The *Drosophila* visual system: a super-efficient encoder

Keivan Razban

A thesis submitted for the degree of Doctor of Philosophy Department

of Biomedical Science The University of Sheffield

October 2022

Acknowledgement

Sheffield. What a journey it has been. If someone had told me five years ago that I would go to the North of England to dissect flies and model their brains; I would have probably looked at them, smiled, and finished my pain au chocolat, ignoring what I had just heard. But here I am, now, happy to have lived such a fantastic experience in this beautiful city, proud of the work I have done, and lucky to have met so many amazing people during these four years of great science. These people, I would not dare to thank in person, so I will do it on this page, hidden inside this long thesis that I hope they will all read thoroughly.

To start with, I would like to thank a person who I first met out of the blue, through Skype, on a Sunday morning, while my mum was cooking fried eggs in the background. His love for science, motivation, energy and happiness inspired me a lot. Thanks Mikko for this great time in Sheffield as well as in China, I would not have been where I am now without you!

Thank you Katie for being such a good friend, all the gossips we had inspired me a lot for my PhD. I am so glad I met you and I know our friendship will last forever.

Being a PhD student can be hard sometimes because of the lack of social interactions. I was lucky to work in the same lab as James, Katie, Neveen, Joni, Ben, Anthi, Gregor, Tom, Jouni, Hadi, Yunqi and Yifan. Pranking them, playing ukulele together, and Nerf gun shooting each other made the workplace a beautiful ecosystem.

Laura and Dhial, the random walks to the supermarkets and the UNO games are crucial moments that I will never forget. Thank you both for being such good and reliable friends. See ya!

My last flatmates in Sheffield were maybe odd people, but it always felt like a second family and I always felt at home living with them. Thank you James and Emma for being so easy to live with.

Thank you Max for those unforgettable moments spent together at Coffee Rev's, those football passes we did in front of Arts tower, and more…

Thank you Adriano for introducing me to ice swimming, for the beers we had at the jetty, and for those chess games you lost badly at. Keep swimming!

Thank you my dear Clement for visiting me six times from Paris! Not many people would do such long trips to see a friend, and it was always worth it, every time a new story, new encounters and new jokes.

I cannot but be grateful to you, Léïla, for your attention, support and love during my last months of writing-up. It really meant a lot to me.

Thank you Jean Valjean for being yourself. You deserve the bread you "stole" and, despite the rumours, I believe that you are a good and loyal person. I am sure our brotherhood will thrive whatsoever.

Even though, when visiting me in Sheffield, my mum broke her ankle, my sister was ill, and my dad's flight got cancelled because of some peculiar pandemic, I would not have gone through all the work done without my family. Everything that I am, I owe it to them first. They supported me all the way through it and helped me when my knee broke into pieces. They have *shown that* (choundeh) I can really count on them at any time of life.

Finally, thank you Priam for coming into this world, and especially during the toughest moment of my thesis. I cannot wait to explain this thesis to you when you're older! But for now, you can have a look at those colourful figures I made especially for you!

Abstract

In order to survive and reproduce, every animal needs to run accurate and diverse visual processes efficiently. However, understanding how they see is limited by our lack of insight into how evolution optimises resource- and area-constrained neural machinery.

It was shown recently in *Drosophila* that its photoreceptor cells, corresponding to individual "pixels" of the scene, react photomechanically to these light changes by generating an ultrafast counter-motion, a photoreceptor microsaccade. Each photoreceptor moves in a specific direction at its particular location inside the compound eye, transiently readjusting its own light input. These mirror-symmetrically opposing microsaccades cause small timing differences in the eye and the brain networks' electrical signals, rapidly and accurately informing the fly of the 3D world structure. Remarkably, it has been shown that the *Drosophila* can resolve angles finer than 1°, five times less than what the optic laws would predict in a static fly eye.

The results presented in this thesis demonstrate that hyperacute visual information is transmitted from the photoreceptors to the visual pathway and I report a deep learning approach for discovering how the *Drosophila* compound eyes' biological neural network (BNN) samples and represents hyperacute stimuli.

Using *in vivo* two-photon calcium imaging on a transgenic fly, I recorded the responses of 17 flies' L2 neurons, OFF neurons in the early visual pathway, while presenting fine resolution visual patterns. I showed that the *Drosophila*'s visual hyperacute information is transmitted from the photoreceptors to the medulla layer (2^{nd} layer in the visual system). Additionally, I found that the L2 neurons show direction-specific acuity and proved that this is a consequence of the photoreceptors' microsaccades.

Next, I show that an artificial neural network (ANN), with precisely-positioned and photomechanically-moving photoreceptors, shaping and feeding visual information to a lifelike-wired neuropile, learns to reproduce natural response dynamics. Remarkably, this ANN predicts realistic stimulus-locked responses and synaptic connection weights at each eye location, mapping the eyes' experimentally verified hyperacute orientation sensitivity. By systematically altering sampling dynamics and connections, I further show that without the realistic orientation-tuned photoreceptor microsaccades and connectome, performance falters to suboptimal. My results demonstrate the importance of precise microsaccades and connectivity for efficient visual encoding and highlight the effect of morphodynamic information sampling on accurate perception.

Table of Contents

List of figures

L2 neurons - drawn by Katie Greenin-Whitehead

Chapter 1: General introduction

1.1 Motivation and objectives

Human curiosity leads us to try to understand better the beauty of the world we live in. Its harmony is present from atomic structure to planetary motion through the mystery of living organisms. This thesis deals with a small branch of this complex sophistication: brain cognition and, more specifically, visual encoding. An animal's perception of its surrounding world is an interpretation of its brain that receives limited sensory inputs. Yet, this interpretation must be accurate enough to survive and reproduce. Evolution presumably organised and tuned the animal brains to find a compromise between their world interpretation accuracy and energy consumption. Using the well-studied and highly tractable genetic model *Drosophila melanogaster*, we can investigate how a brain perceives the visual world. Deciphering its brain's encoding mechanisms gives us insight into how animals perceive the visual world.

It was long thought that insect photoreceptors remain immobile (static) as they sample visual information (Land, 1997). However, recent studies using *Drosophila* have shown that photoreceptors are not static. Instead, light intensity changes make them move dynamically, enhancing a fly's visual capabilities beyond the optical limitations of its compound eyes (Juusola et al., 2017). These ultrafast photomechanical photoreceptor microsaccades enable *Drosophila* to see the world in hyperacute detail (Juusola et al., 2017). These results demonstrated that the fly brain's neural algorithms to perceive the world differ from the general idea. Moreover, parts of the algorithm enhancing its hyperacuity remain unknown.

In this thesis, I aim to show that the *Drosophila* visual pathway encodes the world neurally in high resolution and that the morpho-dynamism discovered at the photoreceptor level plays an important role in this efficient and powerful encoding. Furthermore, to help the science community regarding fly vision analysis, I present the custom-written stimuli and analysis algorithms, as well as the experimental techniques, all described in detail in each Chapter's material and methods section. Finally, I hope the model built and described in Chapter 3 will give insight into computer vision algorithms in general and robotic paradigms.

My PhD study consisted of two main parts, each of which took about two years. The first part entailed intense experimental learning and recordings, and response analysis. During

the second half of my PhD, I profited from my data science background/knowledge to model the responses recorded during the first two years.

1.2 Outline of the thesis

This thesis presents findings, reviews and discussions about how the visual world is perceived and encoded in the fly visual pathway. It consists of a general introduction chapter, followed by four result chapters in which specific questions are addressed, methods are detailed, and findings are discussed.

In the remaining, I summarise each result chapter and provide a brief literature review to contextualise and present relevant information for the performed research.

In Chapter 2, I measure and analyse if and how the *Drosophila* visual hyperacuity, as discovered previously at the photoreceptor level (Juusola et al., 2017), is maintained at the first neuropile, lamina, and, specifically, at the input layer of the second neuropile, medulla. I performed two-photon Ca^{2+}-imaging in transgenic flies expressing GCaMP6f-calcium-indicator in the first visual interneurons, L2 monopolar cells (L2 LMCs). The output of these neurons emphasises environmental dark-contrast changes. L2 axon terminals' calcium activity was recorded in the medulla input layer-2 while visually stimulating the fly with ultra-fine moving patterns. The results revealed how visual hyperacuity was maintained in the lamina and medulla neural responses. Additionally, L2 medulla terminals showed orientation sensitivity correlating with the corresponding photoreceptor microsaccade directions. These results indicate that microsaccades play an important role in the high-resolution visual encoding in the fly brain.

Chapter 3 describes a model of the *Drosophila* visual pathway's first two processing layers, which uniquely includes a realistic, dynamically-operating photoreceptor microsaccade sampling matrix (made by Dr Jouni Takalo, a senior researcher in our laboratory). This combined model is a connectome-constrained artificial neural network trained on the recorded L2 responses (Chapter 2) to fit them. The model highlights the importance of specific bio-physical properties and synaptic interactions of the fly visual pathway, such as the saccades and the local circuitry, for accurate neural network functions and predicting individual neural responses in the lamina and medulla input layers (at the L2 axon terminals). Furthermore, it is possible that this model can be used as an encoder for computer vision paradigms that are similar to the tasks solved by flying animals, such as visual processing in drones or general high-resolution video compression.

Chapter 4 attempts to extend the findings in the previous chapters to the whole visual pathway, using a pan-neuronal GCaMP6f-expression in the *Drosophila* brain. Visual stimulation with ultra-fine moving patterns while two-photon calcium imaging neural activity revealed hyper-acuity at every recorded layer of the visual pathway, which in these recordings encompasses at least the complete lamina and medulla network. Automatic response-clustering showed differences in directional and orientation sensitivities between the lamina and medulla neurons. These results open the door for discussions about the possible mechanisms enabling the most efficient and high-resolution visual encoding in the fly brain. In addition, this Chapter brings about ideas on tackling the challenging issues of pan-neuronal response analysis.

1.3 The visual pathway of the *Drosophila*

To study and make sense of how neural activity represents visual stimuli in the fruit fly brain, one needs to understand how its visual pathway is structured. An accurate connectivity map of its different processing layers enables one to choose where to study visual information processing, identify the neurons, and interpret their activity *a posteriori*. Many animals' visual pathways, particularly those of genetic model organisms, such as *Drosophila*, have been precisely dissected, and a wealth of published details are available. As shown below, these connectomes provide useful 3D circuit diagrams of where light information is converted to neural signals and how such neural activity may drive visual behaviours.

Moreover, modern genetics has enabled detailed investigations into synaptic connectivity and its role in local circuit computations. For example, neuronal photoactivation allows desired input from a group of neurons so that an investigator can identify their axonal branching and pathways to route information (Sun et al., 2017a).

Three types of fly eyes sample light coming from the surrounding environment. Each of the two principal eyes, the left and right compound eye, is made of about 800 small lens-capped modules (ommatidia) that contain the light sensors, the retinal photoreceptor cells. Together, the ommatidia map a panoramic visual field covering almost every direction around the fly's head. Four synaptic processing layers, called optic neuropils, further process information captured by the photoreceptors: the lamina, medulla, lobula and lobula plate (**Figure 1**) (K. F. Fischbach & Dittrich, 1989; S. Y. Takemura, 2015). As in the mammalian brain, the fly visual pathway is hierarchically layered: information is first processed locally before being integrated more complexly as it goes through each layer (A. Borst & Egelhaaf, 1990; Reichardt, 1987). The fly visual pathway is also retinotopically structured (Braitenberg, 1967): each point in the visual space stimulates only a column (*aka* module or cartridge) of neurons. Such a module is repeated across the eye, and electron microscopy cross sections have enabled detailed connectome reconstructions for the first layers.
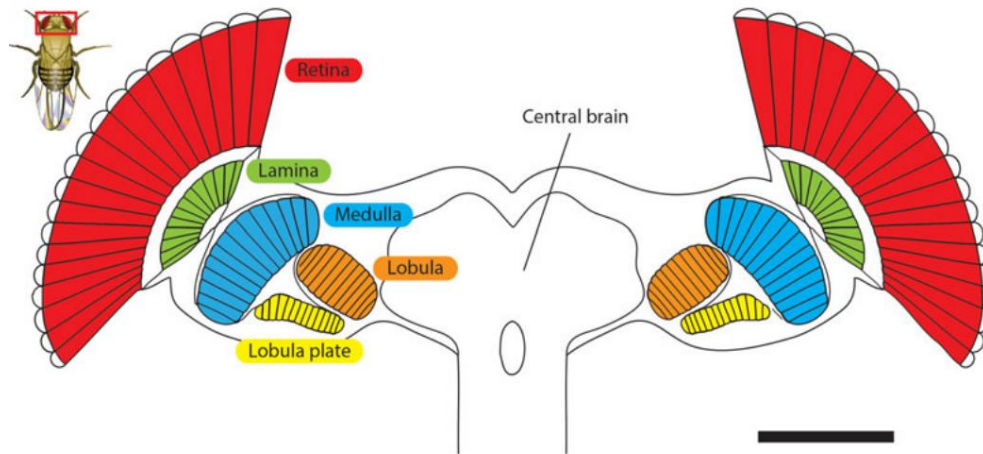
Figure 1: Diagram of a horizontal section of the fly's brain and visual system.

A fly's optic lobe consists of four retinotopically organised neuropils: the lamina, medulla, lobula and lobula plate. Scale bar = 100 µm. Image from (S. Y. Takemura, 2015).

The first layer, made of photoreceptors, samples light through their light-sensitive protein called rhodopsin. The outer R1-R6 photoreceptors (**Figure 2**) express green-sensitive rhodopsin (Rh1) and UV-sensitive 'antenna pigment'. The inner R7-R8 photoreceptors express slightly different rhodopsin variants, responding to yellow, pale and can be sensitive to light polarization (Clandinin & Zipursky, 2002; Wardill et al., 2012; Wernet et al., 2003).
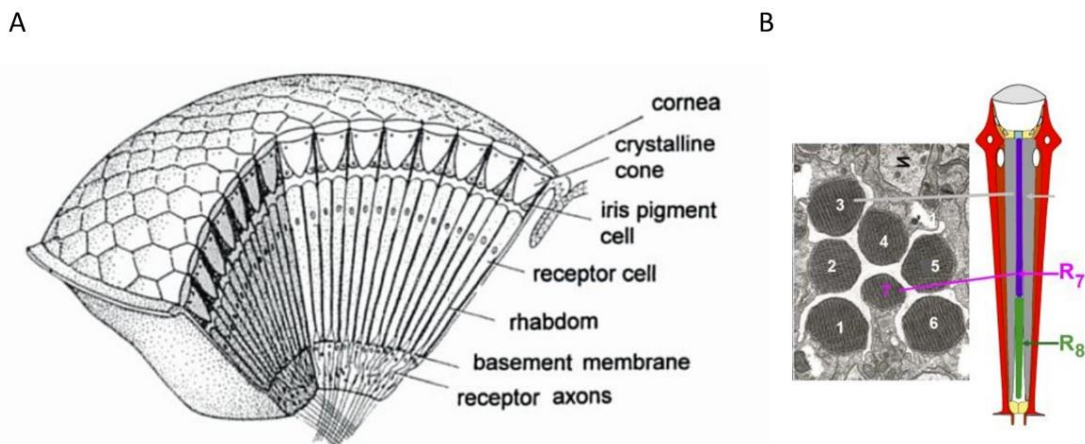
Figure 2: Fly eye description.

A) Schematic structure of the compound eye. Figure from (Land and Nilsson, 2002). **(B)** Electron micrograph (left) and schematic structure (right) of an ommatidium in the fly eye. Inside an ommatidium under a faceted lens, eight photoreceptors (R1-R8) align in seven different directions. R7 and R8 photoreceptors lie on top of each other in the centre and share the same axis Images from (Elyada, 2009).

The photoreceptors connect to the second layer following the neural superposition principle (Braitenberg, 1967; K. Kirschfeld & Franceschini, 1968). This arrangement means that the eight photoreceptor cells (R1-R8) inside each ommatidium share the same facet lens (**Figure 2**A) but point in seven different directions (**Figure 2**B; R7 and R8 share the same receptive field). On the other hand, neural signals of eight different photoreceptors (R1, R2, R3, R4, R5, R6 and R/8), belonging to the seven neighbouring ommatidia looking at the same point in space, are pooled together at one neural cartridge in the next two neuropils: the lamina and medulla (**Figure 3**). While the six outer photoreceptors R1-R6 project their axon terminals to neural columns in the lamina (**Figure 3**), R7 and R8 cells bypass this layer and make synaptic contacts with their corresponding medulla column (Karl Friedrich Fischbach & Hiesinger, 2008; Kuno Kirschfeld, 1967; Morante & Desplan, 2005).
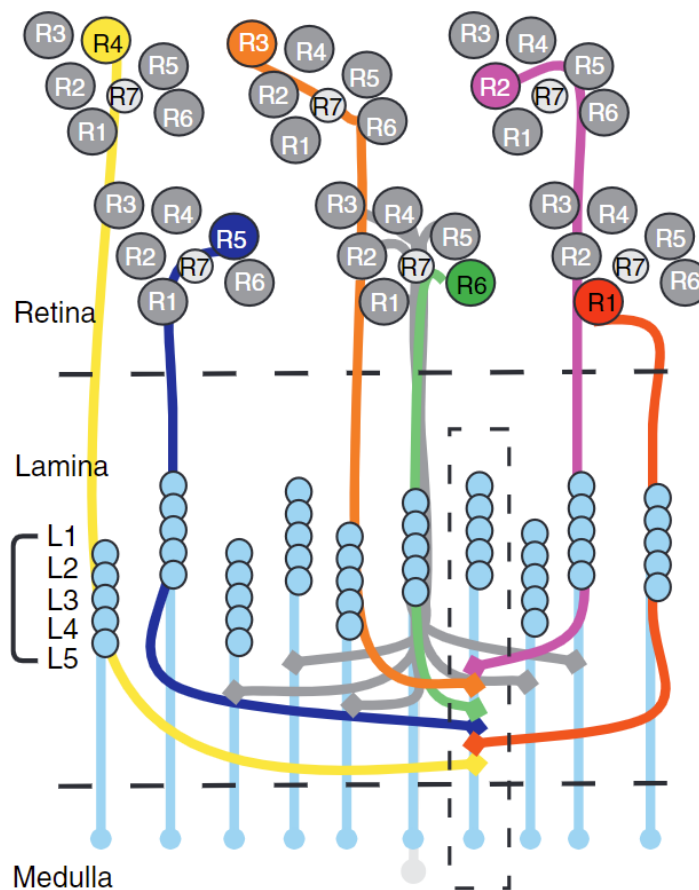
Figure 3: Wiring diagram of the retinotopic mapping in the fly lamina.
Six outer photoreceptors R1-R6 - belonging to six neighbouring ommatidia but point in the same direction - send their neural signal downstream to the same lamina column. Figure from Morante and Desplan, 2005.

Neural superposition wiring may benefit the fly visual system in different ways. For example, single photoreceptor light stimulation can evoke unstable oscillations at the postsynaptic target neurons, the large monopolar cells (LMCs), to the detriment of transmission (van Hateren, 1987). In contrast, having six simultaneous inputs from the superimposed R1-R6s of the neighbouring ommatidia deliver a more stable response (van Hateren, 1987). The results further suggested that this layout, together with synaptic feedbacks (Dau et al., 2016; Nikolaev et al., 2009; Zheng et al., 2006), may prevent over-excitation of the LMCs, normalising transmission in variable conditions. Furthermore, intracellular R1-R6 and LMC recordings suggest that presynaptic membrane properties (Weckstrom et al., 1992) and synaptic feedback interactions accelerate and accentuate synaptically transmitted information about light contrast changes through balancing the synaptic loads (Dau et al., 2016; Nikolaev et al., 2009; Zheng et al., 2006). Besides these advantages, I will show below that neural superposition is also involved in enhancing visual acuity.

The outer R1-R6 photoreceptors feed motion information to the LMCs (L1-L3), whilst the inner R7-R8 photoreceptors innervate specific medulla neurons, subserving colour vision (Heisenberg & Buchner, 1977; Pichaud et al., 1999; Troje, 1993; Yamaguchi et al., 2008). Interestingly, however, R7 and R8 also form gap junctions with R1 and R6 in the lamina, broadening the spectral range of motion vision (Wardill et al., 2012).

Serial electron microscopy has revealed the connections distributing photoreceptor information in the lamina, the most peripheral neuropil in the fly eye (Meinertzhagen and O'Neil, 1991; Rivera-Alba *et al.*, 2011). Each lamina cartridge contains twelve cell types: seven narrow-field elements (5 types of monopolar cell, L1–5; 2 centrifugal neurons C2, C3), two wide-field feedback neurons (Lawf1 and Lawf2), lamina intrinsic amacrine neurons (Lai) and the T1 basket cell and a wide-field intrinsic or amacrine cell. Laterally, L4 cells provide inputs from two neighbouring cartridges (posterodorsal and posteroventral), contributing to the local connection matrix (Meinertzhagen and O'Neil, 1991).

In the lamina, L1, L2 and L3 LMCs start segregation in the visual information processing through parallel ON and OFF pathways (Joesch et al., 2010b; Reiff et al., 2010). These pathways contain visual neurons in which responses selectively amplify bright or dark moving edges. Whilst the relations between ON and OFF pathways and the questions about their utility are still actively debated, the current consensus is that the correlation of signals coming from these two pathways with different temporal dynamics enables visual

motion detection (Behnia et al., 2014; Alexander Borst et al., 2020; Alexander Borst & Helmstaedter, 2015; Clark et al., 2011). Additionally, studies on these pathways show that their specific structural segregation is crucial for high-contrast stimuli but essential for low-contrast stimuli (Joesch et al., 2010a; Rister et al., 2007).

L1 and L2 are the two main inputs to the ON and OFF pathways, respectively. Targeted expression of genetically encoded calcium indicators has revealed these cell types do not carry information about motion direction yet, but rather respond to light increment/decrement respectively (Reiff et al., 2010). Furthermore, genetic manipulations have demonstrated that L2 neurons provide inputs to a pathway specialised for detecting moving dark edges (Clark et al., 2011; Joesch et al., 2010a; Rister et al., 2007). Finally, some studies on L2 neurons showed that they are inhibited laterally in an 'antagonist-centre-surrounding' way from their input photoreceptors (Dubs, 1982; Freifeld et al., 2013a; Järvilehto & Zettler, 1973; Laughlin & Hardie, 1978; van Hateren, 1992).

The subsequent processing layer, the medulla, also has a well-studied connectome (Campos-Ortega & Strausfeld, 1972; K. F. Fischbach & Dittrich, 1989; Rein et al., 2002; S.-Y. Takemura et al., 2008; S. Y. Takemura et al., 2013) and contains the T4 and T5 cells that receive inputs from the specific parallel ON and OFF pathways, respectively (Shinomiya et al., 2019). The current evidence suggests that these two neuron types (projecting axons onto Lobula Plate layers) are the first direction-selective neurons in *Drosophila*'s visual system (Schilling et al., 2019; Strother et al., 2017). Strikingly, they also respond strongly to static features orthogonal to their preferred axis of motion (Fisher et al., 2015).

Finally, in the fourth layer, the local motion stimuli of the previous layers are integrated and non-linearly filtered by lobula plate cells that are stimulated primarily by global motion stimuli (Barnhart et al., 2018; A. Borst & Haag, 2002; Joesch et al., 2008; Schnell et al., 2010).

To understand the physiology of the visual pathway, some researchers focused on the communication between neurons and the pathways they reveal. The primary inhibitory neurotransmitters in the visual pathway are histamine and gamma-aminobutyric acid (GABA) (Raghu et al., 2013), whilst glutamate, acetylcholine and serotonin provide the excitation in the network. Photoreceptors release histamine (R. C. Hardie, 1987) in the LMCs and amacrine cell synapses. GABA is a major inhibitory neurotransmitter widely present in the visual pathway, creating centre-surrounding receptive fields and separating

the ON and OFF pathways (Freifeld et al., 2013a). L1s release glutamate and L2s acetylcholine in their medulla terminals (S. Y. Takemura et al., 2011), and serotonin is used as a neuromodulator in the lamina (Gschweng et al., 2019). Moreover, to optimise the visual information flow, balancing the synaptic loads at the photoreceptor-LMC synapses, the major feedback connections to the photoreceptor axon terminal are excitatory (Nikolaev et al., 2009; Zheng et al., 2006).

1.4 Recording neural activity

The brain of every living organism comprises neurons that form networks via synapses and gap junctions. Their electrochemical communication is enabled by a controlled transfer of potassium, calcium, and sodium ions across their plasma membranes through specific ion channels. Hence, one can measure a neuron's membrane potential with intracellularly placed microelectrodes. It is also possible to record the neural activity by measuring the related ion concentration changes via specific genetically-encoded or pharmacologically-introduced indicators. These two main methods are electrophysiology and calcium imaging, respectively.

Although microelectrode penetrations invariably damage the cells, transiently changing the ion concentrations across its membrane, the cells typically recover within seconds as the membrane seals to the electrode. Electrophysiology is arguably the most precise recording method, enabling accurate recordings of ultrafast neural activity, including action potentials. However, visualising more than one neuron at the same time is difficult with this method.

Conversely, calcium imaging is an advantageous technique for investigating activity across neuron ensembles. It visualises Ca^{2+} concentration changes within (a) neuron(s) by using specific calcium indicator molecules that evoke concentration-dependent fluorescence changes. Hence, the brighter the image, the higher the local calcium concentration. So, more calcium ions are concentrated inside the observed intracellular cell compartment at a given time. Using well-established standard techniques, calcium indicators can be genetically expressed in many model organisms, such as fruit flies. The first genetically encoded calcium indicators (GECI), called Cameleons, were generated in 1997 (Miyawaki et al., 1997). Since then, this method has dramatically improved, and there are now tens of new GECIs available, differing in several aspects, such as colour, reactivity, or brightness. Review articles explain in detail how GECIs work and discuss a large panel

of the different GECIs used today, with their advantages and drawbacks (Grienberger & Konnerth, 2012).

Although synthetic dyes have better temporal resolution and kinetics, only GECIs enable targeting specific neurons by restricting their expression using the Gal4-UAS approach, a well-established genetic method for investigating gene expression and function in *Drosophila* (Brand & Perrimon, 1993).

One can visualise activity-dependent fluorescence in neurons using *charge-coupled devices* (CCD-cameras based upon complementary metal-oxide-semiconductors, CMOS) coupled to a fluorescence microscope. However, it is also possible to excite the GECI to generate fluorescence using a two-photon laser (using the 2-photon excitation principle (Wilson, 1985)), using x/y/z-laser scanning with highly-sensitive photomultiplier tubes (PMT) (Denk et al., 1990). This method is precise, and when the laser power is not too high, it does not damage the visualised cells.

Nowadays, genetic tools on *Drosophila* enable us to generate flies in which specific neurons express a particular gene. Moreover, this gene can be engineered or taken from another species (transgene). Hence, it is possible to have flies expressing GECI(s) almost wherever we want within the brain so that we can analyse only specific neurons and neural regions. Or, we can express the selected GECI in all neurons (pan-neural expression) (e.g. Russell, 2011).

Well-detailed protocols are openly available for two-photon calcium imaging in the *Drosophila* brain. These explain what transgenic mutant flies to choose, the preferred age of the fly, how to tether it properly, and the experimental conditions for robust and clean recordings (Chiappe & Jayaraman, 2012; Seelig et al., 2010)

When using two-photon calcium imaging to measure neural activity, one must ensure calcium transients correlate well with the neural responses. For example, it has been recently shown how serotonin is involved in calcium transient as a response to visual stimuli in L2 neurons (Gschweng et al., 2019). Although these results indicate that calcium indicators are a good approximation for neural activity measurements, novel genetically encoded voltage indicators appear to measure neural processing in vivo at a faster resolution (Yang et al., 2016). Hence, this method could be relevant to investigating fast responses induced by hyperacute visual stimuli.

