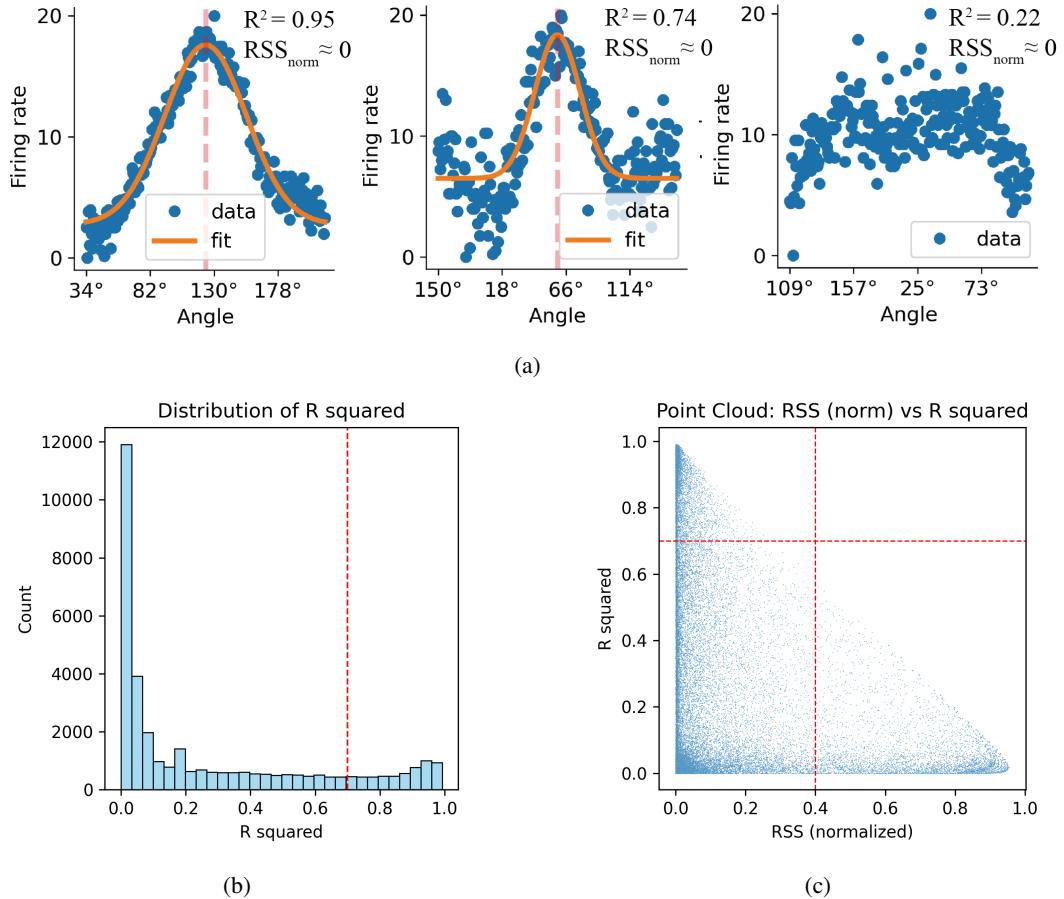


Figure S4: **Model Gaussian fit quality.** (a) Representative neurons where their Gaussian fit is considered good (left), near borderline (middle) and poor (right). (b) Distribution of R^2 values, where the red line indicates the threshold. (c) Distribution of R^2 and normalised RSS_{norm} of all neurons in the right optic lobe, including neurons in Figure S4a.

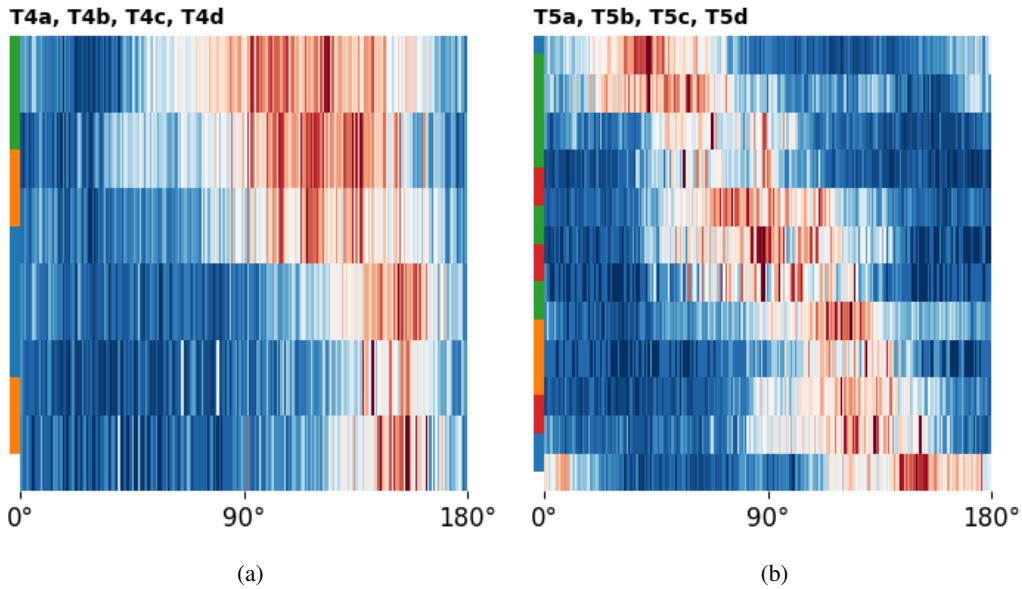


Figure S5: **Orientation tuning profiles of some T4 and T5 neurons.** (a) Heatmap showing the orientation tuning responses of T4 neurons that are well-fitted. (b) Heatmap showing the orientation tuning responses of T5 that are well-fitted.

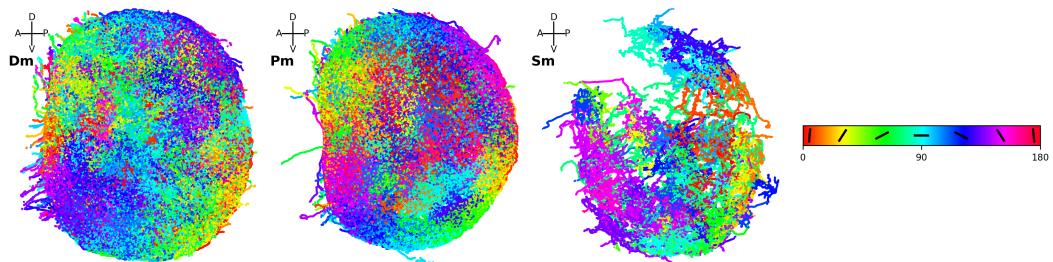


Figure S6: **Orientation preference structure across different neuron types.** Preferred orientation maps across optic lobe layers before smoothing, colour-coded by angle.