The problem of experimental artefacts, such as muscle- or photomechanically-induced brain movements (Kemppainen et al., 2022) inside the head capsule, can be a real issue affecting calcium imaging. Indeed, it is difficult to prevent the fly brain from moving, with these movements resulting in fluorescence changes regardless of the local calcium ions concentration. Some specific preparation protocols have been developed to reduce these effects. One way to reduce brain movements is to cut the muscles inside the head capsule during dissection (Chiappe & Jayaraman, 2012; Seelig et al., 2010). Unfortunately, this procedure may eliminate essential aspects of the fly's active vision, potentially leading to massively underestimated signalling performance. Strikingly, activity-independent dyes are not usually used to monitor how the brain moves, although these would precisely highlight the presence of motion artefacts. For this purpose, two-colour two-photon calcium imaging could be used in the appropriate setup (having enough PMT channels). Such a system can also be used to visualise neural activity simultaneously in two different neural regions (Sun et al., 2017b)

Moreover, one can also correlate brain motion with other physiological processes, in which activity can be conveniently recorded, such as respiration or heartbeat (Paukert & Bergles, 2012). Various motion correction algorithms have been used to tackle these issues (Dombeck et al., 2010).

1.5 *Drosophila*'s visual acuity

1.5.1 Photoreceptor light activation

Phototransduction is the process in which a biochemical reaction cascade converts a photon's energy into a photoreceptor's electrical response. This cascade has been a topic of intense research for a long time (R. C. Hardie & Postma, 2008). A *Drosophila* R1-R6 photoreceptor's light-sensitive part, the rhabdomere, contains ~30,000 tightly packed tubular microvilli, each of which houses the whole phototransduction cascade reactions (**Figure 4**), characterised in detail by Hardie and Juusola, 2015. Each microvillus expresses light-sensitive proteins, rhodopsin, the most populous of the molecules involved in the transduction cascade. Once a photon activates a single rhodopsin molecule in a microvillus, it takes 8-30 ms (latency distribution) before the phototransduction cascade has fully amplified this signal into a small electrical response of variable size (amplitude distribution), called the quantum bump (Mikko Juusola & Hardie, 2001). These reactions also take a little time to recover in the light-activated microvillus. Thus, it cannot respond to another photon for about 20-300 ms, remaining in the state of refractoriness (refractory distribution) as part of light adaptation (M. Juusola et al., 1994; Song et al., 2012).

Meanwhile, quantum bumps from different microvilli sum up the photoreceptor's macroscopic voltage response to light contrast changes. Experiments and simulations indicate that evolution has tuned the microvilli numbers and their refractory photon sampling dynamics to visual behaviours, to be particularly effective in maximising information from bursty saccadic natural stimuli (Mikko Juusola et al., 2017; Song & Juusola, 2014).
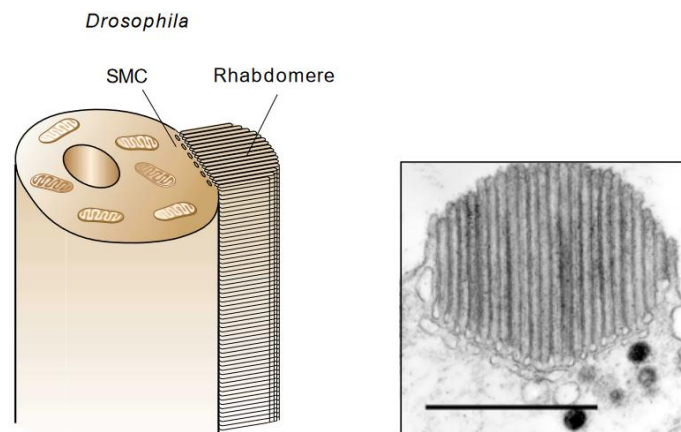
Figure 4: Photoreceptor structure.

In *Drosophila*, as in other dipteran photoreceptors, the photoreceptive membrane is organised into thousands of tightly packed, tubular microvilli. Each is 1–2 µm long and ~60 nm in diameter, forming a 100- µm -long rhabdomere. At the microvilli base, submicrovillar cisternae (SMC) systems are presumed to represent smooth endoplasmic reticulum Ca^{2+} stores endowed with $Ins(1,4,5)P_3$ receptors. However, the SMC may have a more important role in phosphoinositide turnover. The inset shows an electron micrograph of one rhabdomere (scale bar, 1 µm).Figure from Hardie and Raghu, 2001.

1.5.2 Old theory: a static eye for a poor resolution

Visual acuity is defined as the minimum angle that the eye can resolve. For a long time, because the compound eyes are an extension of the rigid cuticle, it was believed that their inner workings would also be static. Their visual acuity would then be limited by the photoreceptor spacing, as set by the interommatidial angle, and further affected by the photoreceptors' finite integration time and receptive field sizes (Mikko Juusola & French, 1997; Land, 1997; Srinivasan & Bernard, 1975). Hence, we could calculate an upper limit of the compound eye's spatial resolution from the following logical statements.

Since each ommatidium forms the primary sampling unit, the upper limit of a compound eye's spatial resolution is set by the density of its ommatidial array, the eye's 'pixel density' (Snyder et al., 1977). Suppose a regular pattern of black and white stripes is presented to the fly. According to this theory, and following the Nyquist–Shannon sampling theorem,

the maximum spatial frequency that the fly can resolve, νs, is achieved when one ommatidium points to a black stripe and its adjacent ommatidium points at the next white stripe (**Figure 5A**). Thus, the interommatidial angle $\Delta\varphi$ is the critical parameter in determining νs (Snyder et al., 1977). For compound eyes with a hexagonal layout, as in the case of most flies, the effective interommatidial angle $\Delta\varphi_e$ (**Figure 5B**) can be calculated by the following equation:

$$\Delta\phi_e = \cos(30°)\,\Delta\phi = \frac{\sqrt{3}}{2}\Delta\phi \qquad\qquad \text{eq. 1-1}$$

Thus, the upper limit of the fly eye's visual acuity is given by:

$$\alpha = 2\Delta\phi_e = \sqrt{3}\Delta\phi \qquad\qquad \text{eq. 1-2}$$
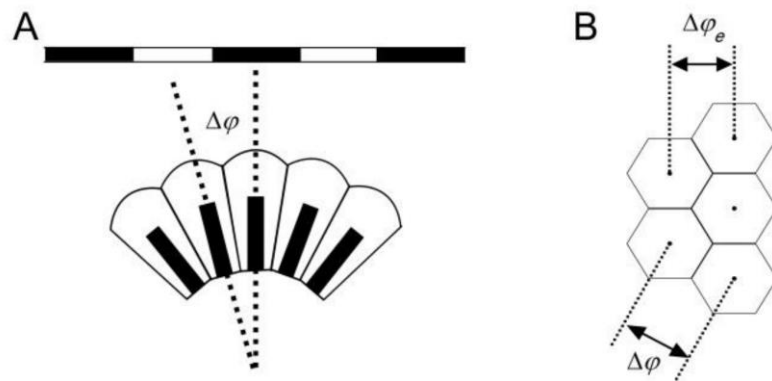
Figure 5: Interommatidial angle.

(A) The minimum angle a compound eye can resolve is its interommatidial angle. **(B)** The effective interommatidial angle of the fly eye with hexagonal layout, $\Delta\varphi e$, is smaller than its actual interommatidial angle $\Delta\varphi$. Their geometrical relation is described by equation (1-1). Images are redrawn from (Land, 1997)

On the other hand, whether this limit is achieved or not depends on the spatial performance of a single photoreceptor (Snyder et al., 1977). Several factors are considered in the sophisticated estimation of a photoreceptor's receptive field, which is quantified by its width at half-maximum, or acceptance angle $\Delta\rho$ (Warrant & McIntyre, 1993). Firstly, since the lens of ommatidium and rhabdomere of the photoreceptor are very small, optical quality is strongly affected by the diffraction of light, of which airy pattern is a function of light wavelength λ, lens diameter D, rhabdomere diameter, d, and focal distance, f. For theoretical estimation, the blurring functions at the ommatidium lens and rhabdomere tip

are broadly Gaussian and, therefore, can be combined to yield a simple approximation of $\Delta\rho$ (Snyder et al., 1977):

$$\Delta\rho = \sqrt{\left(\frac{\lambda}{D}\right)^2 + \left(\frac{d}{f}\right)^2} \qquad \text{eq. 1-3}$$

However, this formula could not be safely applied in many cases, owing to the complication of waveguide theory in small-diameter structures. (Stavenga, 2003; van Hateren, 1984) found that along the rhabdomere of fly photoreceptors, only a limited number of light patterns (modes) could be formed and that this number depends on the incident angle of light, leading to a smaller actual $\Delta\rho$ than what equation (1-3) implies.

Spatial cross-talk is another phenomenon to be considered: a photon escaping the rhabdomere first travels in and enters an adjacent rhabdomere (Warrant and McIntyre, 1991). Such an effect is likely to happen when the cross-talk index of the ommatidia/rhabdomere structure is less than three (Wijngaard & Stavenga, 1975). Indeed, this was found to be the case in *Drosophila* (Gonzalez-Bellido et al., 2011). Therefore, the static resolution of their neural images could be potentially lower than theoretically calculated.

1.5.3 New theory: A saccadic sampling for a high acuity

However, a new dynamic theory of visual acuity considers the *Drosophila* photoreceptors' photomechanical microsaccades. According to this theory, the microsaccades could improve visual acuity about four times over the static compound eyes' expected optical constraint, the interommatidial angle (Mikko Juusola et al., 2017). Optomotor behaviour, intracellularly recorded photoreceptor voltage responses and theoretical simulations matched the predictions of this theory. The phototransduction reactions (PIP_2 cleavage in and out of the rhabdomere membrane) directly cause the microsaccades, physically moving the photoreceptors whenever activated by photon flux changes (Roger C. Hardie & Franze, 2012; Mikko Juusola et al., 2017). The microsaccades' movement directions have been since mapped across the whole compound eye (Kemppainen *et al.*, 2022). They occur mirror-symmetrically in the left and right eyes and follow the optic flow of a forward-flying fly. Simultaneously, the light changes make the photoreceptor contract also axially, away and towards the ommatidium lens, narrowing and broadening their receptive

fields (Kemppainen *et al.*, 2022). Solid behavioural evidence, lamina 2-photon imaging and modelling effectively prove that microsaccades enable hyperacute-3D vision.

Moreover, a theoretical analysis using machine learning implies that *Drosophila*'s hyperacuity is necessary and sufficient to visually distinguish its conspecifics (Schneider et al., 2018). Their calculations show that hyperacuity, combined with current convolutional neural networks similar to the *Drosophila* visual system, allows us to accurately re-identify other flies.

Visual acuity has been estimated in different insects using various combinations of experimental and theoretical approximations. These studies' general trends indicate that compound eye architecture has evolved to process visual information according to the species' ecological demands (Gonzalez-Bellido et al., 2011; Land, 1997).

In Queensland fruit flies (Lawson & Srinivasan, 2020), optomotor behaviour to sinusoidal gratings in a flight simulator revealed that they have an acuity of about 1.9° inter-bar distance, from which they inferred the acceptance angle. They also inferred the interommatidial distance using the aliasing angle (grating wavelength that induces a reverse in the optomotor responses). However, those values only match the *Drosophila* eye optics if photoreceptors' microsaccadic receptive field movements are considered.

Classic studies in *Locusta* and two species of *Diptera* showed visual acuities around 0.3°. These estimates are 3-to-4-times smaller than the static eye optics' prediction (Burtt & Catton, 1954).

Furthermore, higher-order neurons such as LC18 in the Lobula columnar neurons can detect objects smaller than the lens resolution (Klapoetke et al., 2022). However, theoretically, this would not necessarily mean that the photoreceptors' resolution is smaller than the lens resolution because integrating visual information, if processed layer by layer, could potentially improve acuity through some diffraction phenomenon (E. T. Burtt and W. T. Catton, 1962).

1.6 Modelling neural activity in the fly brain

Behavioural experiments, calcium imaging or electrophysiology in living organisms' brains have enabled researchers to investigate neural activity and synaptic connectivity. To understand the findings more thoroughly, it is usually relevant to make a model based on the data and predict future data. The resulting models not only give a better insight of visualised behaviours or dynamics but are sometimes a way to answer questions when the modelled behaviour was able to solve them without us understanding why.

As an example of computer vision, many models for motion predictions have been made using information about the *Drosophila* visual pathway. Its characteristic ON/OFF pathway structure was used in some models such as the Hassenstein-Reichardt Correlator (HRC) (Clark et al., 2011) or non-linear enhancers/suppressors microcircuits (Ramos-Traslosheros et al., 2018), which can predict very well the behaviour of neurons in the medulla layer. However, HRC models for motion prediction struggle to predict T4-T5 responses to static orientations (Fisher et al., 2015). A possible reason for this is that such simple models lack some properties of the fly vision, such as microsaccadic eye movement.

One of the most recent models (Fu & Yue, 2020) successfully decodes local motion direction in front of a cluttered background. Moreover, the model reproduces the responses of all the layers, from the photoreceptors to the lobula plate layer. More specifically, they use, at the lamina layer, a variant of the 'difference of gaussian' spatial mechanism to enhance the motion edge selectivity and maximise the transmission of useful information from visually cluttered environments.

Another example in Artificial Intelligence (AI): electrophysiology in the cat cortical neurons has shown that some neurons respond to bar stimuli moving in a specific direction (Hubel & Wiesel, 1962). This experiment is one of the preliminary findings for hierarchical convolutional neural networks, widely used in artificial neural networks (ANN) (Y. Lecun et al., 2015).

This new era of AI has been used more and more in neuroscience, which inspired it. An ANN is a function approximation algorithm made of artificial neurons connected (or not) to each other as an oriented graph with weighted connections. The input of the model is carried by some neurons of the ANN. Each neuron sends its output to the neurons that are connected to them. The input of the intermediate neurons (those not connected to the input directly) is the weighted sum of all connected 'presynaptic' neurons passed through

a non-linear function (see **Figure 6**). To approximate a function, an ANN optimises some of its parameters (usually the weighted connections) using optimisation algorithms (usually derived from the gradient descent technique) (Amari, 1993). Thus, ANNs are made of units that individually have a simple rule (aggregating values and sending an output) but collectively can solve complex problems, i.e. converge to a highly intractable function (e.g. image recognition, speech generation). The similarity they have with the animal brain motivated researchers to infer biological properties of the brain by measuring the resulting parameters of trained ANNs, usually after biologically constraining them. In a sense, modelling the brain using ANNs is a way to focus attention on the combinatorial and distributed function of the brain rather than on a single neuron as an independent unit.
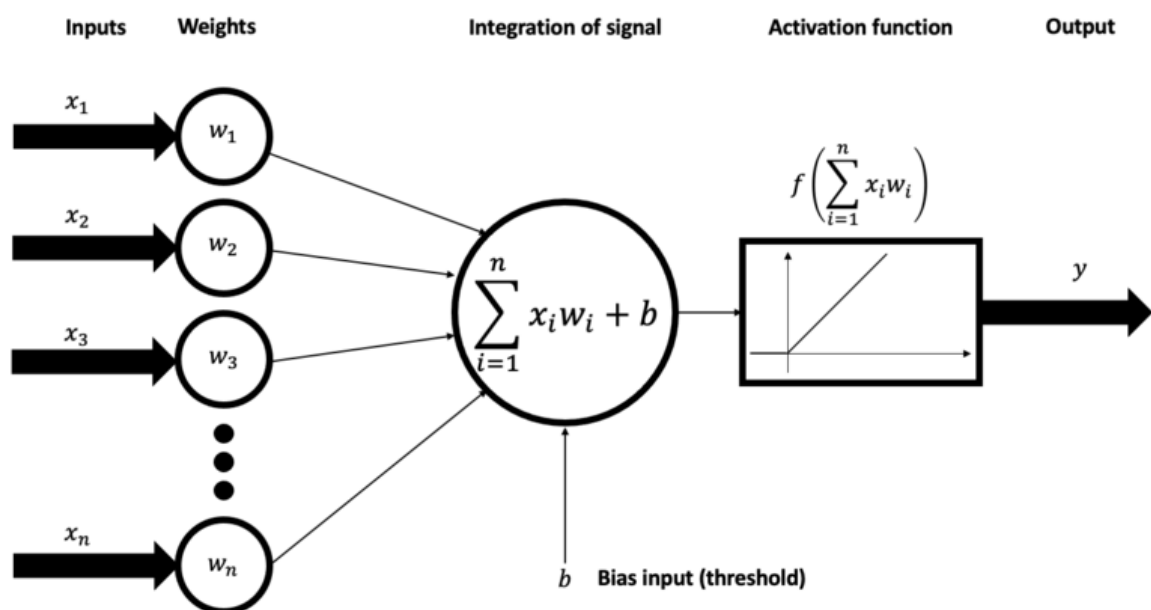
Figure 6: A neural network's unit.
In an artificial neural network, a neuron (centre) integrates the weighted outputs of the connected neurons (left), which, added to a bias, are passed through a non-linear function (activation function) to create the output.

These biological constraints enable shaping the ANN structure before training it. To do so, researchers use the connectome of the *Drosophila* brain, which, as discussed above, is well known. Many studies have described the wiring diagrams of different *Drosophila* brain centres, which show important properties that a biologically constrained ANN should have (Scheffer, 2020).

An anatomically-constrained ANN, trained to detect motion, learns some functional properties of the neurons in the *Drosophila* visual pathway, such as the ON and OFF pathways (Mano et al., 2021). Furthermore, biologically-constrained ANN training also revealed that the *Drosophila* visual pathway is robust to noise by artificially adding it to training visual scenes

As another example, training ANNs to reproduce the diversity of colour-coding in the bee ganglia showed that their presynaptic connections are likely to be randomly wired (MaBouDi et al., 2020).

Some models based on the lamina wiring diagrams in *Drosophila* (Rivera-Alba et al., 2011) have been proposed for natural scene object tracking. Studies testing the connectome-constrained convolutional networks against randomly-initialised convolutional networks showed that well-known neuronal output properties of medulla layers could not be discovered automatically if the network were initialised randomly. They concluded that the function of a biological neural network could be understood from its structure alone (Tschopp et al., 2018).

When the connectome is poorly known, or the parameters needed to describe it are not trivial, biologically-constraining ANNs become more difficult. Some researchers proposed a technique using Generative Adversarial Networks (GAN) to fit circuit models such as ANNs to recorded neural responses. GANs generate a circuit from latent variables (such as individual synaptic connections or single-cell biophysical properties) that are inferred by the adversarial training process (Creswell et al., 2018) instead of being given highly computationally expensive calculations (Arakaki et al., 2019).

Besides, training ANNs to solve a problem can be seen as a way to find the best combination of weights (parameters) to outputs; i.e., the best possible solution given the constraints (hyper-parameters). For example, Tseng *et al.*, 2021 found the best way to produce a high-quality image using ultra-thin lenses. We follow this idea in Chapter 4, as we aimed to reproduce hyperacuity at the Lamina level, given the complex biological constraints of the fly visual pathway.

Sometimes, it is impossible to do experiments in animals' brains at the level of single neurons. Such is the case for many cognitively more complex animals than *Drosophila*, to have neural activity modelled from limited electrophysiological or imaging recordings. Indeed, electrophysiology or calcium imaging on bigger brains enables monitoring of only

a small proportion of their neurons. These approaches give too little information if we believe the brain processes information as a vast interdependent and coordinated group of neurons, not as a stack of independent neurons. Moreover, because of technical and ethical reasons, most big animals are harder to modify genetically, needed for monitoring the brain activity at the level of single neurons. In these cases, people have used more macroscopical data to model the brain. For example, fMRI gives less accurate (Logothetis, 2008) but broader information about brain activity, yet it has enabled us to understand more sleep phases or emotions in the human brain (Greene et al., 2001; Redcay et al., 2007). Furthermore, animal behaviour data can indirectly enable brain modelling: MaBouDi, Marshall and Barron, 2020 showed that bees' cognitive strategy suggests particular central brain structures.

1.7 A virtual reality environment for the fly to enhance the responses

The animal preparation and its brain has to be healthy for one to record consistent neural activity. But it is also thought that the impression of being in a natural environment may enhance neural activity. For example, researchers found that the behavioural state of *Drosophila* modulates how sensory neurons (such as photoreceptors) and motor neurons are coupled (Ache et al., 2019), which implies that fly vision might be affected by the act of walking on a trackball.

More specifically, 2-photon calcium imaging showed that lobula plate tangential neurons respond with higher intensity when the fly is walking and that their peak intensity correlated to the walking speed (Chiappe et al., 2010). Additionally, whole-cell patch clamp electrophysiology and extracellular field potential recordings revealed that these neurons respond with a higher gain in the visual motion processing when the *Drosophila* is flying (Maimon et al., 2010; Tang & Juusola, 2010). The patch-clamp recordings further revealed some neuron types gated by flight behaviour through the octopamine neurotransmitter. Fujiwara and colleagues also showed that, during locomotion, global motion processing neurons in the fly brain responded even without visual stimuli. The same dynamic was seen in blind mutant flies (Fujiwara et al., 2017).

Additionally, Paulk and colleagues showed that a closed-loop system, which is more similar to a real (i.e., interactive) environment than an open-loop, also enhanced the

coherence between different brain regions in a fruit fly (Paulk et al., 2015). More precisely, they used electrophysiology on the optic lobe and central brain neurons in closed- and open-loop visual environments and compared the coherence for each result. It has also been shown (Taylor et al., 2015) that the responses of insects in a closed-loop environment were significantly different according to the realism of the stimulus feedback (fast feedback and accurate sensor). They compared two methods: one with a computer mouse sensor and the other with a computer vision algorithm (called 'FicTrac'). They also measured the general difference in the behaviour between both cases. Hence, several researchers have developed techniques to display good virtual reality stimulation for insect experiments (Lott et al., 2007; Takalo et al., 2012).

However, some findings suggest that behaviour-dependent neural activity is not always positively correlated. Indeed. in some parts of the fly brain, visual responses have been shown to diminish during flight (Seelig & Jayaraman, 2013).

The behavioural state has also been found to alter the baseline activity levels and the temporal tuning of neurons in the medulla layers (downstream to lamina) (Strother et al., 2018). These results show that fly visual acuity could depend on whether the fly is active or not.

Chapter 2: Hyperacuity in the medulla input layer

2.1 Introduction

Having discovered that *Drosophila* sees the world in super-resolution (Juusola et al., 2017), one could wonder if and how this super-resolution is neurally encoded to maintain an accurate perception of the visual world. Furthermore, it was recently found that the fly's photomechanical photoreceptor microsaccades make their receptive fields trace the direction of a forward-flying fly optic flow. Therefore, I ought to test directly whether such optic-flow-tuned microsaccadic sampling improves the acuity of moving stimuli directionally, as suggested experimentally (Kemppainen et al., 2022) and predicted theoretically (Juusola et al., 2017). To do so, I recorded neural responses of specific large monopolar cells (LMC), L2 interneurons, to a novel stimulus consisting of panoramic black-and-white dynamic-wavelength gratings, in which velocity and direction were changed systematically. These recordings were primarily done in so-called ultraviolet (UV) flies, using a bespoke two-photon Ca^{2+}-imaging system while presenting UV stimuli in an ultrafine spatiotemporal resolution to a fly walking on a trackball.

I first show that *Drosophila*'s hyperacute photoreceptor information is transferred to this visual brain layer. Additionally, I show that the L2 neuron terminals' visual acuity is orientation-specific. This orientation strongly correlates with the corresponding presynaptic photoreceptor micro-saccades orientations, showing the importance of morpho-dynamic sampling for acutely encoding the visual field. Finally, I found that the L2 neurons move similarly to the photoreceptors; these movements correlate with their activity.

2.2 Materials and methods

2.2.1 Two-photon system

I used a 2-photon resonance-scanning microscope (TrimScope, La Vision Biotech, Germany) with 40XW objective of numerical-aperture=1 (Zeiss C Achroplan 83 NIR 40x/0.8 w, ∞/0.17, Germany). The 2-photon excitation source was a mode-locked Ti:Sapphire Mai Tai SP Laser tuned to 920 nm. Fluorescence was collected by a photomultiplier (Hamamatsu H7422-40-LV, Japan) after bandpass filtering by a 525/50 nm emission filter. Images (approximately 150 x 1024 pixels) were acquired with

ImSpectorPro software (La Vision Biotech, Germany), typically 20-25 frames/s. Besides, when imaging smaller areas (e.g. 32 x 512 pixels), the used sampling rates were considerably higher (~50-200 frames/s). The laser intensity was kept below 240 mW (measured at the back aperture) to avoid heat-induced artefacts.

2.2.2 Fly genetics

The UV-flies were generated using rhodopsin *ninaE8,* also known as Rh1, with rescued UV-rhodopsin (Rh3) insertion. The *ninaE8* (ninaEP334) mutation reduces the expression of the rhodopsin *ninaE* to 0.0004% of wild-type levels (Wardill et al., 2012; Washburn & O'Tousa, 1989). This particular mutation was chosen as some level of expression of *ninaE* is required for normal rhabdomere development (Kumar & Ready, 1995). The UV-fly genotype used in 2-photon Ca^{2+}-imaging was UAS-GCaMP6f/CyO; L2-Gal4, UV/TM6B and UAS-GCaMP6; L2-Gal4, UV/TM6B. Origins of its different parts: R1-R6 photoreceptor UV-sensitivity resulted from P(Rh1:Rh3)[4303],*ninaE[8]*/TM6B, see supplementary material (Wardill et al., 2012). L2-Gal4 was 21D-Gal4, a gift from Martin Heisenberg (Washburn & O'Tousa, 1989). 21-Gal4 insertion was recombined to chromosome III together with the UV genetic set P(Rh1:Rh3)[4303],*ninaE[8],* using my UV-line stock and the 21D-Gal4 insertion line. The resulting lines were crossed to UAS-CD8-GFP and tested for GFP presence in L2 neurons using fluorescence microscopy. The presence of the UV genetic set was verified in positive lines by ERG testing for UV sensitivity (Wardill et al., 2012). UAS-GCaMP6f was BS46747 P[20xUAS-IVS-GCaMP6f] at P40 2L. Their eyes' structural integrity and photoreceptor microsaccade dynamics were within the normal range.

2.2.3 *Drosophila* preparation

2-to-4-day-old cold-anaesthetised flies (usually males) were prepared for the experiments much as described in (Chiappe & Jayaraman, 2012). A fly was waxed to a 0.001-inch-thick folded stainless steel shim holder, allowing access to the head's back through a 0.8 mm opening (**Figure 7B**). The head was tilted forward approximately 60°, exposing its back at the opening, and left the retina below the shim (**Figure 7A**). I cut a small hole at the back of the head cuticle with a fine tungsten needle and removed connective tissue, including the g trachea, to obtain optical access to the left and/or right medulla L2 axon terminals (**Figure 7A**). The fly was positioned over an air-suspended 6.13 mm ∅ polypropylene ball within the 2-photon imaging system, facing panoramic visual stimulation screens to enable motor activity recording (**Figure 7C**). Closed-loop temperature-controlled (25 °C) oxygenated fly ringer solution (containing mM: 120 NaCl, 5 KCl, 10 TES, 1.5 $CaCl_2$, 4

MgCl₂, and 30 sucrose) was perfused over the back of the head, keeping the preparation alive/healthy for hours-long experiments.

Figure 7: *in vivo Drosophila* preparation for 2-photon Ca²⁺-imaging.
(**A**) Using a bespoke stereomicroscope/preparation micromanipulation system, I fixed a fly in the predetermined position and orientation to a 0.001-inch-thick folded stainless steel shim (of a disk-shaped fly-holder) to access the back of the head through a 0.8 mm gap, comparable to (Chiappe et al., 2010). A small opening was cut at the head's back cuticle through an oxygenated fly ringer bath that covered the back of the head only, giving a visual view of the left or right medulla L2-terminals. (**B**) The fly's positioning inside the portable disk-shaped fly-holder. The fly-holder was transported to the 2-photon imaging system and rotationally adjusted by hand to centre the fly facing the panoramic visual stimulation (**Figure 8D**). Notice the semi-transparent beeswax droplet underneath the fly's eyes, immobilising its proboscis. (**C**) A tested fly could walk on a trackball during the 2-photon imaging of its L2-terminals' neural responses (Ca²⁺ fluorescence signals). In the experiments, the fly faced the visual stimulation screen inside a black-fabric chamber, which blocked outside light leakage and minimised scatter and internal reflections.

2.2.4 Visual stimulation

Figure 8 shows the method of how ultrafine, 7.5-11.25-times finer than the *Drosophila* eyes' ~4.5° interommatidial angle, video-stimulation was presented to *Drosophila* during their medulla L2-terminal's 2-photon Ca²⁺-imaging. I used a digital light projector (EKB DLP® LightCrafter™ Fiber-E4500MKII™ development module, EKB Technologies, Israel), equipped with a powerful 380 nm UV-LED, to provide 360 Hz UV-video stimulation with native 912 x 1140 pixel resolution to flies (**Figure 8A**). The UV-video images were projected on a back-projection (diffuser) screen. The whole system was inside a black, fluffy-fabric enclosed cage (**Figure 8B**) to block outside light and minimise internal reflections and scattering. Three short focal length achromatic doublet camera lenses (ref: MVL6WA, Thorlabs, USA) were then used to focus the projected images onto one end of three 7 x 7 mm coherent bundles of optical fibres (IB ASSY QA x 24", Schott, USA), with ~108 x 108 pixels (as the counted average) projecting onto each bundle (**Figure 8C**). The images were transmitted and magnified on the other end of the bundles by three optical tapers (Schott, USA). The tapers formed three Parafilm-capped panoramic fibre-optic screens (virtual reality stimulation screens) surrounding a tested fly frontally (**Figure 8D**).

Parafilm diffused light and damped reflections related to the numerical aperture of the taper/bundle fibres. The three fibre-optic screens accurately reproduced the video images into three angled vertical sections, positioned 38 mm from the fly eyes, filling large central parts of their left and right visual fields (total area: 135° x 45°). Therefore, with 108 x 324 pixels spread across the three screens, the angular resolution was ~0.6° at the point closest to the eyes and ~0.4° near the corners.
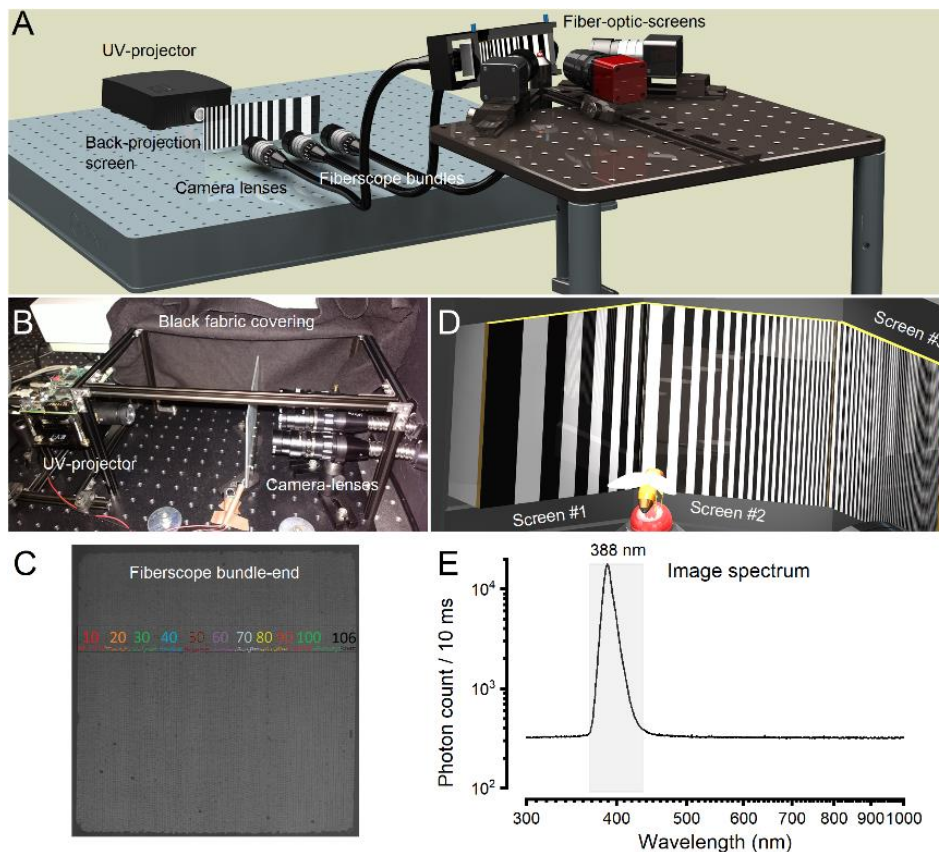
Figure 8: The bespoke high-resolution UV-video-display system (attached to the 2-photon imaging system) used for stimulating L2 neurons visually.
(A) The optical path, from the high-speed UV-projector to three high-resolution fibre-optic screens (taper ends), for presenting *Drosophila* with UV-video stimuli. **(B)** UV-stimuli were projected on a UV-preserving back-projection screen. Three camera lenses sampled the focused back-projected video images on three high-resolution ordered fiberscope bundles. This optical path was kept inside a light-proof cage (covered by a thick, fluffy black fabric) to minimise light scatter and internal reflections. **(C)** One fiberscope bundle end, with the highlighted fibre count for one of its rows. **(D)** The panoramic visual stimulation screen assembly was made of three high-resolution optical tapers (fibre-optic-screens), in which angles and position could be precisely and freely adjusted and fixed around the tested fly (by the instrument design). **(E)** The video-display system's spectral output as directly measured at the visual stimulation screen facing the tested *Drosophila*. The visual stimulation was dominated by UV light, peaking at 388 nm.

Visual stimuli were created using custom-written Matlab code, partly using the Psychophysics toolbox, in which the renderer updated images at 360 Hz, with a nominal 8-bits of DLP intensity at each pixel, and accurately projected them onto the three taper-

screens. Additional UV-band-pass filters (Semrock, USA) and adjustable apertures, interposed between the back-projection screen and the bundles, allowed me to cut-off long (non-UV) tail wavelengths of the images and adjust their overall intensities. The spectrum used in experiments is shown in (**Figure 8E**). By integrating it, one can estimate that R1-R6 photoreceptors that faced the optic taper screens were presented with 10^5-10^6 UV-photons/s, causing moderate-to-high light adaptation. Notice that because of the refractory photon sampling and intracellular pupil, which cause a dramatic drop in quantum efficiency (Hardcastle & Krapp, 2016; Juusola et al., 2017; Song & Juusola, 2014), most photons are lost during light adaptation. Consequently, an R1-R6 photoreceptor's effective photon absorption rate is actively maintained at ~1.5-8.0 x 10^5 to maximise its information transfer rate for high-contrast stimuli (Juusola et al., 2017).

2.2.5 Measuring L2-terminals' activity

The images about medulla L2-terminal fluorescence responses were analysed by custom-written python-scripts. The fluorescence intensity variations were quantified after background subtraction. Ca^{2+}-signal variations were obtained by subtracting the basal fluorescence, F_0, calculated as the mean intensity before the visual stimulation, from the observed intensity, F, ($\Delta F = F - F_0$) and giving this difference as the relative fluorescence change ($\Delta F / F_0$).

2.2.6 Measuring L2-terminals' resolvability

L2 neurons' medulla terminals respond strongly to light-OFF stimuli (Courgeon & Desplan, 2019; Freifeld et al., 2013). Therefore, a bright moving bar crossing an L2 neuron's receptive field (RF) evokes a transient response. Here, I used two types of moving stimuli to measure L2 speed and orientation sensitivity.

One stimulus type contained two parallel bars crossing an L2 neuron's RF. These bars induced a two-peaked change in the observed L2-terminal calcium fluorescence as a response. I can measure how well this intraneural calcium response resolved the two moving stimuli using the Rayleigh criterion:

$$R = \frac{P_{min} - T}{P_{max}}$$
eq. 2-1

, where T, P_{min} and P_{max} are the trough, smallest peak and highest peak, respectively.

We further measured single L2 neurons' resolvability to dynamically narrowing bar gratings (of continuously decreasing wavelength; **Figure 9A**) using a novel four-parameter bar-grating stimulus, as constructed in Matlab. The stimulus parameters were the speed, motion direction, initial wavelength and final wavelength (s, θ, λ_0 and λ_1, respectively). The inter-bar wavelength, which entered the tested *Drosophila*'s field of view, followed the geometric sequence update:

$$\lambda(t + dt) = (\frac{\lambda_1}{\lambda_0})^{\frac{1}{D}}\lambda(t) \qquad \text{eq. 2-2}$$

, where D was the duration of the stimulus (**Figure 9BC**). This way, the wavelength was divided by a constant factor, frame after frame, enabling an accurate estimate of the wavelength/time point when the L2 neuron could no longer resolve the adjacent moving bars. A more intuitive formula representing the wavelength over time is the following:

$$\lambda(t) = \lambda_0 \left(\frac{\lambda_1}{\lambda_0}\right)^{\frac{t}{D}} \qquad \text{eq. 2-3}$$

Importantly, this spatiotemporal stimulation enabled me to simultaneously monitor how the neighbouring L2-terminals encoded the same directional motion stimulation in different angular resolutions.

Like the moving two-bar stimulation, the dynamically narrowing bar grating stimulation induced a Ca^{2+}-fluorescence signal, showing a succession of peaks. To each pair of peaks, I can attribute a resolvability. Since this stimulus induces a response with a dynamic baseline (e.g., **Figure 10A** or **Figure 14C**), I applied the Rayleigh criterion to the relative peak heights:

$$R = \frac{P_{min} - T}{P_{max} - T} \qquad \text{eq. 2-4}$$
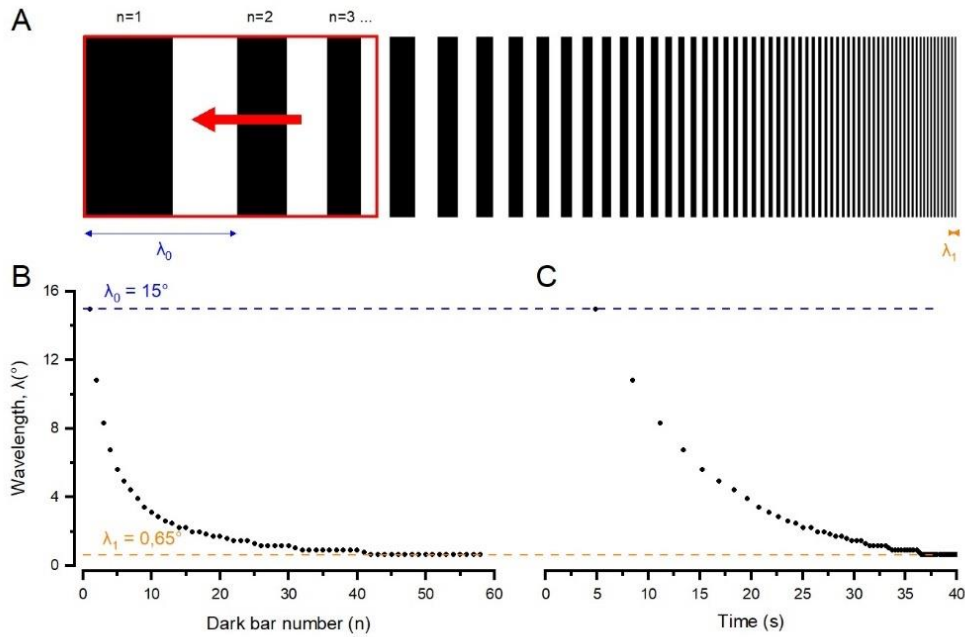
Figure 9: Description of the 4-parameters narrowing grating.
Description of the 4-parameters narrowing grating. **(A)** Grating design. The wavelength narrows from λ_0 to λ_1. Red rectangle: Screen (what is seen by the fly). Red arrow: Grating motion direction, speed is constant. **(B)** Wavelengths at each dark bar. **(C)** Wavelength over time.

2.2.7 Smallest resolvable angle (SRA) calculation

Given the resolvability of individual L2-terminals for every inter-bar distance, I want to compute the SRA. To make a consistent estimation of it and free of human observer bias, I built a six hyper-parameter algorithm in Python that takes the Ca^{2+}-fluorescence signal as input and returns the SRA (**Figure 10**). Two of the parameters enable accurate peak detection, considering the noise in the data (calculated either on the last 5 s of the stimulus during which the stimulus is comparable to a grey frame (**Figure 10A**); or on the first five seconds if the stimulus starts with 5 s of a grey frame (**Figure 10B**)). One parameter is the noise threshold: $R = 0$, if $P_{min} - T$ is smaller than the threshold. The other parameter is the inter-peak noise threshold: $R = 0$, if the inter-peak noise is higher than the threshold. Two separate parameters were used to detect false negatives. The last pair of peaks where $R \neq 0$ is taken as the SRA. In separate tests, the algorithm generated highly similar resolvability estimates to those provided by trained experimentalists.
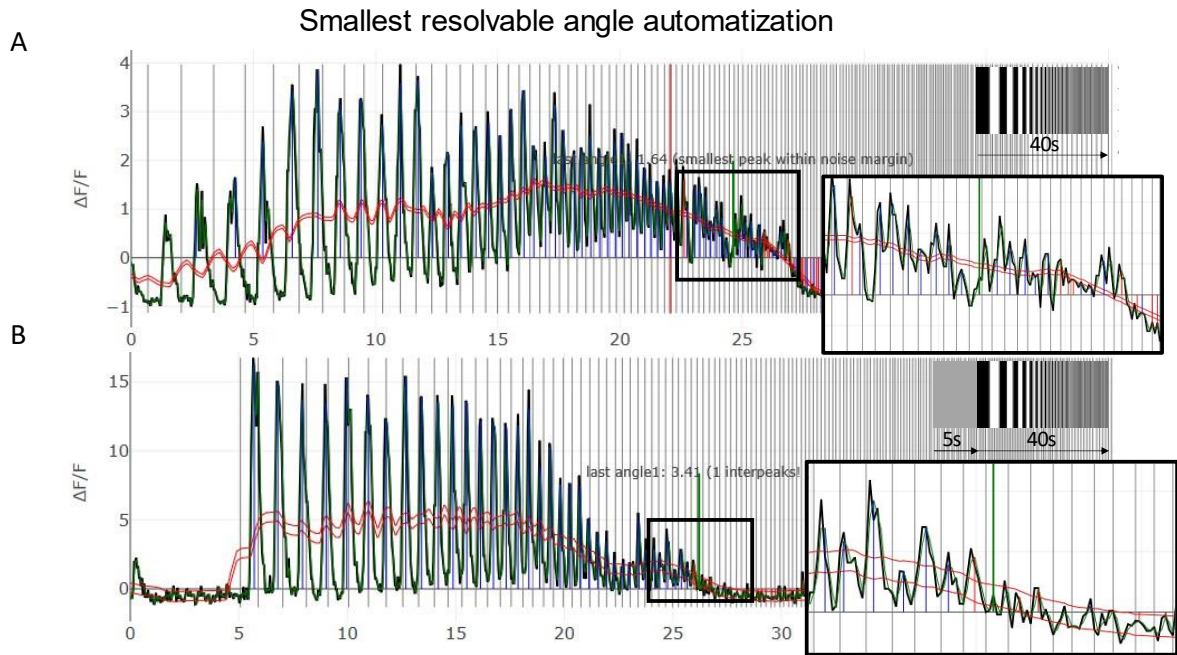
Figure 10: Two examples of smallest resolvable angle automatisation.
(A) Noise calculated on the last 5 s of stimulation (the angle is small enough, so I can assume that the stimulus is comparable to a grey frame). Close-up window: R = 0 because a peak is smaller than the noise level. **(B)** Noise calculated in the first five seconds, during which a grey frame was displayed to the fly. Close-up window: R = 0 because of a significant 'inter-peak' (located between two stimulus-locked peaked). **(A)** and **(B)** Vertical grey bars: times of a bright bar leaving the L2-terminal's RF. Red traces: noise margin (displayed only for the expert to check if the noise calculation is consistent or not). Vertical green bar: time of the smallest resolvable angle.

For each recording, I could simultaneously monitor several (between 1 and 12) L2-terminal responses (**Figure 11B**). The stimuli were presented multiple times to the fly by varying the speed (usually $s = 20, 30, 60°/s$) and the motion direction (usually every 15° or 30°, covering 360°). Hence, this gave me an SRA polar heat map (acuity map) for each recorded neuron in the fly preparation (**Figure 11A**).

2.2.8 L2 preferred orientations' calculation

These SRA polar heat maps almost always suggested the best-resolved orientation (the orientation of highest acuity; or the stimulus orientation for which SRA is smallest). To calculate it accurately and quantify the accuracy, I fitted the SRA (modulo 180°) using a 180° fixed-wavelength sine-function with the Levenberg-Marquardt iteration algorithm (see examples in **Figure 12A**). The reason for this choice is that I expect periodic SRA values with minima at an angle α and $\alpha + 180°$, and a maximum at $\alpha + 90°$ and $\alpha - 90°$. The phase (subtracted by 45°) of the fitting gives me the "preferred" highest-acuity orientation.

I used the Levenberg-Marquardt error values as error margins (**Figure 20B, D**). I also evaluated these fits with the R^2 value (**Figure 20A, C**). Given that Gaussian noise sinusoidal fitting has an R^2 distribution with the mean = 5.8% and rarely reaches 15%, I considered that a clear preferred orientation for L2 SRA fitting was when $R > 25\%$ ($\sim Err < 12°$).
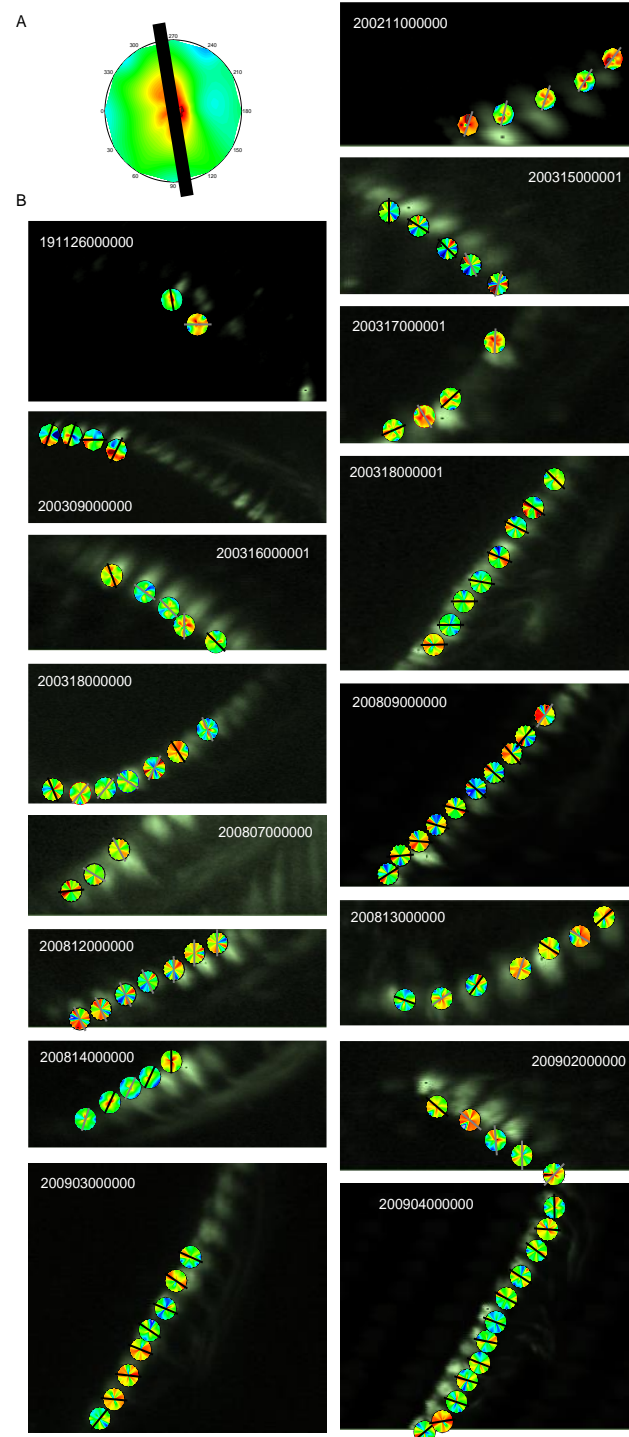
Figure 11: Images of L2 axons terminals with their acuity heat maps.
(A) Acuity heatmap calculated from 16 different stimulus directions and 2 different stimulus speeds.
(B) 16 recordings of different flies' ROIs. Labels: day of recording.
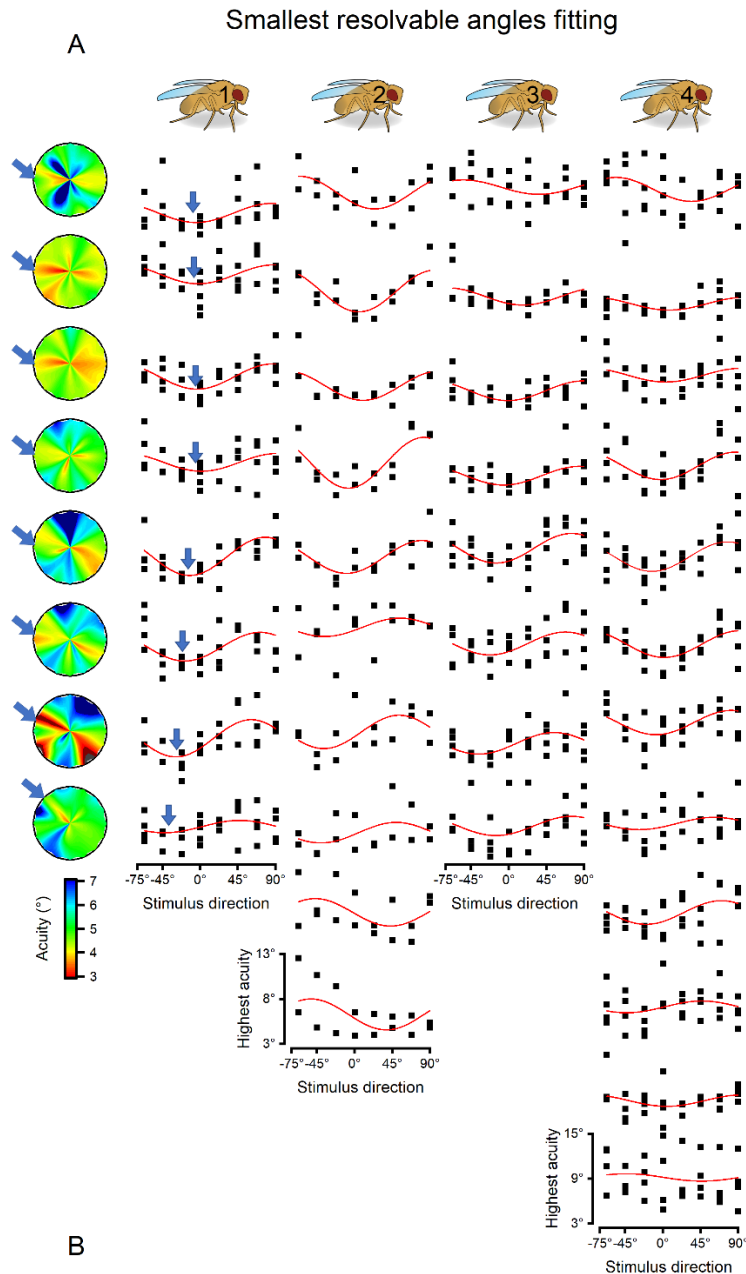
Figure 12: Determining L2-terminal's highest-acuity orientation by sinusoidal fitting (examples from four healthy flies).
(**A**) The obtained SRA is plotted against the stimulus orientation for each recorded L2-terminal. Red curves indicate 180°-wavelength sinusoidal fitting, shown for each L2-terminal (rows) of the four flies (columns), one column per fly. Blue arrows are the "preferred" highest-acuity orientations; chosen as the sinusoidal fits' minima. Heat maps (SRA acuity maps) are shown for the #1 fly every L2-terminal. (**B**) The minima of each fit for the consecutive (neighbouring) L2-terminal, plotted per fly. Error bars give the Levenberg-Marquardt error margin for each fit.

2.2.9 L2 receptive fields' locations calculation

We calculated each recorded L2 neuron's receptive field location using two stimuli: a single light bar moving back and forth horizontally and another vertically. I considered that the

31

peak responses of each terminal were induced by the bar leaving its receptive field (characteristic of an OFF-response). This approach enabled me to reconstruct a good approximation of the receptive field boundaries.

Hence, for each fly, I have a map of the L2 terminals' most sensitive orientation positioned at the corresponding receptive field location.

2.2.10 Motion cancelation technique

Both photomechanical photoreceptor microsaccades and spontaneous intraocular muscle contractions can move the fly brain during 2-photon imaging. I used a computer vision and machine learning library (open-cv) to write a stabilisation algorithm in Python. Two main functions were needed: one (goodFeaturesToTrack) finds the most prominent corners in the image or the specified image region, as described in a proposed algorithm that uses Newton-Raphson style search methods (Shi & Tomasi, 1994). The other (calcOpticalFlowPyrLK) calculates an optical flow for a sparse feature set using the iterative Lucas-Kanade method with pyramids.

2.3 Results

I performed 2-photon Ca^{2+}-imaging from L2 monopolar cells in UV or transgenic flies, which had wild-type photoreceptor pigments (**Figure 14A-right**). The fly's walking behaviour was also tracked on an air-floating ball. The visual stimuli were shown on three screens, spanning about 50°x150° as seen by the walking fly (**Figure 14A-left**). Simultaneously, GCaMP6f was expressed selectively in L2 neurons, and activity changes (fluorescence signals) to visual motion stimuli were imaged at L2 medulla terminals using a two-photon laser.

I used UV flies because their R1–R6 photoreceptors express only Rh3 (UV rhodopsin) and therefore see UV but not green (Wardill et al., 2012), while their L2 neurons express the green-fluorescent Ca2+-reporter GCaMP6f. Critically, UV flies show normal photomechanical microsaccades (**Figure 13**). And as their L2 green-fluorescence Ca^{2+} responses cannot activate the UV-sensitive R1–R6s through orthodromic green-light transmission, UV flies enable naturalistic low-noise conditions for recording high-precision neural signals (**Figure 14C and D**). Even so, the wild-type eye L2-GCaMP6f controls' Ca^{2+} responses showed consistently similar general dynamics; thus, both results were pooled (**Figure 14E**).
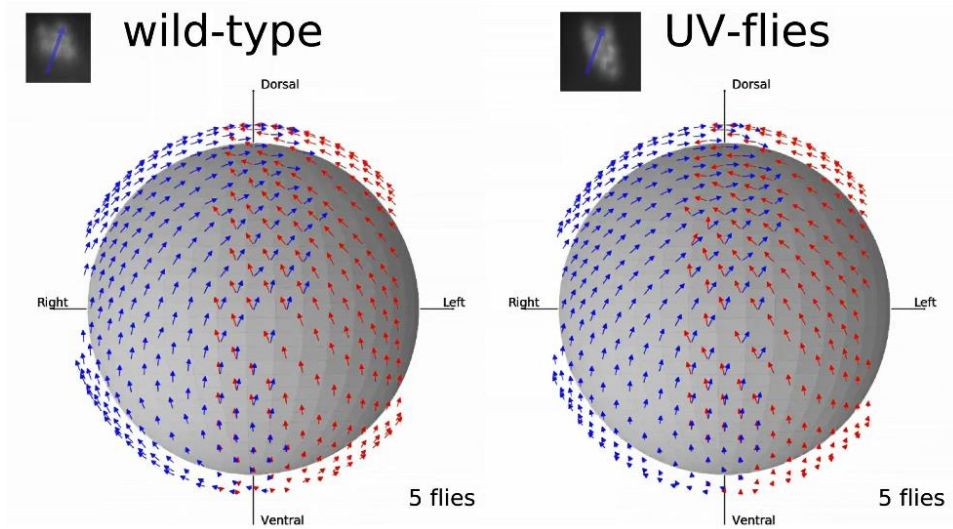
Figure 13: UV flies show standard photoreceptor microsaccade directions across their eyes. The left (red arrows) and right eye (blue) photoreceptor microsaccade movement trajectories of wild-type (left) and UV-flies (right) flies match. The trajectories were calculated through image cross-correlation from light-triggered high-speed deep pseudopupil motion video recording. Average directions are shown; data were recorded from five tethered flies.

2.3.1 L2 neurons' highest acuity can be measured with high accuracy

L2 neurons receive inputs from 7 photoreceptors (R1-R6 synaptically + R7/R8 via cap-junctions; Wardill et al., 2012) branching in the lamina and transmitting their output in the medulla (**Figure 14B**). Initiating the OFF pathway, they depolarise when their receptive fields (RF) luminance decreases (Joesch et al., 2010b). Hence, I can measure their visual acuity by making two light bars cross their RF (**Figure 14C-top**). The resolvability of the two peaks of fluorescence induced tells me what the neurons' acuity is for a specific stimulus. To be more efficient and accurate, I built a wavelength-dynamic grating that narrows in time (**Figure 14C-bottom**) as a stimulus. This stimulus induced wavelength-dynamic oscillations in the neuron's response. I could calculate the resolvability for each pair of peaks (**Figure 14D**). The smallest resolvable inter-bar distance then becomes the smallest resolved angle (SRA). Given that one SRA corresponded to a single orientation and speed of stimulus, I repeated this experiment over a range of orientations and speeds. This approach gave me a sensitivity map, for which one example is shown in **Figure 14F**.
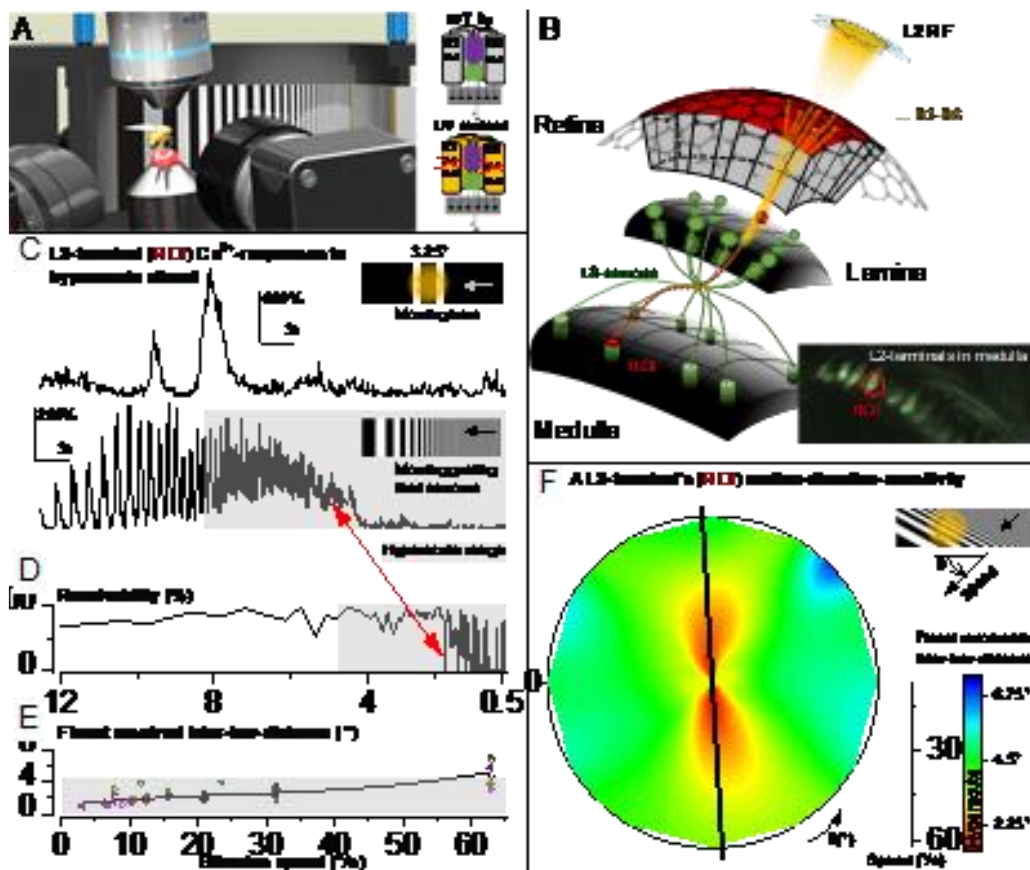
Figure 14: L2 neurons show hyperacuity and have a directional sensitivity
(A) Two-photon imaging setup (the holder is not presented here). The fly head was fixed within a holder. **(B)** L2 neuron (green and red) anatomy. Each neuron is connected to 7 photoreceptors (R1-R6 synaptically and R7/R8 via gap-junctions; yellow) from 7 different ommatidia that roughly see

the same point in space. I visualise the L2 axon terminal (bottom-right). **(C)** Responses to (top) 2 moving bars or (bottom) dynamic-wavelength moving gratings. (C-D) Red arrow: highest inter-bar distance with null resolvability (or the peaks and trough are smaller than the average noise) gives the finest resolved inter-bar distance. **(D)** Resolvability of each pair of peaks for the dynamic grating response. **(E)** Each dot represents the minima of the recorded neurons' finest resolved inter-bar-distances over all the responses to all stimuli directions at one speed for one fly. **(F)** A neuron's sensitivity map shows all the finest resolved inter-bar distances for every stimulus direction and speed. Black bar: best sensitivity(preferred) orientation.

2.3.2 L2 neurons are hyperacute

I managed to record 101 neurons over 17 flies (**Figure 15A**). Most of the SRA measured are in the hyperacuity range (lower than the inter-ommatidial distance) (**Figure 14EF, Figure 15AC**). I found that L2 neurons robustly respond to hyperacute 1 to 4° moving gratings. Thus, by encoding spatial information in time, akin to photoreceptors Juusola et al., 2017), L2s can transmit finer image details than the compound eye's optical limit, 4.5° interommatidial angle (Gonzalez-Bellido et al., 2011), improving vision.
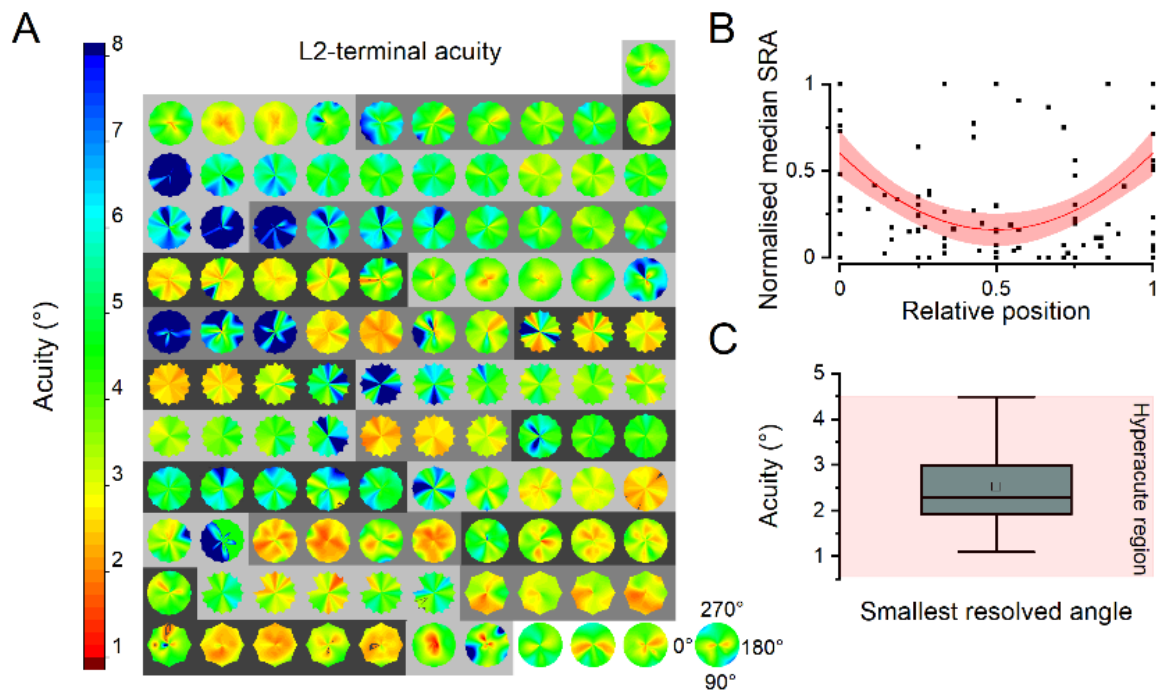
Figure 15: L2-terminal Ca^{2+}-fluorescence responses show hyperacute speed and orientation sensitivity to moving bar-grating stimulation.
 (A) The figure presents the collective motion direction sensitivity maps (*i.e.* acuity maps) of 20 flies, tested with the four-parameter bar-grating stimulation protocol. Each fly displayed at least one and at most twelve actively responding neighbouring L2 medulla terminals. Each fly's maps are shown in their physical order, following their terminals' medulla positions, plotted on the same grey background. The L2-terminals closest to the edge of the imaging window, bordering the dissected tissue area, typically showed the least sensitive responses (from cyan to blue). Whilst the L2-terminals in the centre showed the highest sensitivity and hyperacute stimulus resolvability (from light-green to dark-red), with many neurons encoding less than 1.5° apart bars moving along a specific direction(s). **(B)** Each point represents the median value of a single L2-terminal's *smallest resolved angle* (SRA) map vs its relative position in the recording window. The values are

normalised within each fly. The red line shows a quadratic fitting of the points, with 95% confidence and prediction intervals (red and pink areas, respectively). Two possible reasons for a better acuity in the centre than on the edges: the peripheral terminals might be closer to the dissected tissue, implying health issues; the recording plane and the neural plane are not parallel, and the local fluorescence signal-to-noise ratio (SNR) along the axon is not constant: implying that only the central axons can be sectioned at the highest SNR region. **(C)** Distribution of all the L2-terminal' minimum SRA (highest acuity). The mean (dot) and median (line) are shown. Box range: 25-75%. The pink background indicates the hyperacute stimulus motion resolvability range (<4.5°).

2.3.3 L2 neurons show orientation-specific hyperacuity

Surprisingly, these 2^{nd}-3^{rd} layer visual interneurons do not show the same acuity for all stimulus directions. We can see that some directions are resolved better as the acuity heatmap are not symmetric by rotation. In **Figure 14F,** for example, vertical stimuli induced much smaller SRAs than horizontal ones. However, for most of the recorded terminals, we can see symmetry with respect to the centre of the corresponding acuity heatmap. This result shows that, for each L2 neuron, there is an orientation for which the neuron responds with higher acuity. I will call those orientations the L2 preferred orientations. I assessed the strength of this orientation specificity using the residuals of a sinusoidal-fitting algorithm (See material and methods (**Figure 12**) and discussion (**Figure 20**)).

2.3.4 L2 neurons' hyperacute motion sensitivity and microsaccade directions show a similar shift

I used data from the best-dissected (or healthiest) fly preparations in **Figure 14**, which displayed at least eight consecutive neurons with consistent activity (See discussion). I found that the preferred motion orientations shifted systematically about 5° from neighbour to neighbour (**Figure 16, B to D**), similar to the gradual shifting of the photoreceptor motion directions (**Figure 16C**). This assessment excluded the most peripherally recorded terminals. Such outliers typically showed inconsistent responses (see discussion for more details). Collectively, these results indicate that directional motion information from microsaccadic photoreceptor sampling is retained at the medulla input layer.
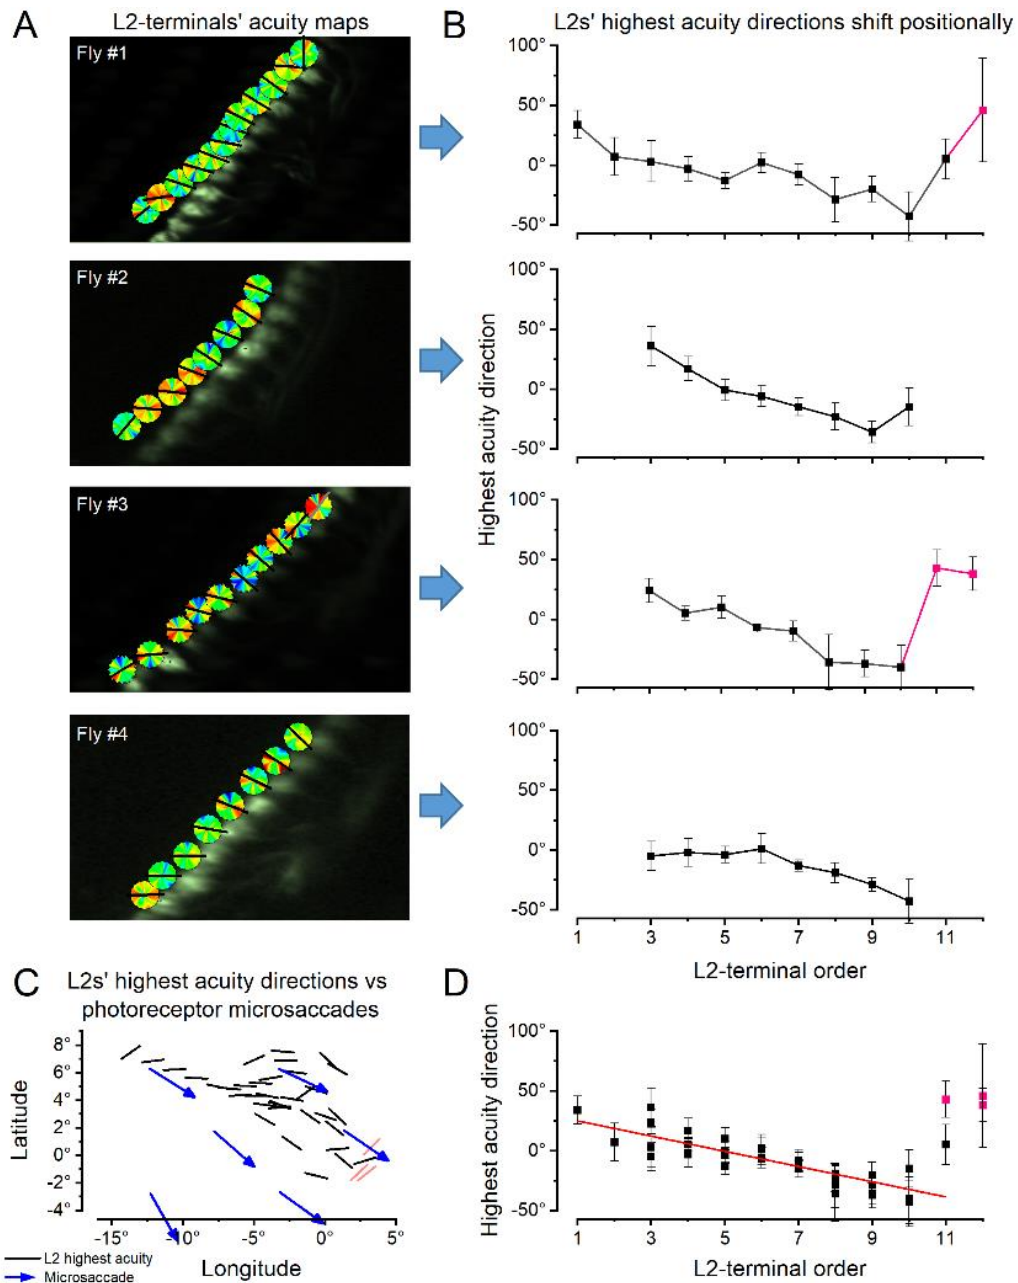
Figure 16: In healthy fly preparations, L2-terminals' motion orientations with the highest acuity shift similarly to their photoreceptors' microsaccade directions.
(A) Two-photon GCaMP6f-fluorescence images of L2 axon terminals from four *Drosophila* (#1-4) that provided long-lasting stable recording conditions. Next to each terminal is its corresponding acuity map, with a black line indicating its "preferred" highest-acuity motion orientation. **(B)** L2-terminals' preferred motion orientations shift gradually and systematically across their retinotopically organised medulla layer. Only a few peripheral terminals (red), closest to the surgically prepared recording window's edge, showed inconsistent, possibly dissection-affected responses. The error bars give the Levenberg-Marquardt error range for each fitted highest-acuity orientation; see **Figure 20** below). **(C)** The locations of the L2 receptive fields (RF; shown for the fly's right eye) with their respective highest-acuity motion orientations (black lines) aligned broadly with the corresponding photoreceptor microsaccades' biphasic motion directions (blue arrows; cf. **Figure 17** below). **(D)** L2-terminals' highest-acuity motion orientations aligned regarding their RF locations.

2.3.5 L2 neurons' hyperacute orientation sensitivity is collinear with the connected photoreceptors' microsaccade directions

I mapped those preferred axes across the whole eye. To know where the receptive field of an L2 neuron is, I make a single light bar across the visual field. I consider that the maximum of the neural response corresponds to the bar leaving the neuron's receptive field. By making the bar stimulus cross back and forth horizontally and vertically gave me a good approximation of the neuron's receptive field boundaries. Their preferred orientations mostly aligned with the photoreceptors' micro-saccades directions that share the same RF locations (**Figure 17B and C**). This finding indicates angular conservation of synaptic information from R1–R6 to L2, consistent with preserving the downstream optic flow processing (Krapp & Hengstenberg, 1996). Future experiments need to test whether this is also true for L1 (ON-channel) and L3 (Fisher et al., 2015; Joesch et al., 2010b; Silies et al., 2013) LMCs, as asymmetric microanatomical adaptations (Meinertzhagen & O'Neil, 1991; Rivera-Alba et al., 2011; Tuthill et al., 2013) may further influence local motion computations. These results demonstrated that L2s collectively convey a high-resolution neural representation of the moving world, maximising visual information.
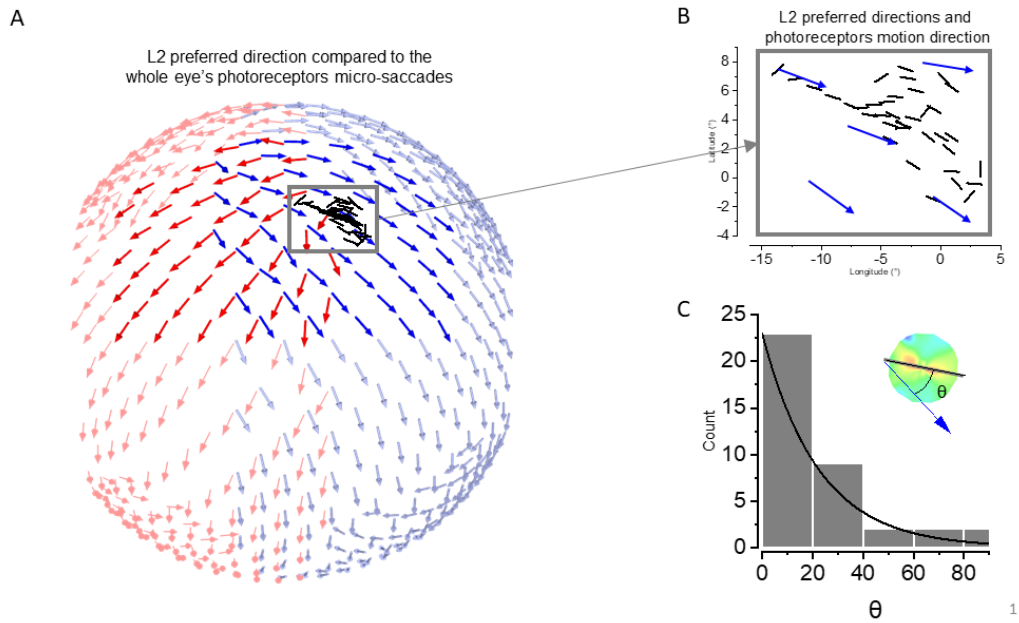
Figure 17: L2's preferred orientations align with the photoreceptor microsaccades.
(A) Sensitivities for 8 neighbouring neurons of one fly plotted with the photoreceptors' receptive fields' saccadic movements. Red/Blue: Left/Right eye, respectively. **(B)** Comparison between L2 preferred orientation and their connected photoreceptors' microsaccade directions. Left: distribution of neurons according to their alignment with the photoreceptor motion. Right: Location of each L2's receptive field with their preferred orientations and the photoreceptors' motion directions. **(C)** View (from the inside) of the whole eyes' photoreceptors' microsaccades. Grayscale bars show the L2's receptive field with their preferred orientations. The darker the bars are, the more colinear the L2's preferred orientation with the motion direction of the photoreceptor microsaccades. Dotted rectangle: photoreceptors whose receptive field overlaps with the screens in front of the fly.

2.3.6 L2 axon terminals physically move in synchronisation with the stimulus

We used a motion artefact correction technique on recordings where the motion artefacts moved the L2-terminal away from the region of interest (ROI) window (typically ~2µm). This technique enabled the ROI fluorescence average to be coherently correlated with the neural activity and not affected by physical displacement. (**Figure 18**) shows the resulting displacements for some cases. Interestingly, the displacements were sometimes stimulus-locked: **Figure 18B** shows slower displacement at the beginning but faster around the end of the stimulus. A high sampling rate (~85 Hz) showed a robust synchronisation between the displacement and the stimulation (**Figure 18A**). These fast stimulus-locked L2-terminal displacements are likely induced by the photoreceptor microsaccades, analogous to the photomechanical tissue displacement recorded during the X-ray imaging experiments. Indeed, photoreceptors move photomechanically back and forth along the central axis each time a bar crosses their receptive fields (Kemppainen et al., 2022). Such a motion could similarly drive L2-terminal displacement in **Figure 18A**.
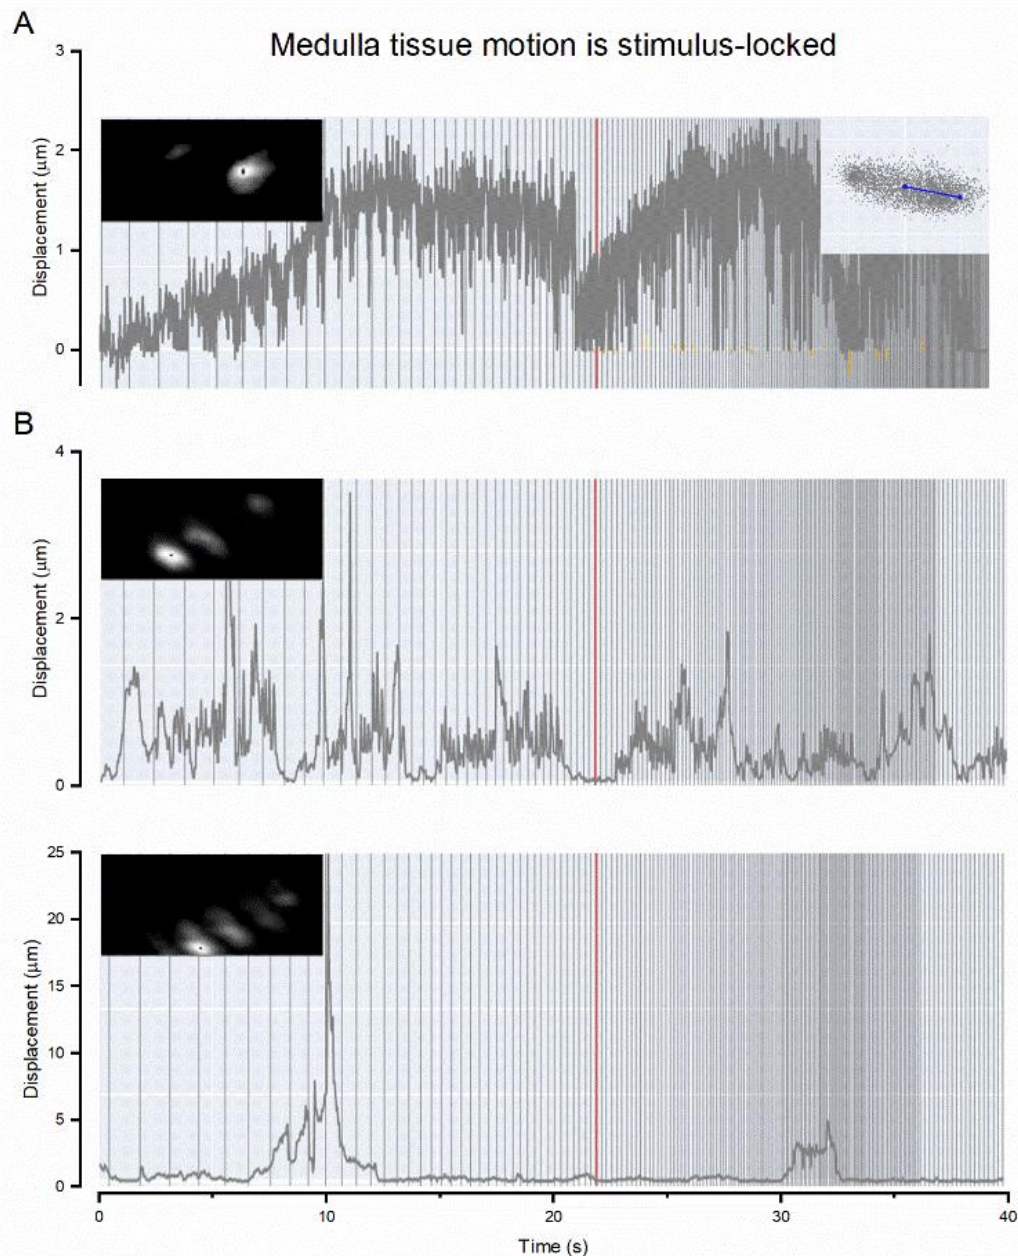
Figure 18: During 2-photon imaging, L2-terminals can show mechanical stimulus-synchronised jitter. I used a stabilisation algorithm to subtract this jitter from the fluorescence video recordings if it was deemed too large.

(A) An L2-terminal (ROI; region of interest) displacement during 85 fps imaging. In the inset, the L2-terminal's position is projected in the principal direction (blue line). Given the regularity and size of these small movements (<<1 μm), they likely resulted from the photoreceptor microsaccades bouncing the optic lobes in a stimulus-synchronised manner. Similar optic-lobe-displacement dynamics were seen during the X-ray imaging (see Kemppainen et al., 2022) **(B)** Two examples of larger mechanical displacements of the medulla L2-neuron terminals, obtained with low (~20 Hz) sampling rates (20 fps). The larger movements (>1 μm) are likely caused by intraocular muscle activity (3) that can move the retina in slow bursts. The smaller movements (<<1 μm) superimposed on the bursts are likely caused by the stimulus-synchronised photoreceptor microsaccades moving the retinal tissue. The three images depict the studied ROI pixels' standard deviation, i.e. showing how the L2-terminal physically moved during the dynamically narrowing bar grating stimulation. The

red vertical lines in A and B indicates the GCamp6f resolvability limit, as obtained from separate flash-stimulation tests.

So far, I have focussed on the synchronisation between the L2 axon terminals' physical movement amplitudes and the stimulus dynamics. This correspondence implies that photomechanical photoreceptor contractions induce tissue movement downstream in the lamina. One can ask if the direction of the movement is random or follow a specific direction (within limits of the recorded 2D plane). I hypothesised that if the microsaccades caused the L2 axon terminal motion, these movements could follow the same axis as their connected photoreceptors. I present below one fly's L2 axon-terminal motion directions as stimulated by dynamic gratings moving in different directions and speeds (every 23° at 20 or 30°/s).

I found that the movement direction is not random but consistently around an average value of 33° (circular standard deviation: 0.29) (**Figure 19**AB). In the first two layers of the visual pathway inside each optic lobe, the cartridges connect the photoreceptors to the lamina terminals following a chiasma structure (cf. the section about the visual pathway in Chapter 1). The centre of the chiasma is roughly in the centre of each optic lobe. Thus, in the 2D-recording plane, a photoreceptor located at an angle α is connected to the L2 neuron located at the angle $90 - \alpha$ (see the transformation in **Figure 19**C). Hence, by applying this transformation to the average axon-terminal movement direction value, I find the expected L2 axon-terminal alignment if their movement was induced by their connected photoreceptors (**Figure 19**C). This conclusion assumes that the photoreceptor microsaccades occur mostly tangent to the eye sphere, as seen in high-speed X-ray imaging (Kemppainen et al., 2022). Accordingly, I found this predicted movement alignment very close to the L2 terminal movement alignment in the 2-photon imaging data (**Figure 19**DE). This result confirms that the L2 terminals in the medulla neuropile move along their connected photoreceptors' microsaccade movement direction axis and provides further evidence that the photoreceptor microsaccadic movements are mostly tangential to the eye sphere. Interestingly, the orientation of these L2 terminal movements slightly correlates to the stimulus direction. Indeed, **Figure 19**B shows that visual stimuli moving more towards 180 ° (i.e. right) induce L2 terminal movements, whose orientation is slightly more toward the forward axis. In contrast, stimuli moving in the opposite direction (i.e. left) cause L2 terminals to move slightly more longitudinally.
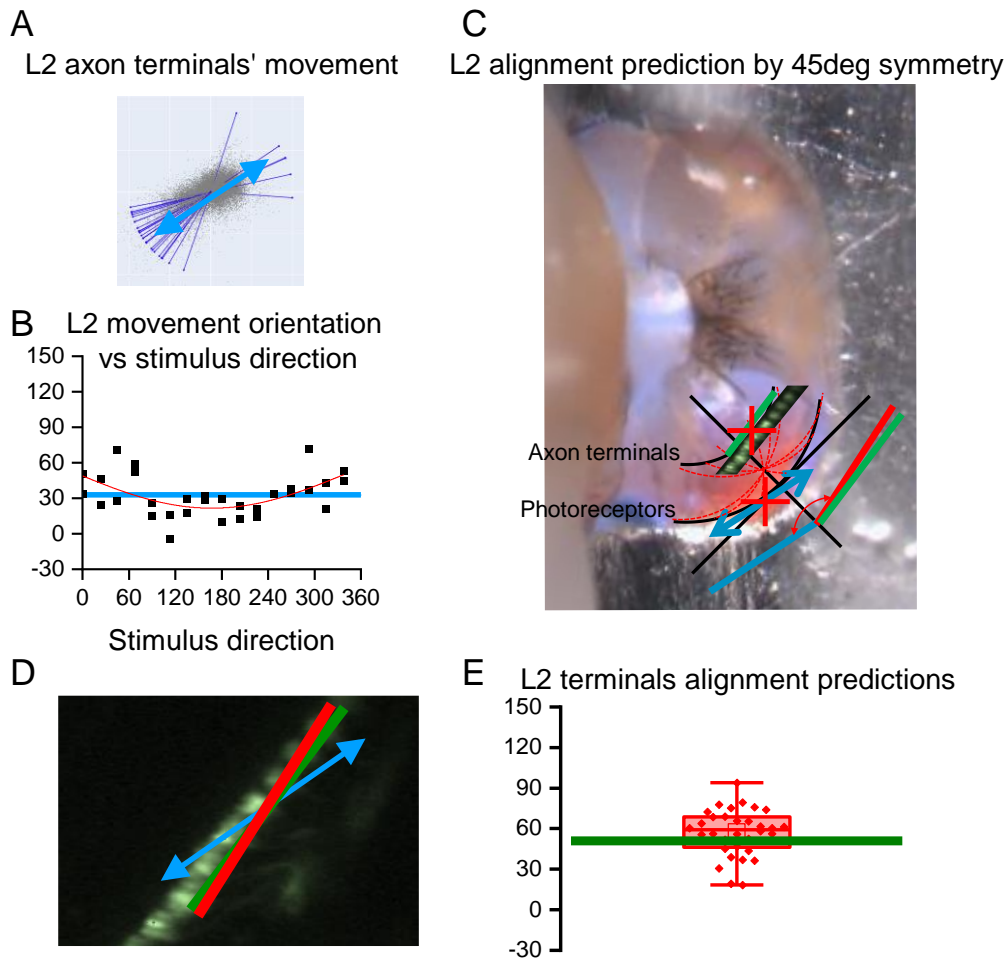
Figure 19: L2 axon terminals' movement orientation axis aligns with their connected photoreceptors' microsaccade direction.
(A) Each point represents the displacement vector of each frame with respect to the stabilisation algorithm reference frame for all stimuli presented to the fly. Each blue bar is the Principal Component Analysis (PCA) orientation of a single stimulus' displacement points. Blue arrow: angular mean of all PCA orientations. **(B)** A plot of the PCA orientations for each of the 32 stimuli. Blue bar: angular mean. Red line: Levenberg-Marquardt 360°– wavelength sinusoidal fitting (middle point: $167°$; $R^2 = 0.31$). **(C)** Black curves show the theoretical alignments of the photoreceptors and the axon lamina terminals. The black line is the 45° symmetry line splitting the optic lobe into two equal parts. The bottom red crosse shows where the L2 axon terminals' movement orientation (blue arrow) is tangential to the eye circle. The top red cross shows the location of the bottom red cross transformation after chiasma symmetry (shown with the red dotted curves between the photoreceptor and axon terminals curves. Green lines represent the actual recorded L2 axon terminals alignment axis. **(D)** Close-up of the recorded L2 axon terminals with their physical movement orientation (blue arrow), predicted alignment (red line) and actual alignment (green line). **(E)** Distribution of the L2 axon terminals alignment predictions for every stimulus. Green bar: actual alignment.

2.4 Discussion

The difference between 'healthy flies' and the other 12 flies

I call healthy flies all the flies where I could simultaneously record Ca^{2+} signals to simple flashes from more than 8 L2 terminals. Thus, when only a few L2s responded, those were next to an inactive neuron (probably due to dissection imperfections). I suspect such tissue damage would affect how they respond, explaining the outliers (acuity and orientation wise) at the edges of the recorded windows. In **Figure 15**, I presented all the recorded terminals' preferred stimulus movement orientations (for the finest resolvability as SRA maps) and displayed the error of their sinusoidal fittings.

However, another possible reason for the edges being outliers can be that because the recording plane and the L2 plane are not parallel, I cannot record the same areas of each L2 neuron. This nonconformity could induce a difference in noise-to-signal ratio and hence, outliers at the edges.

Figure 15B quantifies the outliers' issue by fitting quadratically the normalised medians of all the SRA maps: the neurons at the centre of the recorded screen tend to resolve better than those at the edges.
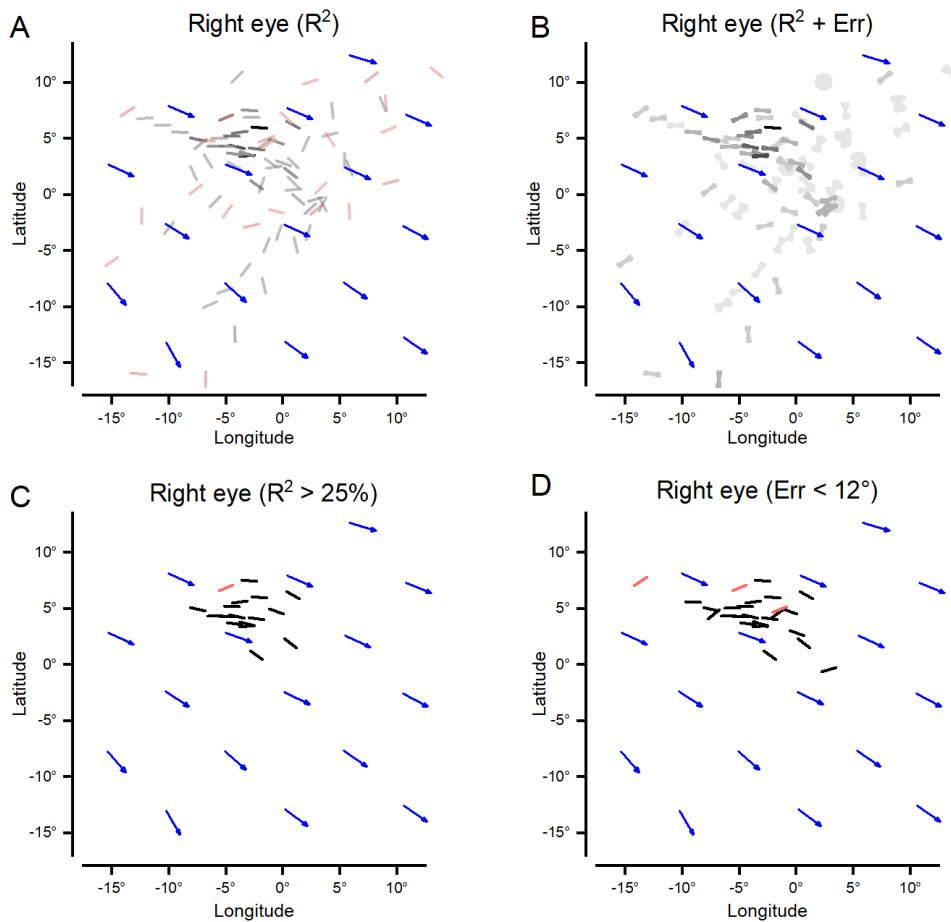
Figure 20: All the measured right eye L2-terminals' highest acuity orientations, plotted regarding their receptive field (RF) locations.

(**A**) R^2 of the sinusoidal fits are shown in a linear grayscale: the lightest 0%; the darkest 100%. The darker the line, the better the reliability and predictive value of the estimated highest-acuity orientation. Red: peripherally recorded L2-terminals (possible outliers). Blue arrows: the photoreceptor micro-saccade directions at their corresponding RF locations. (**B**) R^2 of the sinusoidal fits shown with the same reliability-dependent colouring as in **A**. The error of the sinusoidal fits is shown with their orientation margins. (**C**) L2-terminals' highest acuity orientations when the sinusoidal fits' R^2 is >25%. (**D**) L2-terminals' highest acuity orientations when the sinusoidal fitting error is <12°.

Some neurons do not show a preferred orientation

Almost half of the L2 terminals showed no apparent orientation specificity (**Figure 20AB**). Several possibilities could contribute to this conundrum. First, the orientation specificity arises from responses to very fine patterns when the grating moves to a specific orientation. Hence, the resulting fine pattern responses cannot be resolved if they do not have a high enough signal-to-noise ratio. This limit, which directly reflects the combined imaging and Ca^{2+}-fluorescence system sensitivities, makes it harder to see the preferred orientation in

44

the fly preparation is suboptimal. Moreover, this limit might be further degraded by motion artefacts, decreasing the observed neural activity resolution.

Notice that I expected L2 neurons to have a direction (instead of orientation) specific hyperacuity. Indeed, each connected photoreceptor moves fast in one direction and slowly in the opposite direction. Why this asymmetry does not appear in the L2 acuity properties remains unknown. However, I suspect this characteristic might underlie a presynaptic mechanism that encodes photoreceptors' directional stimulus motion resolvability in a way that needs to be studied in the future.

The locations of the receptive fields cover a small portion of the fly eye

The combined receptive fields of the healthy flies' L2 neurons span roughly 15°x15°, which is very small compared to the fly's panoramic visual field. Several possible reasons can explain these limitations. First, the optic cables used to display the stimuli to the fly do not distribute the light intensity uniformly across each screen. Instead, the centre of each screen is slightly brighter than near the edges. Hence, L2 neurons with receptive fields around the centres of the screens might respond better than the others. This predicament, however, does not explain why the L2 neurons with receptive fields at the centres of the left and right screens do not show clear responses. This issue could be related to dissections imperfections. Because the chiasma twists the photoreceptor axons, the L2 terminal activity measured at the top part of the brain is actually driven by the photoreceptor receptive fields at the centre. Since the lateral parts of the fly head are lower than the top part of the head, this part of the fly brain is more in contact with hot wax, which might damage the tissue and consequently induce weaker responses to this area of the visual space.

Hence, more investigation into other parts of the visual field is required to conclude with more confidence that their preferred orientations follow the whole eye's photoreceptor microsaccade orientations.

Sampling aliasing issue

I used relatively low frame rates of 20-25 fps to concurrently image many L2-terminals with a high signal-to-noise ratio (*i.e.* each complete image frame was sampled at ~20 Hz). Such sampling could be prone to aliasing; if the actual light-stimulus-induced fluorescence

changes happened faster than the sampling. However, several factors ensured that aliasing effects on the data were minimal:

Each image frame is not an instant snapshot but is built up by scanning its pixels line-by-line at ultra-high speed (each pixel in ~50 ns). Thus, both the used resonant scanner's line scan rat and the recorded local Ca^{2+}-signals' (pixel-wise) spatiotemporal correlations are much faster than the full image frame rate and the underlying Ca^{2+}-fluorescence dynamics.

The Shannon-Nyquist sampling theorem states that no information is lost if the sampling rate is higher than twice the signal's maximum frequency. Hence the minimum consistent value for SRA (smallest resolvable angle) follows the rule:

$$\alpha_{min} > \frac{\omega}{f_s}$$

eq. 2-5

, where α_{min}, ω and f_s are the minimum inter-bar distance used for the SRA, stimulus motion speed and sampling rate, respectively. Those minimum values for the SRA were rarely reached, so the risk of aliasing was minimal.

The sampling rate was never kept constant in the recordings, thus minimising any systematic aliasing effects. Theoretically, aliasing causes central symmetrically spreading patterns in the recorded images, such as fake rigs or harmonic ringing (Juusola et al., 2017), which never occurred in the SRA maps.

Control experiments with much higher frame rates (85-145 fps) generated even higher L2-terminal acuity maps than those with 20 fps sampling. Yet, these showed similar orientational selectivity trends, showing clear hyperacuity and specific highest acuity motion orientations. The acuity map trends for the 20 fps and >85 fps sampling started to differ only at the highest tested velocity stimuli (60°/s). One acuity map for 85 fps sampling was included in **Figure 15A**. Overall, I found a suggestively higher L2-terminal hyperacuity for the higher sampling rate data **Figure 21B**:

- *High fps*: 2.20° ± 0.25° (mean ± SD); SRA = 1.93°, Median = 2.17°, Max = 2.5° (n = 6 L2-terminals)
- *Low fps*: 2.53° ± 0.82° (mean ± SD); SRA = 1.09°, Median = 2.31°, Max=6° (n = 117 L2-terminals)

Therefore, in light of all this evidence, together with *Drosophila*'s striking hyperacute visual behaviours in a flight simulator system (Juusola et al., 2017) and faster intracellular voltage responses (Nikolaev et al., 2009; Zheng et al., 2006), I am confident that I present reliable and conservative estimates (lower bounds) of the L2-terminals' orientation-selective hyperacuity (for the given experimental conditions, instrumental noise and sampling limitations). In further predictable support of these findings, behavioural experiments in a classic flight simulator system demonstrate hyperacute *Drosophila* vision (Juusola et al., 2017; Kemppainen et al., 2022). A freely flying *Drosophila*'s visual acuity can only be better in natural environments and could even be significantly higher.
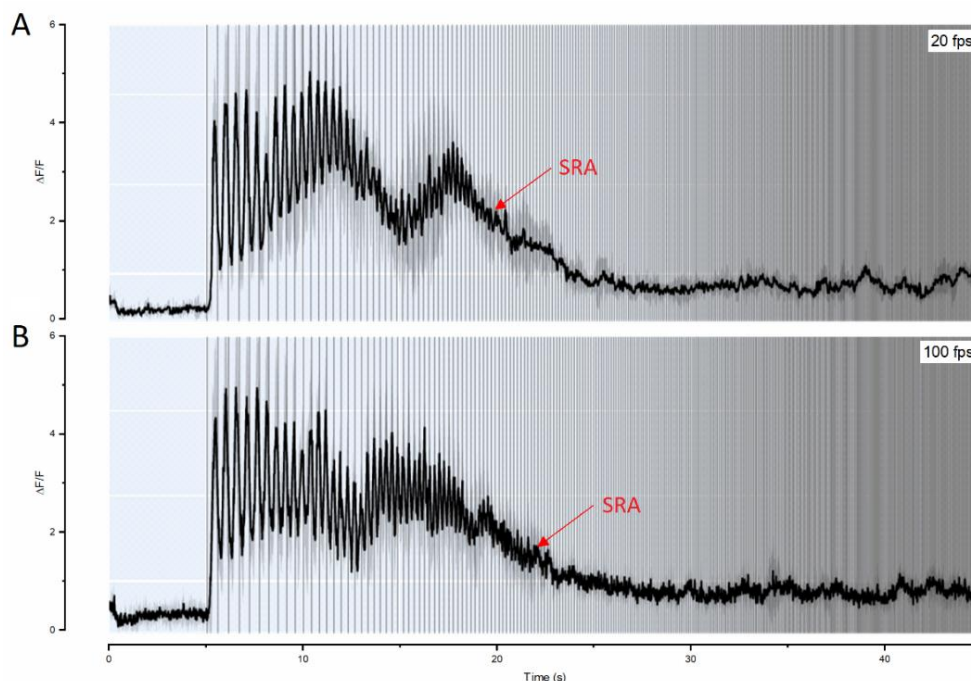
Figure 21: Using a higher imaging frame rate (i.e. Ca^{2+}-signal sampling rate) increases the recorded L2-terminals' resolvability for fast-moving stimuli.
 (A) Average of 4 repeats of the same L2-terminal (ROI) responses at 20 fps to a 60°/s narrowing (13° to 0.65°) grating stimulus moving upwards following 5 s of a grey frame. (**B**) Average of 4 repeats of the same L2-terminal (ROI) responses at 100 fps to the same stimulus as in **A**. Grey margins: ± SD. Red arrows: the first pair of peaks with null resolvability, edging the smallest resolved angle (SRA). Gray vertical bars: times when a dark bar of the stimulus crosses the L2's receptive field.

The use of "UV-flies" minimise antidromic sampling artefacts.

Because the basement membrane between the lamina and retina lacks screening pigments, photoreceptors can be stimulated antidromically by shining light through the fly brain (Juusola & Hardie, 2001). Equally, during Ca^{2+}-imaging, fluorescence signals from the brain circuits propagate toward the photoreceptors. Therefore, in *Drosophila* with wild-

type spectral sensitivities, the green-light-activated R1-R6s and R8y photoreceptors inadvertently multiplex light stimuli from the world with the L2 green-fluorescence signals from the lamina, potentially obfuscating downstream visual processing (as recorded by two-photon imaging). I used "UV-flies" (see material & method section) to overcome this problem.

Time adaptation

We believe that neural responses, despite the low order of L2 neurons, are not stationary. Hence, because of the small number of recordings (compared to the conclusion I want to make), I must be aware that this dynamism can affect the acuity of the cell over time. One solution is to randomise the stimuli order, which is what I have done. Another way would be to check whether there is any correlation between the neural responses and time.

L2 displacement

We should bear in mind that the stimulus-induced fluorescence variations themselves may fool the stabilisation-algorithm by faking a motion. However, this phenomenon is unlikely because applying the stabilisation-algorithm on the stabilised video only resulted in small and noisy motion residuals. Hence, even though the collective evidence from separate experiments using different assays is already compelling, an additional displacement analysis on activity-independent fluorescence (such as Tomato dye) can be done in the future for conclusive proof. Note that the small L2-terminal displacements, such as the one seen in **Figure 18**, had no real effect on the recorded fluorescence signals, so subtracting them made no difference in the analyses.

Our analysis regarding the L2 terminals' movement orientations is limited because my prediction is valid only if the photoreceptors connected to the recorded terminals move on the vertical cylinder surface. Otherwise, the projection of their movement on the recording section (a horizontal plane) would incorporate a non-null component that would mislead the prediction through the 45°-axis symmetry. For example, a photoreceptor located at the bottom of the eye (and nearest to the centre of the head) would move roughly towards the back (according to Figure 7, left). The 45°-axis symmetry transformation shown in **Figure 13C** would then consider photoreceptors on the side of the head and predict L2 terminals near the centre. This simplification might be a reason for my prediction error (51° instead of 57°). However, since the screen viewed by the fly is located only a little below (<10°) the recorded section (see **Figure 22**), considering the photoreceptors moving within a vertical cylinder is not an aberration.
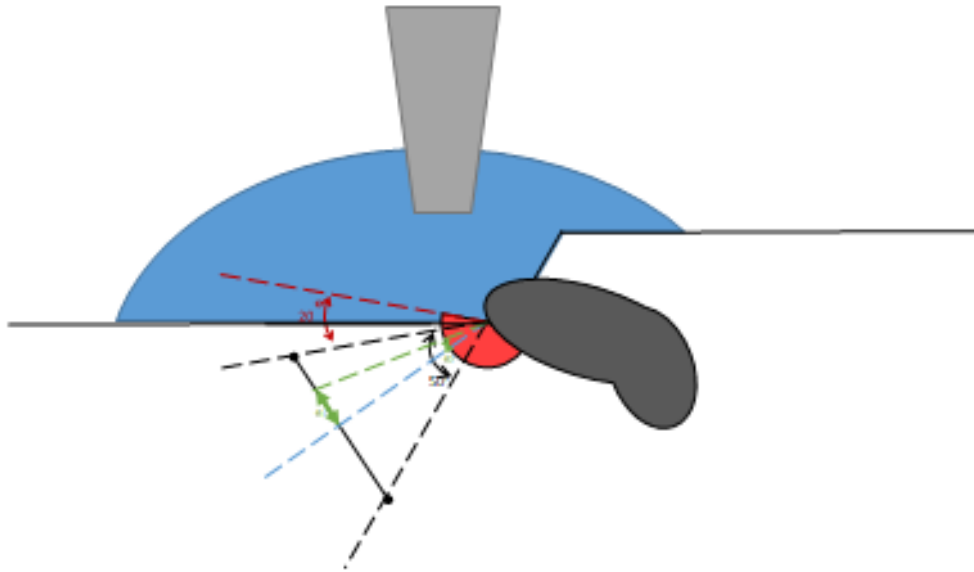
Figure 22: The photoreceptors stimulated by the screen are roughly in a vertical cylinder.
 In blue is the saline covering the brain above which the cuticle was opened and removed.

The terminals movement orientation is slightly different for a grating stimulus moving towards the left compared to a grating moving towards the right was not expected. I have not found any reasonable explanation for this phenomenon so far. Notice that I analysed the terminals' movement directions only on one fly for the sake of time limitations. Further analysis should be done on the other flies to confirm the related results and conclusions.

Chapter 3: Biologically-constrained deep-learning model of the L2 neurons' output

3.1 Introduction

The world we live in is too complex to be understood thoroughly. Instead, we build our knowledge and take decisions based on models of this reality, such as models of the forces of nature, chemical interactions, or living things. They help us understand the mechanisms underlying the reality we observe, predict future events, and improve our intuition for designing novel technologies. Modelling living organisms has, for a long time, sparked our interest. From the idea of souls controlling bodies to cells containing DNA, our curiosity regarding what/who we are has always persisted. The concept of Artificial Intelligence emerged especially with the advent of computers, in which input/output systems many scientists considered to be analogous to our brain. In the last century, these analogies became more apparent with the emergence of artificial neural networks (ANN) that solved many human-level tasks and outclassed humans in complex games such as Chess or Go. Based on how our brain is structured, this last generation of AI owes a lot to neuroscience research. However, ANNs are still very limited compared to our brains' good sense and sound judgement in practical matters, task generalisation, or energy efficiency.

Understanding the brain by modelling the neural activity of the preferred model organisms, such as *Drosophila*, has been a common approach to improving this branch of AI. For example, most *Drosophila* brain connectome has been mapped to the details of individual synapses (see the General Introduction chapter), providing a plausible wiring diagram of neural information flow. This motivated building connectome-constrained ANN to get insight into the usefulness of their brain structure by playing with hyper-parameters.

In this chapter, I propose an ANN model of *Drosophila*'s photoreceptor and lamina layers. This ANN is inspired by the anatomy of the real biological neural network (BNN). Some of its parameters, such as synaptic strengths, threshold values, receptive field locations, etc., are learnt during a training process, during which the ANN tweaks them so that the output resembles the recorded L2-axon-terminals calcium fluorescence of a real fly. I think the similarity with the BNN helps untangle the complexity and cleverness of the mechanisms happening in the fly brain beyond the well-known simplified high-order models for fly vision, such as the Hassenstein-Reichardt correlator model for motion estimation. Additionally,

this ANN brings insights into the role of particular connections or mechanisms and enables predicting neural behaviour and synaptic strengths.

3.2 Material and methods

3.2.1 Sampling by photoreceptors: neural network inputs

The ANN model inputs pass first through a layer of photoreceptors that sample light information in four possible ways:

1. A static (positionally-fixed and non-moving) photoreceptor sampling model. A naive fully-connected layer in which every screen 'pixel' connects to a static photoreceptor with a learnable weight.
2. A static (positionally-fixed and non-moving) photoreceptor sampling model. A microsaccade-less learnable multivariate gaussian filter layer.
3. A simplified version of the microsaccadic model (photomechanically moving photoreceptors; Kemppainen et al., 2022) with one photoreceptor per ommatidium.
4. A simplified version of the microsaccadic model (photomechanically moving photoreceptors; Kemppainen et al., 2022) with seven photoreceptors per ommatidium.

For the 1st and 2nd static photoreceptor sampling models, I pixelate the screen by 3°resolution. The screen values are normalised. The sampling rate is 80 ms. The time step is 2 s.

For microsaccadic photoreceptor sampling (the 3rd and 4th models), I use the full-screen resolution of 0.13° and a 1 ms sampling rate. These values match those used in the differential equations (Kemppainen et al., 2022). Finally, I chose a time step of 50 ms.

Photoreceptor microsaccade

For the latter ANN models, I simulated the observed photoreceptor microsaccade mechanisms.

At each time step t, each photoreceptor input results from a fixed size gaussian filter applied on the current frame, multiplied by a scaling factor of 700.

For the model, this input represents the number $N_p(t)$ of photons absorbed by the photoreceptor's microvilli. The number of photons absorbed by each microvillus is determined by a uniform stochastic sampling with a repetition process.

Following a photon absorption (by rhodopsin) and being light-activated, the microvillus goes through three phases before it can be light-activated again. First, it enters an absorption state for a stochastic number of time steps that follows a gamma distribution of parameters $k = 9.9891$ and $\theta = 1.7160$. It then proceeds into a quantum bump state, generating a unitary elementary electrical response, as described further below. Finally, it reaches a refractory state for a stochastic number of time steps that follows a gamma distribution of parameters $k = 3.4710$ and $\theta = 12.6873$. Only at the end of this third state can the microvillus be light-activated again.

At each time step, the number $A(t)$ of microvilli generating a quantum bump contributes to the lateral (sideways) rhabdomere movement. The rhabdomeres move along one axis. On this axis, position x of the rhabdomere relative to its initial position follows this differential equation:

$$\frac{d^2x}{dt^2}(t) = \frac{A(t)^{A_n}}{A_h} + D_c(D_b{}^{-D_e\frac{dx}{dt}(t)} - 1) - S(t)x(t) \qquad \text{eq. 3-1}$$

with the spring function:

$$S(t) = S_0 + S_c\frac{A(t)^{S_n}}{S_h}, \qquad \text{eq. 3-2}$$

where the constants are:

$$A_n, A_h, D_c, D_b, D_e, S_0, S_c, S_n, S_h = \{2, 9500, 8 \times 10^{-5}, 2, 2100, 1 \times 10^{-4}, 1.15 \times 10^{-5}, 1.3, 200\}.$$

The total light intensity driving the photoreceptor is:

$$LI(t) = min\left(LI_{max}, LI_0 + 200 \times \sum_{k=0}^{t} N_p(t)\right), \qquad \text{eq. 3-3}$$

where $LI_{max} = 1.11 \times 10^7$ is the maximum light intensity that a photoreceptor can get and $LI_0 = 4 \times 10^6$.

$$B_{scale}(t) = 3.62 - 0.51 \times log_{10}(LI(t)) \qquad \text{eq. 3-4}$$

$$B_{peak}(t) = \frac{\Gamma(B_{scale}(t)+1)^2}{\Gamma(4)^2} \times \frac{2^{2 \times B_{scale}(t)+1}}{\Gamma(2 \times B_{scale}(t)+1)} \times \frac{2^7}{\Gamma(7)} \qquad \text{eq. 3-5}$$

$$B_t(t') = \frac{B_{peak}(t) \times \left(\frac{t'}{\tau}\right)^{B_{scale}(t)}}{\Gamma(B_{scale}(t)+1)} \times \frac{1}{\tau} e^{-\frac{t'}{\tau}} \qquad \text{eq. 3-6}$$

The duration $B_d(t)$ of the quantum bump state (the second state of a microvillus after absorbing a photon) is the time needed for the bump function to come back close to zero.

Precisely:

$$B_t\big(B_d(t)\big) > \varepsilon \text{ and } \forall t' > B_d(t), B_t(t') < \varepsilon \qquad \text{eq. 3-7}$$

I chose $\varepsilon = 0.02$.

The number of photoreceptors per ommatidia:

I either used a simplified model with one photoreceptor per ommatidium or a model with seven photoreceptors per ommatidium. In the latter model, each of the 7 photoreceptors is responsive to different locations in the visual space (**Figure 23**). In the same way, owing to the eye's spherical geometry, 7 photoreceptors from different neighbouring ommatidia are responsive to roughly the same location in the visual space. The photoreceptors within the same ommatidium move as a group at the same speed vector. I calculated this grouped microsaccade using the sum of the 7 photoreceptor inputs as the microsaccadic model input.
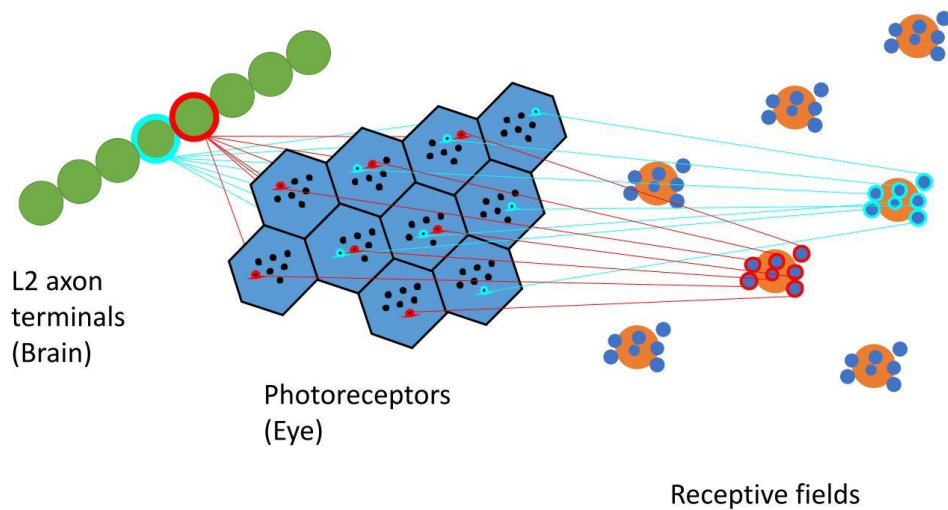
Figure 23: Exact receptive field location per photoreceptor compared to the cartridge approximated receptive field location.
Right: the blue dots represent the photoreceptors' receptive field locations. The orange dots represent the approximate centre receptive field for a cartridge (e.g. light blue and red are two different cartridges). **Middle:** Location of the photoreceptors within the fly eye. 7 different ommatidia of one cartridge (e.g. light blue or red) respond to one centre receptive field and feed to one L2 neuron (**left**).

Note that the 7 neighbouring ommatidia pattern is not symmetric by rotation and flips. Hence, I chose the pattern that best matches the experimentally measured patterns of the deep pseudo-pupils (Kemppainen et al., 2022) (**Figure 24**).
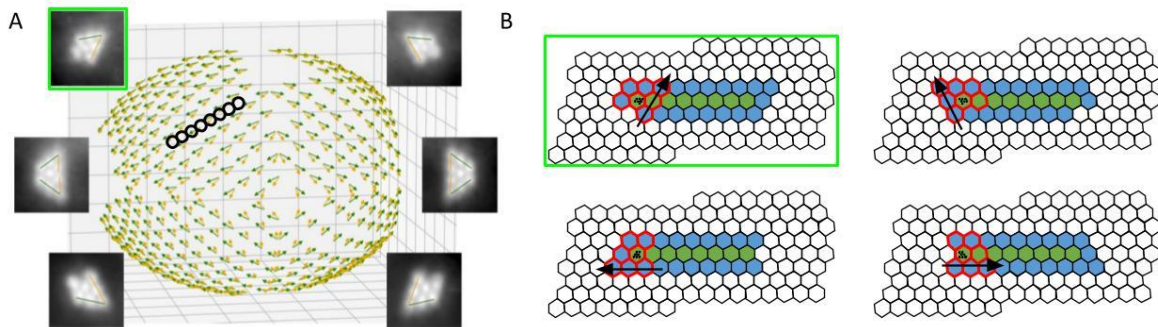
Figure 24: Axis of neighbouring ommatidia feeding the same cartridge.
(A) Axis of the 7 neighbouring ommatidia feeding into the same lamina cartridge across the whole eye. Black circles: estimation of the recorded L2 receptive field locations. Green rectangle: orientation of the 7 neighbouring ommatidia near the recorded L2 receptive field locations. Data from Kemppainen et al., 2022. **(B)** Four examples of all the possible orientations of the 7 neighbouring ommatidia. Green rectangle: Best matching orientation with the deep-pseudopupils' axis at the recorded L2 receptive field locations.

Initialisation:

I chose three possibilities of initialisation for the Gaussian positions.

The first is a stochastic position for each Gaussian that follows a uniform distribution around the centre of the screen (I disallowed the positions less than 120 pixels from the corner).

The second is the positions measured experimentally (see Chapter 2).

The final one uses the fact that each receptive field is shared between the 7 photoreceptors located in different ommatidia and feeding the same cartridge (**Figure 23**). That is why, for the 7-photoreceptor-per-ommatidium model, I added the surrounding ommatidia, in which at least one photoreceptor outputs to a recorded L2 neuron. I also added the receptive fields of every photoreceptor in the same ommatidia as one of the $7n$ photoreceptors that feed to the n recorded L2 neurons, since they contribute to the total ommatidial rhabdomere movement. In this arrangement, the 7 photoreceptors of the same ommatidium move together along one direction (at the average amplitude of their contractions). Inversely, the photoreceptors of one cartridge (sharing roughly the same receptive field but different

54

ommatidia) contribute similarly to different ommatidia that move in different directions (**Figure 25**).

Figure 25: toy example of the 7-photoreceptor model's contractions.
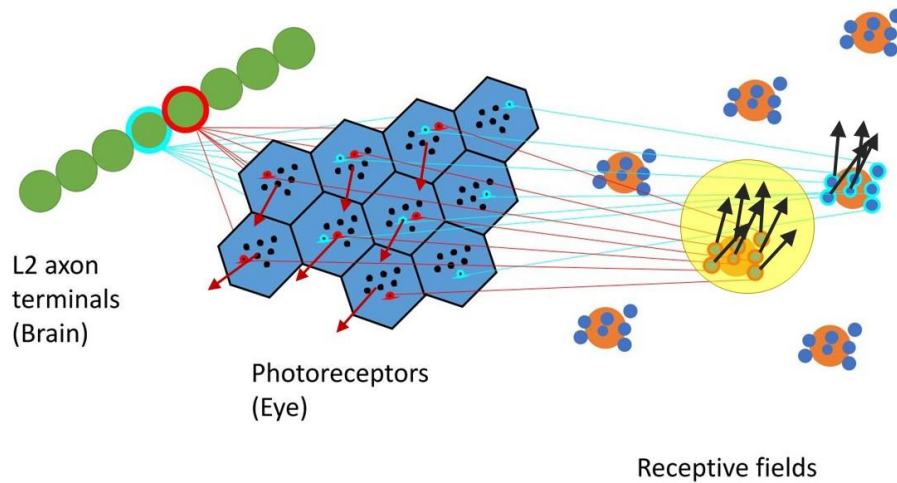The yellow circle represents a transient light input change. The red photoreceptors contract photomechanically to the light change; moving (as a collective) in a particular direction (as set developmentally for this eye location). Because of the ommatidium lens (which acts as a 'lever') and the photoreceptor rhabdomeres' asymmetric positioning around it, the photoreceptors' receptive fields move in the opposite direction (of their microsaccade), even if they did not directly experience the light stimulus. Here, I only consider the red and light blue photoreceptors as toy examples.

Thus, I needed two layers of receptive fields around the n receptive fields: this increased the number n (8 for the fly I used) of receptive fields to $(n + 4) \times 5 - 1$ (59 for the fly I used). Since I only have the experimental values of the recorded L2 neurons' receptive fields, I inferred the initial receptive field positions for the 59 receptive fields using the following technique. I computed the centre of mass (COM) of the n experimental receptive fields. I calculated the principal vector of the distribution of these n receptive fields using the Principal Component Analysis (PCA) algorithm from the *sklearn* Python package. I then computed the average distance $\bar{d} = \frac{1}{n-1} \sum_{i=0}^{n-2} d_i$ where d_i was the distance between the receptive field i and the receptive field $i + 1$. Next, I placed the n new receptive fields along a line. This line is the space generated by the eigenvector corresponding to the biggest eigenvalue of the covariance matrix of the experimental receptive fields. I placed these new receptive fields to have the same COM as the experimental receptive fields. The distance between each new receptive field is $\bar{d}$. Finally, if necessary, I added the surrounding receptive fields around the new receptive fields so that each photoreceptor in

the same ommatidium as a photoreceptor inputting one of the recorded L2 neurons has its receptive field (**Figure 26**). This is why I end up with $5 \times (n + 4) - 1$ new receptive fields.
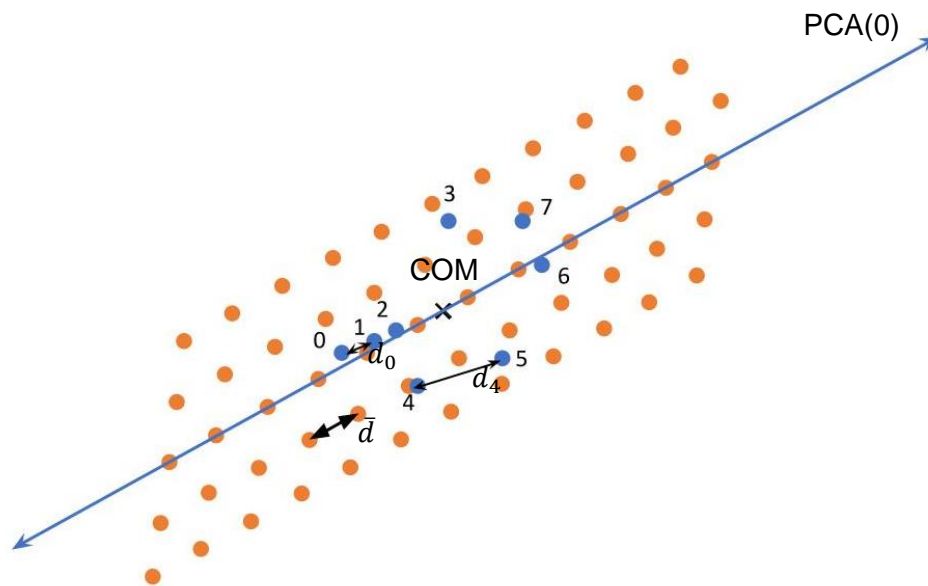
Figure 26: Hexagonal lattice of ommatidia inferred from the experimental receptive fields.
Blue points: recorded L2 receptive field locations, numbered from 0 to 7, in the order of their position in the brain. Orange points: Inferred receptive field locations calculated from the PCA eigenvector, the average distance between two neighbouring L2 neurons' receptive field locations and the experimental receptive fields' centre of mass.

In reality, the 7 photoreceptors of one cartridge have slightly different receptive field locations and half-widths (**Figure 23**).

The 7 locations are $\{(-1.89, -0.722), (-1.56, 1.00), (0,0), (0.39, 1.4), (2.05, -1.50), (1.89, 0.19), (2.39, 2.28)\}$.

The 7 half-widths are, in degrees: $\{5.1, 4.6, 3.12, 4.5, 5.2, 4.6, 4.8\}$.

Figure 27 describes the matrix used in the Tensorflow script for the photoreceptors' connections with the cartridges.
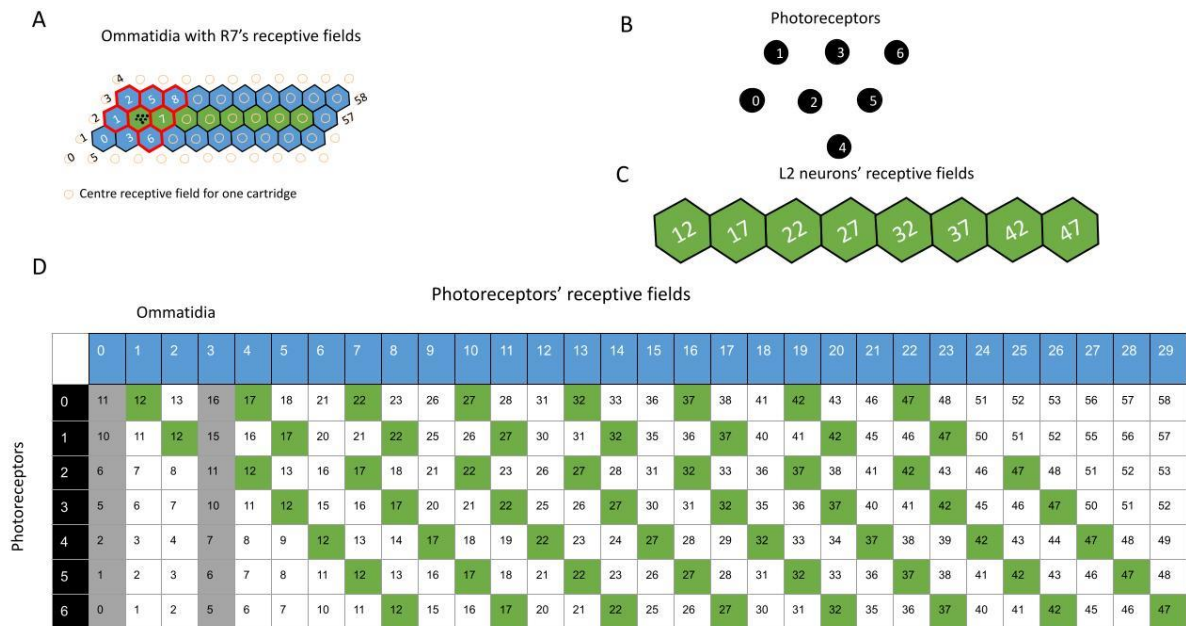
Figure 27: Matrix describing the photoreceptors' receptive fields in the ommatidia.
(A) Orange circles: Receptive fields organised in a hexagonal lattice numbered from 0 to 58. Blue and green: ommatidia placed at their R7 photoreceptor receptive field location, numbered from 0 to 29. Green: ommatidia for which the R7 photoreceptor feeds a recorded L2 neuron. Red: all the ommatidia containing a photoreceptor that feeds the most-left L2 neuron. **(B)** Photoreceptor positions within one ommatidium, ordered from 0 to 6. **(C)** The number of the recorded L2 neurons' receptive fields. **(D)** The matrix describes the receptive field number for each photoreceptor (row) of each ommatidium (column).

Note that this 7-photoreceptor-per-ommatidium model requires a $\frac{7\times28}{8} = 24.5$ times more computation than the 1-photoreceptor model. Since the computation of a contraction is heavy, I decided to reduce the amount of them by only taking into account the photoreceptors that feed the recorded L2 neurons (leaving aside the photoreceptors of the same ommatidia). This simplification makes the 7-photoreceptor-per-ommatidium model require only 7 times more computation.

The output of the photoreceptors was considered to be the dot product of the input and the filter.

3.2.2 Lamina: recurrent neural network

The photoreceptors' output feeds a 'reservoir type' recurrent neural network. The network's output is the value of the L2 neurons, which are part of the reservoir. The network is trained on the recorded L2 terminals' calcium signals. Based on the lamina connectome

(Rivera-Alba et al., 2011), 7 neurons are connected to both the photoreceptors and the L2 neurons (**Figure 28**). Hence, I chose the number of neurons equal to $7 \times N_{L2}$ where N_{L2} is the number of imaged L2 neurons, which also corresponds to the number of cartridges. Five different artificial connectomes, including the real-like connectome, are proposed below. For every connectome, I add self-connections that I assume are relevant since neurons' activities depend on their past activities.
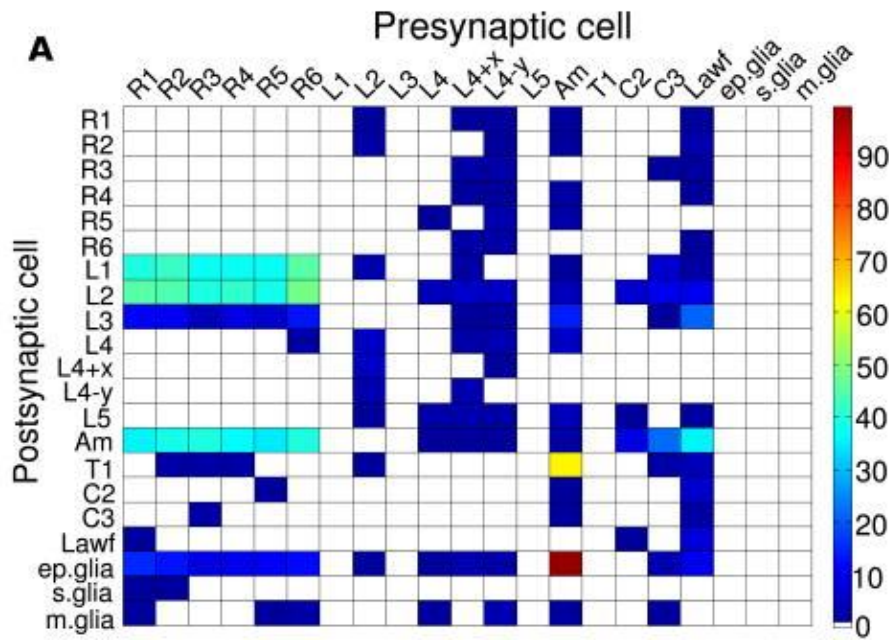
Figure 28: Number of synaptic connections between neurons in the lamina.
(A) Heatmap representing the number of synapses between presynaptic cells (column) and postsynaptic cells (row). Data from Rivera-Alba et al., 2011.

For each connectome C, instead of showing the whole adjacency matrices, I present the graph representation $\mathcal{G}(C)$ of their canonical part. This contains the intra-cartridge connections, the postsynaptic connections to the left (or right) neighbouring cartridge and the presynaptic connections from the right (or left) neighbouring cartridge. $\mathcal{G}(C)$ contains the sufficient and necessary information to describe the whole connectome C since each cartridge is a replicate of its neighbours. However, for the sake of simplicity, I will not show the self-connections since they are all present in every connectome. Moreover, I will not show the large wide-field cells' connections since they output on every neuron, which would make the graphs less digestible.

C_{real} is based on the detailed lamina connectome (Rivera-Alba et al., 2011), describing the number of synapses between the presynaptic and postsynaptic neurons. In this publication, the connectivity is described in 2 dimensions to infer it across the whole fly

eye. For our case, 2-photon imaging only enables monitoring of a lamina surface section corresponding to one line. Hence, to avoid any ambiguity regarding the L4+x and L4-y inter-cartridge connections in **Figure 28**, I considered two possible 'real-like' connectomes: C_{real}^L and C_{real}^R where the "+x" inter-cartridge connection goes respectively towards the left only and towards the right only. For the sake of simplicity, I considered one artificial photoreceptor, instead of the six rhabdomeres (R1-R6), that regroups all the pre/postsynaptic connections to/from. I have:

$\mathcal{G}(C_{real}^L)$:
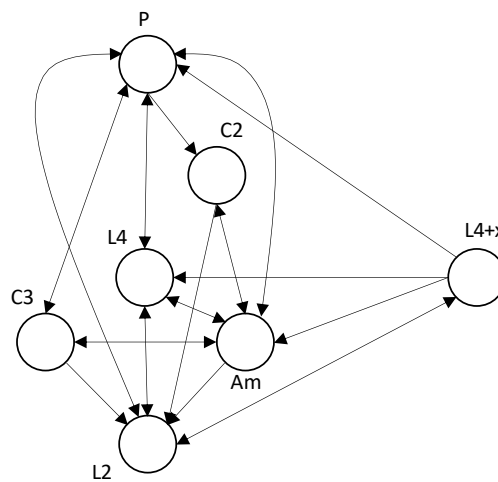
C_{random} is a stochastic connectome where the number of connections is the same as in C_{real} and the input and output neurons are randomly shuffled. In other words, each intra-cartridge-connection is a random permutation of its corresponding adjacency matrix. In addition, each lateral cartridge connection is a random permutation of the presynaptic connections to L4+x and the postsynaptic connections from L4+x combined.

$C_{columns}$ is a cartridge connectome with synaptic connections between every neuron and without inter-cartridge connections.

$\mathcal{G}(C_{column})$:
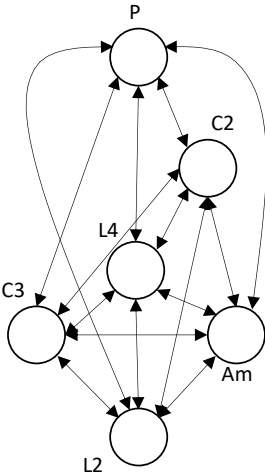
C_{dense} is the connectome where every neuron of every cartridge is connected to each other.

$\mathcal{G}(C_{dense})$:
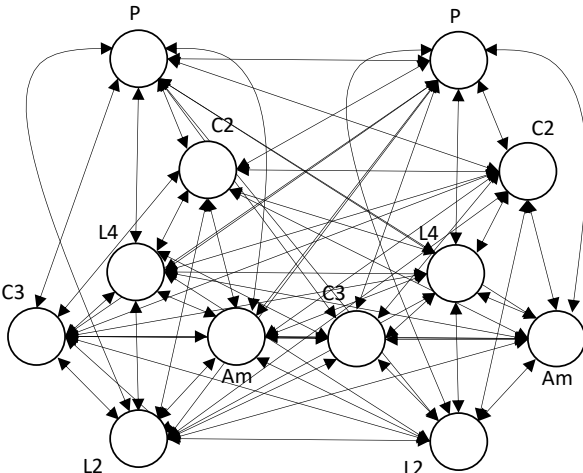
The BNN-like ANN, proposed to reproduce the recorded calcium responses at the L2 axon terminals, is the combination of the microsaccadic photoreceptor layer and the C_{real} lamina layer.

3.2.3 Training dataset

Content: Out of the 50+ flies I had imaged the L2 neurons' calcium activity, I selected the fly whose responses met several important criteria: healthy preparation, high signal-to-noise ratio, a high number of neurons imaged, clear preferred directions, and clear hyperacuity. The dataset used for training this model contains two-photon calcium time series responses of this transgenic fly showing activity at the L2 neuron (accentuated light-OFF responses in the lamina layer - 2^{nd} layer - of the fly's visual pathway) axon terminals located in the medulla (3^{rd} layer). It comprises 32 moving grating stimuli that have a narrowing wavelength over time. The stimuli span 16 regularly distributed orientations at two speeds (20 and 30°/s).

Scaling: I first subtracted the recording plane's background from each neuron's raw fluorescence intensity variations. I then subtracted the basal fluorescence, F_0, calculated as the mean intensity before the visual stimulation, from the observed intensity, F, ($\Delta F = F - F_0$) and giving this difference as the relative fluorescence change ($\Delta F/F_0$). To prevent weight saturation (Y. A. Lecun et al., 2012), I scaled the latter value between 0 and 1. Note that scaling, unlike normalising, enables having a control on the extreme values which might change significantly over time because of the unavoidable signal-to-noise decrease (see Discussion).

Flattening: Finally, because of calcium-indicator induced baseline raise (mostly happening at the beginning of each recording), I flattened the baseline by subtracting each value to a moving minimum with a window of 1 s. This procedure should make it easier for the model to learn the neural responses. **Figure 29** shows an example of relative fluorescence pre-processing (normalisation followed by baseline flattening). Note that this technique aims to mimic the input given to a calcium indicator filter.
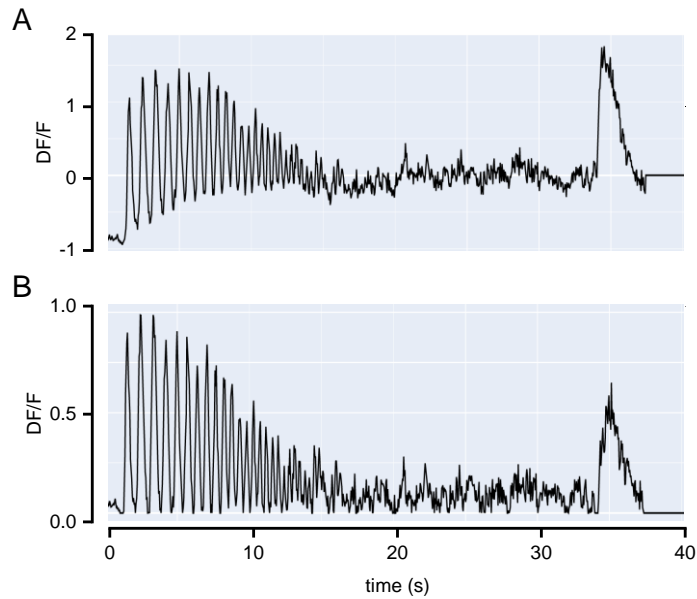
Figure 29: Data pre-processing fed to the model as training data.
(A) Relative fluorescence of an L2 axon terminal as stimulated by a dynamic wavelength grating (**Figure 30**D). **(B)** Trace after normalisation followed by baseline flattening (1s-window moving minimum).

3.2.4 Hyper-parameters

Initialisation: The weights' initial values follow a uniform distribution between -0.05 and 0.05, multiplied with the mask of the connectome used (1 [or 0] if there is [or isn't] a connection).

Loss function: The loss function, which the training process aims to reduce using the optimiser below, is the Euclidean distance between the model output and the experimental calcium response of the L2 axon terminals.

Optimiser: As a well-studied optimiser, I use the Adam optimiser for the ANN's learning rule.

Activation functions: For the photoreceptor layer, the screen input is normalised (or only centred around zero for the microsaccadic sampling) and then passed into a sigmoid function. The activation function for the lamina layer (the reservoir) is the non-linear rectified Linear Unit commonly used in machine learning.

Batch size: When trained on the whole dataset of 32 recordings, I used batches of 4 recordings.

Noise: For ANNs, in which the photoreceptor layer is a fully connected layer or a microsaccade-less multivariate gaussian filter, I added gaussian white noise ($\sigma = 0.02$).

Number of neurons: As explained above, I chose 56 neurons (7 neurons per cartridge, for 8 cartridges).

Trainable parameters: The parameters that the model optimises to reduce the loss function are the followings:

- The photoreceptor output scale that balances the photoreceptor output for a coherent input to the reservoir model.
- The model output scale that balances the model output before passing it to the loss function.
- The microsaccade two-dimension vector that tells where the photoreceptor moves.
- The cartridge's centre receptive field location.
- The reservoir's connection weights and biases.

Recurrent state: This recurrent neural network has an internal state, which is a list of 5 tensors (note that only the first tensor is used as an internal state for the static model):

- A tensor ([Batch x 56]) represents the values of the neurons in the reservoir.
- A tensor ([Batch x 8 x 2]) or ([Batch x 30 x 2]) represents the position and speed of the rhabdomere with respect to the corresponding ommatidial lens (the second dimension changes depending on whether I consider one or 7 photoreceptor(s) per ommatidium).
- A tensor ([Batch x 8 x 2]) or ([Batch x 56 x 2]) represents two characteristics of the photoreceptor microsaccade model. The 1st is the number of photons entering the photoreceptor. The 2nd is the number of microvilli that have absorbed a photon since the beginning of the recording (the second dimension changes depending on whether I consider one or 7 photoreceptor(s) per ommatidium).
- A tensor ([Batch x 8 x Recording_length x 2]) represents two characteristics of the photoreceptor microsaccade model that need storage across the entire recording. The 1st is the photoreceptors' after-latency activation times. The 2nd is the photoreceptors' refractoriness end times.

- A tensor ([Batch x 1]) represents the frame index used for the previous two tensors to update the photoreceptor numbers at the end of their activation latencies or refractoriness.

Output: The model output is the value of the artificial L2 neuron, multiplied by the trainable output_scale parameter.

3.3 Results

L2 neurons' Ca^{2+}-activity was measured at their medulla axon terminal levels (Kemppainen et al., 2022). A head-fixed fly saw moving dynamic-wavelength gratings while I recorded its resulting L2 Ca^{2+}-dynamics with a 2-photon imaging system. The recordings revealed direction-specific hyperacuity.

I made an ANN model trained to reproduce these responses. This model consisted of two compartments: a photoreceptor layer and a lamina layer. The photoreceptor layer was inspired by the actual fly eye anatomy and physiology (**Figure 30**AB). The light sampling was photomechanical (Kemppainen et al., 2022): photoreceptors generated microsaccades (moving and contracting underneath the ommatidial lens) and showed light adaptation and refractoriness. This layer outputs to the lamina layer, which is a bio-constrained ANN. Its connectome was based on the number of synapses within one lamina cartridge layer and the lateral inter-cartridge connections (Rivera-Alba et al., 2011)y (**Figure 30**C). I used this model and its several variants (e.g. naïve connections,

microsaccade-less photoreceptors etc.) to assess the usefulness of different connectome aspects in reproducing the recorded responses, including hyperacuity (**Figure 30**D).
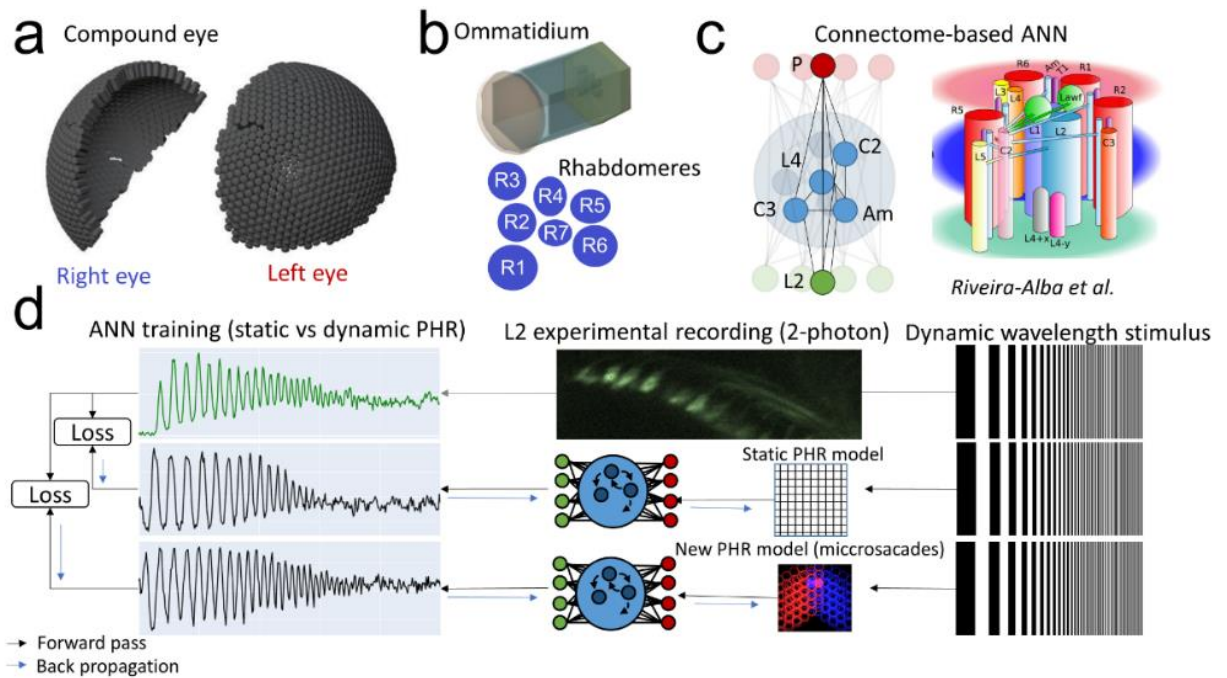

Figure 30: Comparing connectome-based ANN with static vs dynamic photoreceptor inputs. **(A)** The fly compound eyes were used as visual sampling constraints for the model. **(B)** A zoom-in on the ommatidium (top) and rhabdomeres (bottom). **(C)** The connectome-based ANN with one artificial cartridge shown (Left). The red unit (P) was connected to the inputs. The green unit (L2) was the output. The blue units formed a reservoir, in which connections were based on the real lamina connectome (Right). **(D)** Backpropagation through time of the dynamic/static photoreceptor ANN. The training was done on the L2 neurons#39 (see Chapter 2); Ca^{2+}-imaging data.

3.3.1 The static case - BNN-like ANN is better suited to reproduce the experimental data

I first studied the static sampling models, asking whether the BNN-like ANN reproduces the experimental data better than other more naïve ANNs or ANNs used in computer vision. To answer this question, I trained ten different static sampling models: two possible photoreceptor layers (either a naive fully connected layer or a microsaccade-less gaussian filter layer); and five possible lamina layers (C_{real}^R, C_{real}^L, C_{random}, C_{column}, C_{dense}). I call training the process during which a model changes its parameters following an optimizer method in order to reduce the loss function (see Material and methods section). In **Figure 31**, I show some examples of how the loss value evolves during training. We can see that 20 epochs (which we constantly used) is enough for a complete training, as the loss stops reducing after that. To analyse the resulting models, I used two metrics to evaluate each model's capacity to reproduce the experimental data. One was the root-mean-square error

(RMSE) between the visual acuity prediction of the model and the visual acuity measured on the experimental data. The other metric was the mean error between the preferred orientation prediction of the model and the preferred orientation inferred from the experimental data. On both metrics, the results suggested that dense lamina layers best reproduce the experimentally observed acuity characteristics (**Figure 32**). This finding was expected since this model's number of parameters was eight times larger.
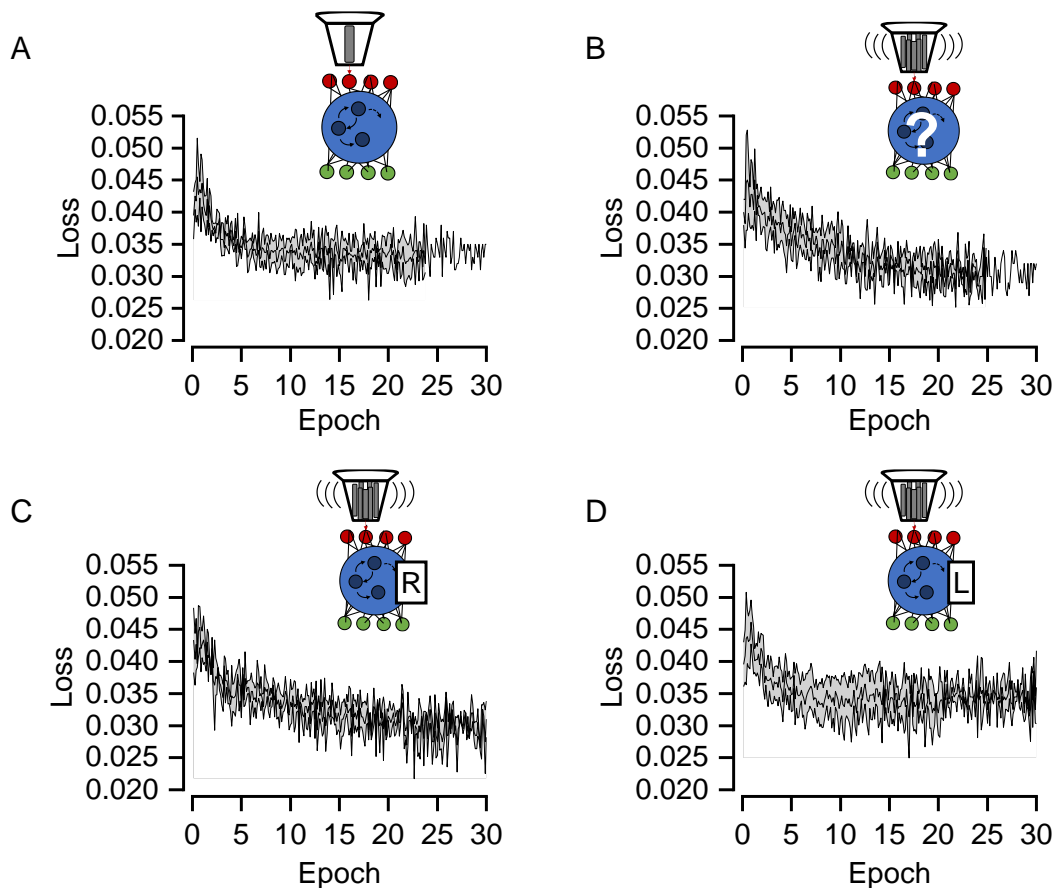
Figure 31: Loss is minimal after 20 epochs.
(A), **(B)**, **(C)**, **(D)**: loss (mean square error between real and predicted L2 fluorescence) evolution for static (n=4), C_{random} (n=5), C_{real}^{R} (n=3), and C_{real}^{L} (n=7) models, respectively.

The results also suggested that the lamina layer with random connectivity is the worst to reproduce the experimental visual acuities. More interestingly, microsaccade-less Gaussian filter photoreceptor layers generally approximated the experimentally observed acuity characteristics better than the naive microsaccade-less fully connected photoreceptor layer. This finding implies that for a static eye (or imaging system), Gaussian sampling at the photoreceptor level will likely be the mechanism by which a fully connected layer will converge. Strikingly, this difference in performance is even significant when the

lamina layer is C_{real}^{R} or C_{real}^{L}, which shows that the fly connectome is specifically tuned to receive inputs from a Gaussian sampling.
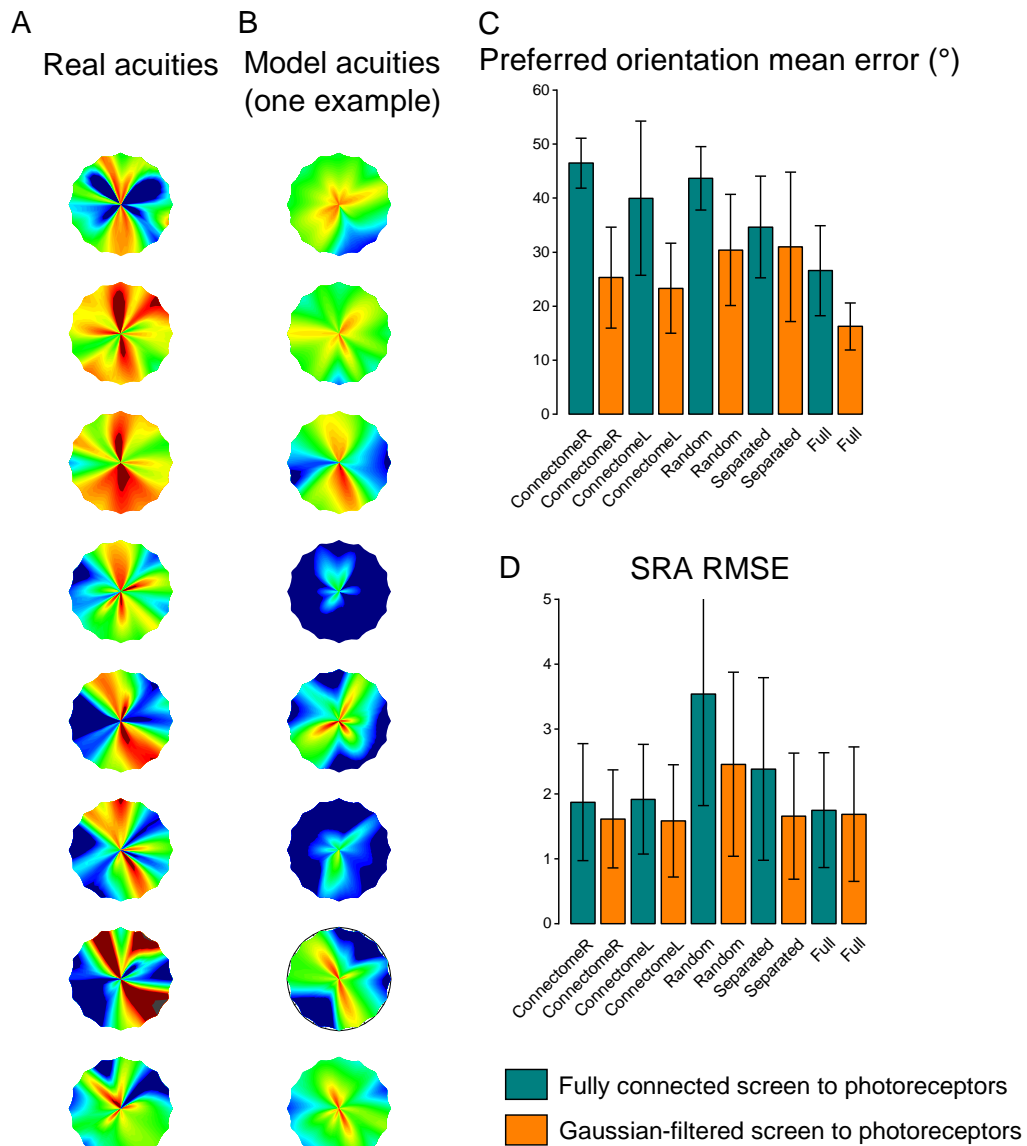
Figure 32: Static models' acuity predictions.
(A) Experimental acuity maps of the 8 L2 axon terminals from one fly, chosen for the model training set. **(B)** Acuity maps were predicted by a static-Gaussian input model with a real-like connectome. **(C)** Absolute mean error of the predicted preferred orientations compared to the real preferred orientations. Each bar corresponds to a different Lamina connectome. Blue bars: fully connected screen to photoreceptors. Orange bars: gaussian inputs to photoreceptors.

3.3.2 The photoreceptors microsaccades help the small fly connectome to reproduce single stimulus L2 responses.

I next asked whether the photoreceptor microsaccades had any utility for the ANN to reproduce the experimental data, and if so, which ANNs? Notably, it has been suggested that microsaccades would help to resolve moving objects (Juusola et al., 2017). To answer these questions, I used this time a dynamic gaussian which contracts, translates (learnable parameter) and shows refractoriness, in a similar way as in Kemppainen et al., 2022 (See **Figure 33**, and Material and methods section). I trained microsaccadic photoreceptor ANNs having three possible lamina connectomes (C_{real}^R, C_{column}, C_{dense}). Each ANN was trained on a single stimulus, which consisted of a dynamic grating moving towards a specific direction at 20deg/s. For each ANN, I trained 16 models (each model being trained with one stimulus moving towards one of the 16 directions used for the global dataset training in the models above).
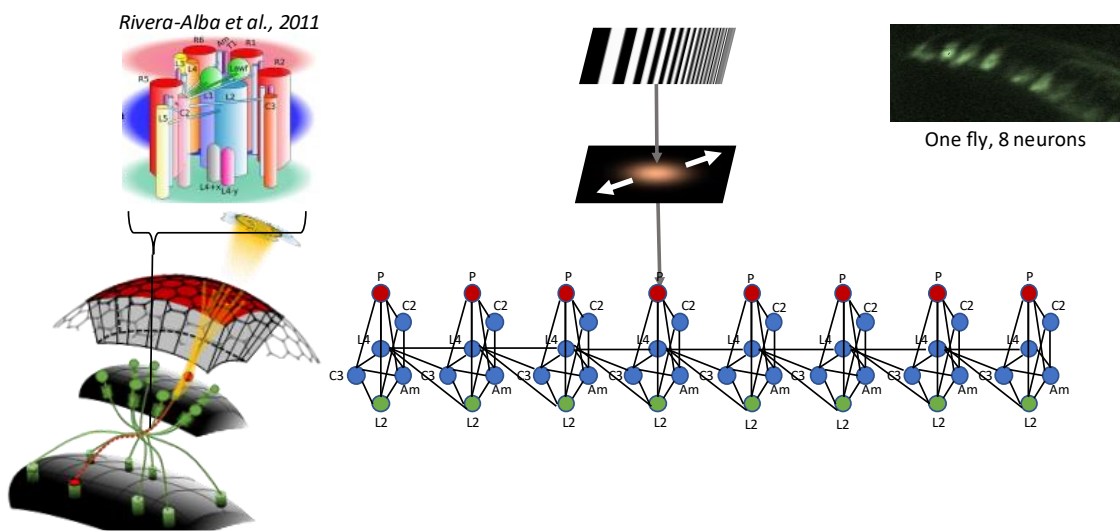
Figure 33: The bio-inspired ANN: A saccadic-gaussians layer feeds eight cartridges.
Layout of the final Bio-inspired ANN model of the photoreceptor/lamina layers. The visual input (top-centre) feeds a dynamic gaussian filter that mimics the biological fly micro-saccades. The output of this layer feeds then the artificial neural network mentioned above and described in details in the Material and methods section. Note that it contains eight cartridges, since it was trained on a fly on which eight L2-terminals were recorded (top-left).

I found that the ANNs' microsaccades tend to converge to the direction of the moving stimulus the ANN was trained on (**Figure 34**). More interestingly, I found that the microsaccades converged more consistently ($S = 0.60$, $R^2 = 0.31$, S being the slope of the linear fitting) to this direction when the lamina connectome is C_{real}^R (BNN-like ANN) (**Figure 34**A). A strong correlation also appeared with the connectome C_{column} ($S = 0.53$, $R^2 =$

0.28) (**Figure 34**B). On the contrary, dense lamina layer ANNs' trained microsaccades had almost a random direction ($S = 0.02$, $R^2 \ll 0.01$) (**Figure 34**C).

One possible explanation of these results is that the microsaccades helped resolve fine patterns by going towards the main direction of their movement. However, when the number of parameters was large enough (C_{dense}), the ANN did not need the help of these microsaccades. Hence, those results showed that photoreceptor microsaccades could play a role in fine-pattern resolving in an efficient, energy-saving way.
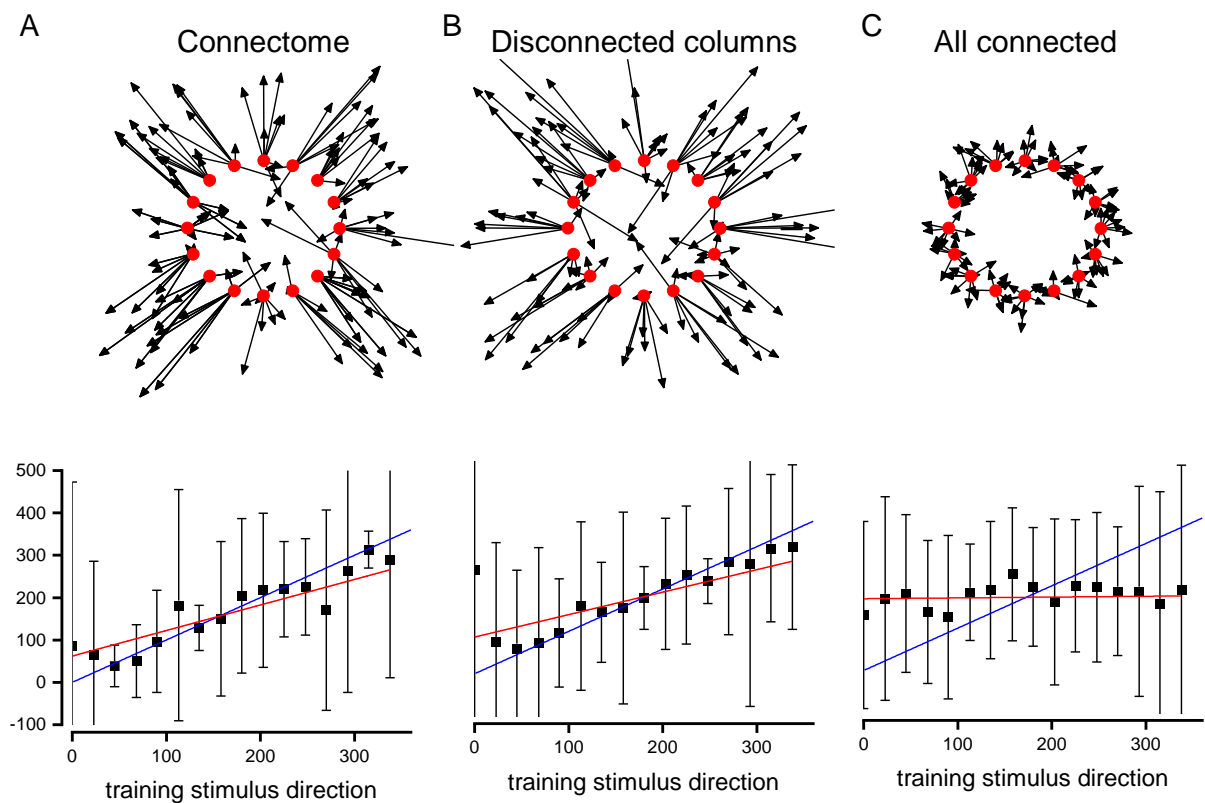
Figure 34: Single-stimulus trained microsaccade directions.
(A) Trained microsaccades for a model with a real-like connectome. Top: Each red dot corresponds to a model trained on a grating moving to the direction of the position of the red dot (relatively to the centre of the red dots). Each arrow starting from a red dot is a trained microsaccade of a photoreceptor of the corresponding model. Bottom: Average (across the eight L2 neuron microsaccades) direction of the trained microsaccades vs the direction of the stimulus used for training. Error bars: standard deviation. Red line: linear fitting ($R^2 = 0.31$). **(B)** Trained microsaccades for a model with disconnected columns. Linear fitting: $R^2 = 0.28$ **(C)** Trained microsaccades for a model with inter-and-intra-connected neurons. Linear fitting: $R^2 \ll 0$. **(A), (B) and (C)** Blue line: y=x.

Hence with much fewer parameters than the naïve ANN, this BNN-like ANN can reproduce well the experimental data.

3.3.3 The real photoreceptor microsaccades help reproduce the real L2 responses.

Kemppainen et al., 2022 measured the photoreceptor microsaccades across the whole fly eye. These are mirror-symmetric between the left and right eye and follow the optic flow of a flying fly (**Figure 35**A). I asked whether such microsaccadic sampling play a role in higher-order visual information encoding. To do so, I trained the microsaccadic-photoreceptors model to reproduce the experimental L2 responses to stimuli going in every chosen direction at 20 and 30°/s and let the model learn the microsaccade directions.

I found that the model predicted microsaccades were very similar to the real ones (**Figure 35**BC). To quantify this, I compared the results with random microsaccades and found the difference significant (**Figure 35**D). Hence, I can conclude that the microsaccades (as measured experimentally across the fly eye) play a role in the L2 axon terminal level activity.

For the rest of this chapter, I will only deal with models whose connectome is the BNN-like C_{real}^{L}.

3.3.4 The trained photoreceptor microsaccades match best the real microsaccades on the BNN-like ANN.

Interestingly, BNN-like models' predicted microsaccades are more similar to real ones than models with randomised connections and separate cartridges. Typically, there was no significant difference between the disconnected-cartridge models' predicted microsaccades and the random ones. These results suggest that the measured photoreceptor microsaccades are used in a process involving cartridge interconnections (**Figure 35**D).
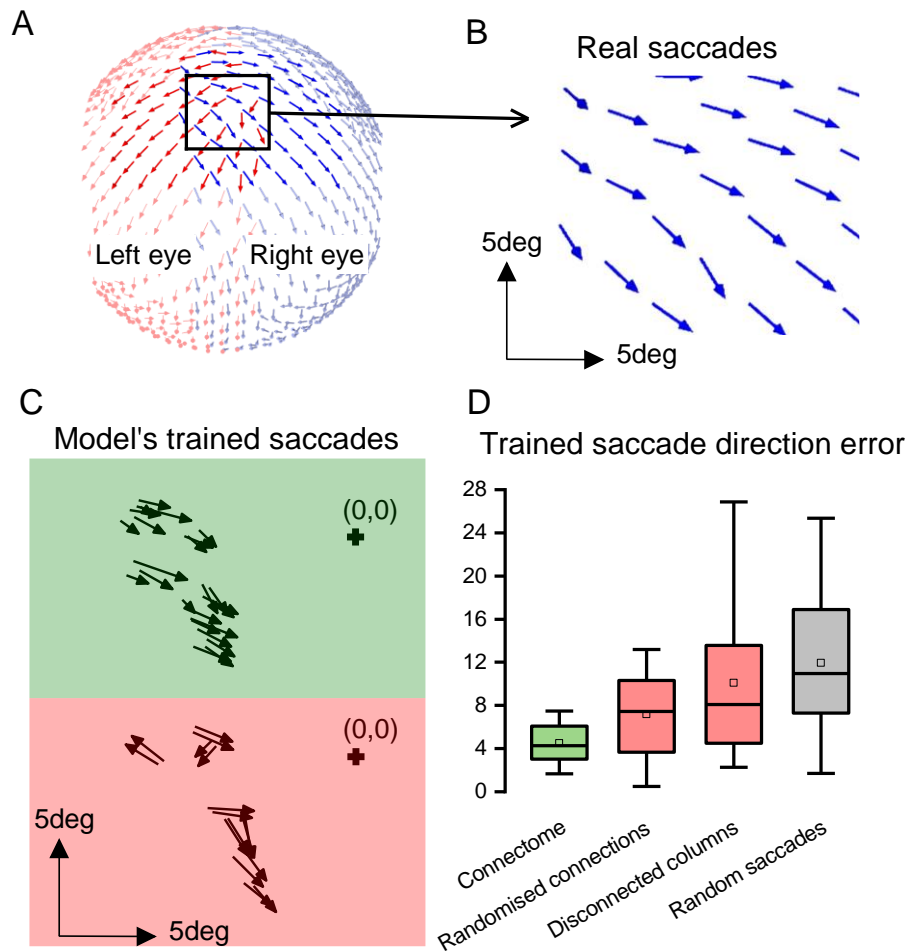
Figure 35: Comparison of the trained microsaccades with the real microsaccades.
 (A) Photoreceptor microsaccade directions across the whole fly eye. **(B)** Microsaccades of the right eye's photoreceptors in which receptive fields are in the black rectangle of **(A)** (centre of the screen). **(C)** Predicted microsaccades of the BNN-like model (green) and the disconnected-columns model (red). The null point refers to the centre of the screen. **(D)** Root-mean-square errors between the trained microsaccade vector and the real microsaccade vectors at the corresponding receptive fields. Green: BNN-like model. Red: Random-connections/Disconnected-columns model. Gray: Random vectors.

I then asked whether the model can reproduce the hyperacuity and orientational selectivity of the experimental recordings and whether these estimates worsen when the microsaccades are blocked. I also assessed the 7-photoreceptor model on the same criteria, hypothesising that the ommatidial 7-photoreceptor structure might play a role in hyperacuity. I trained each model twice (two different initialisations) on the whole 32-recording dataset (apart from the 7-photoreceptor model, which was trained only once because of its prolonged processing time).

3.3.5 The multi-photoreceptor model reproduces better the experimental acuity

I found that the single-photoreceptor models (with and without microsaccades) generate L2-responses that have similar acuities regardless of the speed and direction of the input grating, performing very differently than the real lamina BNN calcium signals. On the contrary, the 7-photoreceptor model's L2-outputs have a broader acuity spectrum, which appears to be more similar to the experimental acuities (fig. **Figure 36**A). This characteristic did not change when fine-tuning (i.e. training again the pre-trained models) on only the hyperacute part of the dynamic grating as input (**Figure 36**B).

Additionally, this model could generate nearly as acute responses as the most acute responses measured experimentally (on the fly used for training these models), as shown in **Figure 36**C.

3.3.6 Microsaccades help to produce hyperacuity

When comparing the microsaccade-less and microsaccadic models, I found that about half of the neurons did not show hyperacuity when the microsaccades were blocked. However, when microsaccades were allowed, almost all the neurons showed hyperacuity (see **Figure 36**C). In particular, when fine-tuning the models on hyperacute data only, every L2 neuron of the microsaccadic models showed hyperacuity.

3.3.7 Microsaccade-less models struggle to reproduce the experimental orientation specificity

Models with blocked microsaccades (i.e. static sampling) showed poorer acuity than microsaccadic models (i.e. dynamic sampling) , and their highest-acuity orientations were very unaligned with the experimental orientations (**Figure 36**E). The other models showed the highest-acuity orientations slightly closer to the experimental orientations but without significance.

3.3.8 Multi-photoreceptor models show naturalistic orientation-specificity when fine-tuned

Surprisingly, fine-tuned models gave some unexpected results. Firstly, the acuity was worse for the 7-photoreceptor model when fine-tuned, while its orientation-specificity settled very close to the experimental orientations (**Figure 36**DF). I suspect the model had found a better overall combination of its parameters (**Figure 41**), representing the experimental data more naturally (hence the very realistic orientation-specificity). Still, more training would be needed to show hyperacuity as well. Secondly, while the single-

photoreceptor microsaccadic model has a better acuity after fine-tuning, the orientation-specificity gets further from the experimental orientations.

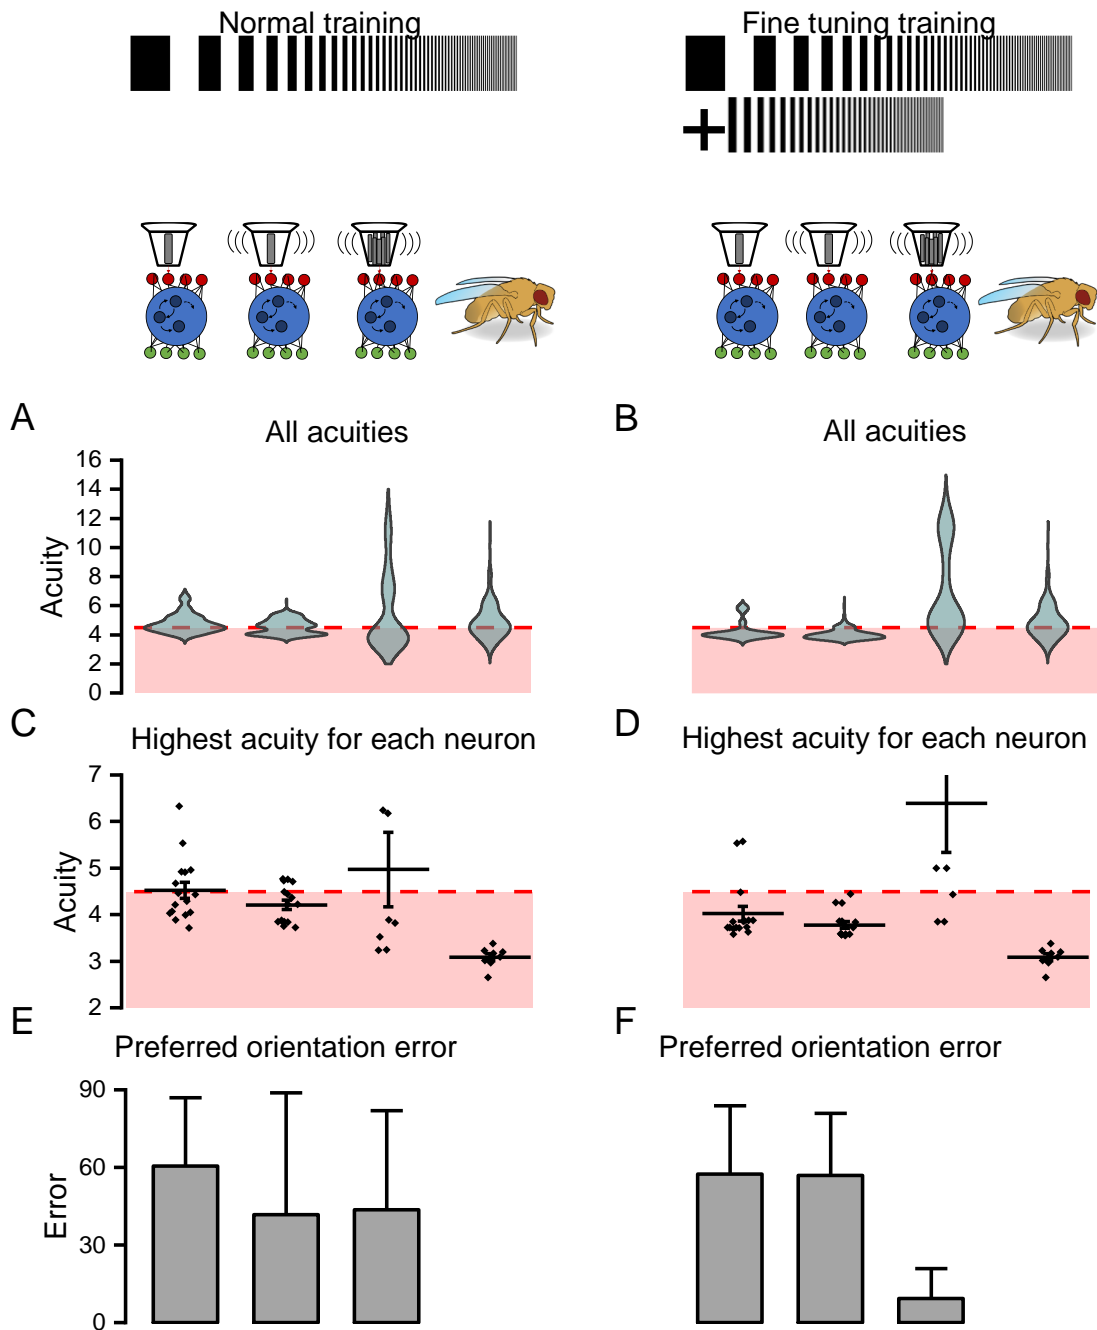**Figure 36: Role of microsaccades and neural-superposition to reproduce the experimental acuities and orientation-specificities.**
 Left column: models trained on the 32-dynamic-grating dataset. Right column: models of the left column after an extra training (fine-tuning) on the hyperacute part (only) of the dynamic gratings. **(A,B)** Acuity distributions of all the responses for all the single-photoreceptor-saccade-less, single-

photoreceptor-saccadic and 7-photoreceptor-saccadic models (ordered from left to right). Red dotted line: hyperacuity limit (4.5°). Light-red region: Hyperacute region. **(C,D)** Same as above, but showing the highest acuity for each model's L2 neuron. Horizontal bar: mean. Error bar: standard deviation. **(E,F)** Absolute differences between the experimental highest-acuity orientations and the models'. Error bars: standard deviation.

3.3.9 Multi-photoreceptor model's responses are shaped realistically

Since the multi-photoreceptor model is the only model capable of showing naturally (i.e. without additive fine-tuning training) similar acuities as the real acuities (**Figure 36**C), I asked whether the shape of this model's response resembled more accurately the real responses. To answer this question, I compared the responses of each of the three models to the first three bars of the dynamic gratings going to the four cardinal directions (**Figure 37**ABC) with the real responses measured experimentally (**Figure 37**D). I chose these first peaks explicitly as their shapes are easier to distinguish (as they are the biggest). I found that:

- When the model is more realistic (**Figure 37** top to bottom), the peaks show a more pointy shape, similar to the actual response shape.
- The model detects the bar stimuli earlier.
- The model's peak responses match more accurately the experimentally recorded real peak times. I quantified these time differences in **Figure 38**.
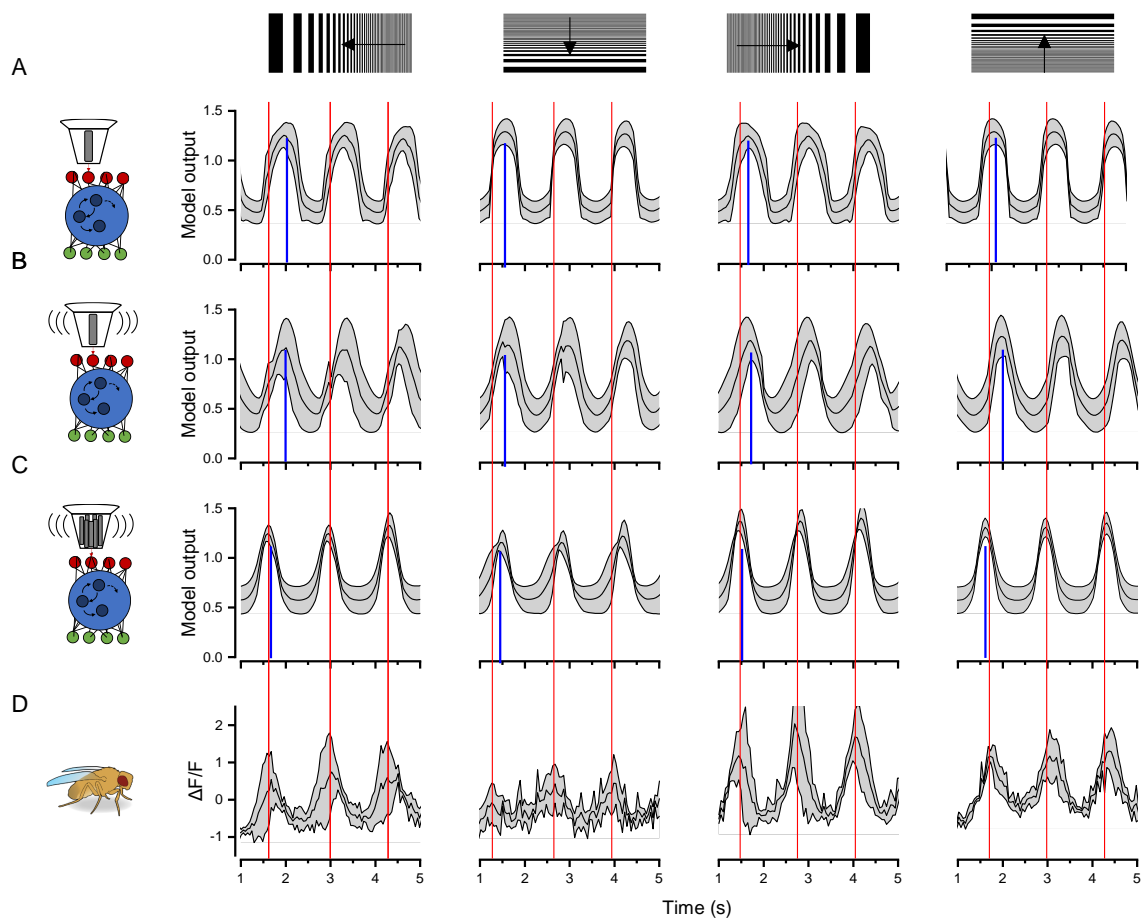
Figure 37: The 7-photoreceptor model predicts the most accurate shape of L2-response.
(A) Single photoreceptor static model responses to seconds 1 to 5 of the dynamic grating stimulus going towards 4 different directions (shown at the top). The line is the mean response of four L2 neurons. The margin shows the standard deviation. **(B) and (C)** Same as (A) but for the single photoreceptor microsaccadic and 7-photoreceptor microsaccadic models, respectively. **(D)** Calcium fluorescence variation of the actual fly to the same stimuli. Red lines: times of experimental peaks in the responses. Blue bars: times of the first peak in the models' responses.

Relative difference between model and real peaks
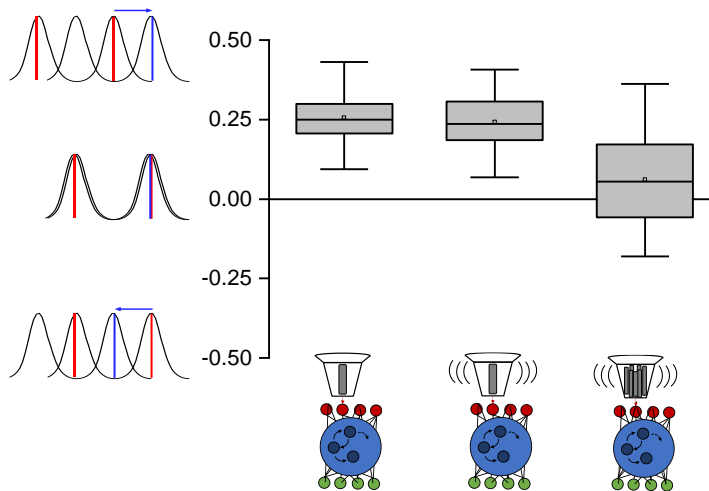
Figure 38: The peaks in the 7-photoreceptor model match most accurately the experimental peaks.
Distributions of relative (to the local wavelength) differences between model output peaks and experimentally observed peaks shown for the single-photoreceptor static (left), saccadic (middle) and 7-photoreceptor models (right). Here, the first three peaks for each response are examined. Positive/negative values mean the peak is late/early compared to the actual peak. Box: quartiles. Line: Median. Dot: Mean. Bar: Standard deviation.

These results show microsaccades of multiple photoreceptors (cf. Chapter 1) in neural superposition accentuating LMC responses. Moreover, photoreceptor movements can be induced by microsaccades of other photoreceptors sharing the same ommatidium but not the same receptive field (**Figure 25**). Thus, some photoreceptors can move even before receiving light input.

3.3.10 The BNN-inspired ANN can predict synapses strengths and types

The model learnt the ANN's connection weights during training. Their values were real numbers, which can be as large or small as needed, negative or positive (see the Material and Methods section). Hence, the model learnt each connection's strength and type (excitatory or inhibitory). I allowed the ANN to learn independent values across different cartridges and synapses of the same neuron. Then, I measured the resulting learnt weights after averaging them across 8 models sharing the same connectome (although different photoreceptor models) to see if some weights consistently converged to the same type and strength. To do so, **Figure 39**A shows the averaged connection weights with their coefficient of variation (the latter value gives an idea of the relative variations of each

weight value across all models). As the lamina is retinotopically (and its cartridges stereotypically) structured, I would expect connections to be the same from one cartridge to another (when ignoring any side-flank effects). Hence, **Figure 39**B shows the average of cartridges across all cartridges of all models.
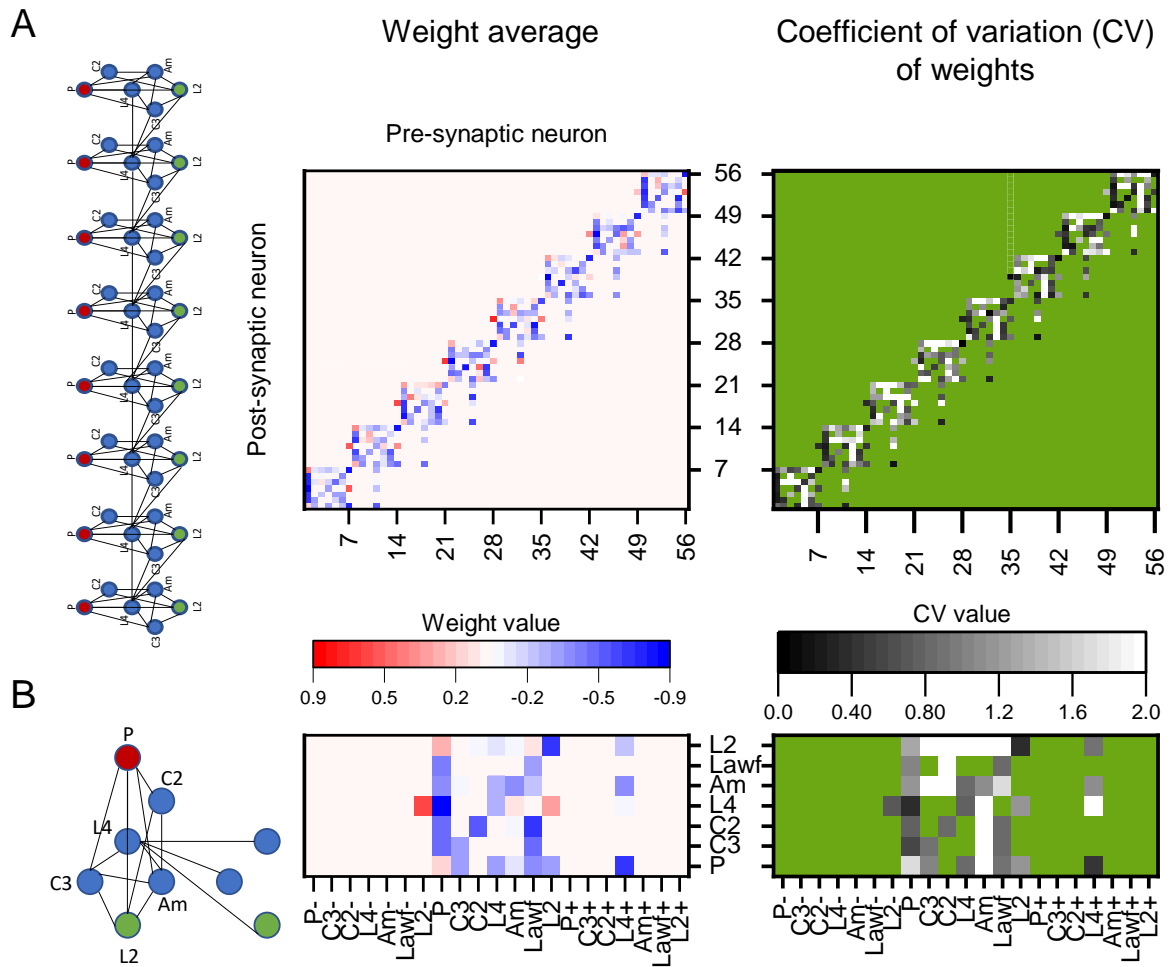
Figure 39: Learnt weights across all models.
(A) Left: schematic of the 8 cartridges with the connections used in the ANN. Middle: Connection matrix of the ANNs weights averaged across all the models. Right: Connection matrix of the ANNs weights' coefficient of variation across all the models. Each matrix's coordinates (modulo 7) represent the neuron name when assigned following this order (P, C3, C2, L4, Am, Lawf, L2). **(B)** Same as above but calculated across single cartridges.

3.3.11 Synaptic type prediction

Neurons communicate by exciting or inhibiting each other, with both functions needed to balance the network's synaptic load. **Table 1** summarises the *Drosophila* lamina neurons' neurotransmitters (considered in the current literature). Histaminergic photoreceptors and GABAergic C2 and C3 neurons inhibit their postsynaptic neurons, while L4 and the Large Wide-Field neurons release the excitatory acetylcholine (Kolodziejczyk et al., 2008). L2 and L4 release Glutamate, which can be either excitatory or inhibitory. However, current evidence implies that Glutamate is excitatory in the early visual pathway (Deng et al., 2019; Zheng et al., 2006).

I asked whether the BNN-based ANN models showed synaptic connection signs (or polarity) consistent with the known *Drosophila* neurotransmitters. As I allowed the network to learn the connection weights and signs for each synapse, some of the final weight signs differed across different cartridges and different postsynaptic neurons. Thus, I compared real neurotransmitter types after averaging the signs of the postsynaptic connection weights of each neuron learnt across all BNN-based ANN models.

Generally, the model weights seemed primarily inhibitory (**Figure 39** shows mostly blue colours). L2 monopolar cells were the only ones with positive postsynaptic connections. This finding is consistent with their Glutaminergic excitatory feedback (**Table 1**). However, L4 and the Large Wide-Field neurons showed negative postsynaptic weight signs that did not match the real transmitters released by those cells. Additionally, the model makes Am and C2 postsynaptic weights converge to zero, giving less importance to these cells. The caveat here is that 2-photon Ca^{2+}-imaging can only track relatively slow (integrative) synaptic events. Thus, this method is unlikely to resolve fast (derivative) synaptic dynamics accurately, including the polarity of signals driving the dominant voltage transients (Zheng et al., 2006, 2009).

Figure 15 shows that nearly all connection weights consistently converged to somewhat similar values across different model initialisations. In particular, 3 connections were learnt with strong consistency:

- Photoreceptors to L4
- L2 to L4+ (i.e. L4 of the neighbouring cartridge)
- L4+ to the photoreceptors

In contrast, significant variations after convergence were seen in:

- The connections from/to Amacrine cells
- Almost all the connections to L2
- The connection from C2 to the Large Wide-Field neurons

Overall, the BNN-based ANN model found combinations of connecting weights that differed from the real connection types (with some additional high-variation convergences). Hence, this finding indicates that multiple combinations of neurotransmitters can generate responses similar to the L2 responses to the dynamic gratings, as long as the photoreceptor outputs and L2 feedback match the actual neurotransmitters.

Table 1: Neurotransmitter predictions.
The prediction is the average postsynaptic weight from **Figure 39**B. Second column. References: Kolodziejczyk et al., 2008;Deng et al., 2019; Zheng et al., 2006.

	Neurotransmitter	Type			Model prediction
P	Histamine	Inhibitory			Inhibitory
C3	GABA	Inhibitory			Inhibitory
C2	GABA	Inhibitory			Weak
L4	Acetylcholine	Excitatory			Inhibitory
Am	Glutamate	Excitatory			Weak
Lawf	Acetylcholine	Excitatory			Inhibitory
L2	Glutamate	Excitatory			Excitatory

3.3.12 Synaptic strength prediction

A more accurate and quantifiable way to compare the learnt connectome with the actual connectome is to compare the relative ANN weight values within one cartridge with the relative number of synapses available (Rivera-Alba et al., 2011). **Figure 40** compares both, attributing light colours to large absolute weight values or many synapses for the ANN and real connectome.

Overall, the connections the ANN values the most after training seem to be very different from the actual number of synapses in the real connectome. For example, among the connections going to the ANN's L2 neuron, the photoreceptor inputs were by far the strongest, whilst the rest of the neurons had much weaker connections (**Figure 40**B).

However, the model learnt the importance of the Large Wide-field neurons and the neighbouring L4, as in the real connectome (**Figure 40**BC).
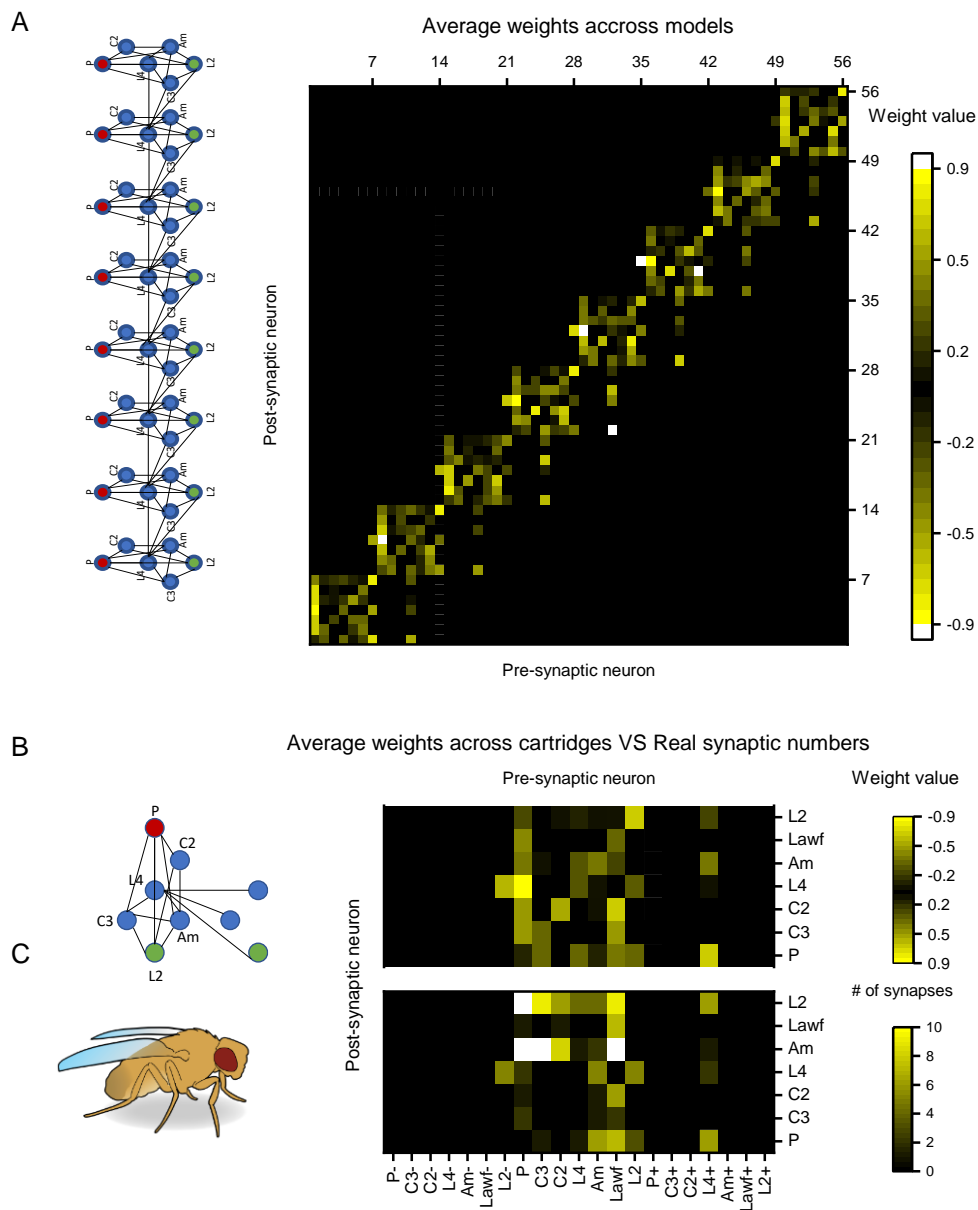
Figure 40: Learnt weights' strengths compared with the real connectome.
(A) The connection matrix of the ANNs weights averaged across all the models, for which the colour shows the absolute weight value. Each matrix's coordinates (modulo 7) represent the neuron name when assigned following this order (P, C3, C2, L4, Am, Lawf, L2). **(B)** Same as above but calculated across single cartridges. **(C)** The number of synapses in the lamina. Data retrieved from the Material and Methods section of Rivera-Alba et al., 2011.

In **Figure 39** and **Figure 40**, we can see the lateral L4 connections' side-flank effects and the photoreceptors' postsynaptic connections. Indeed, on the sides, the lateral L4 connections (particularly L4+ to photoreceptors and L2- to L4) were weaker than in the

centre, whereas the photoreceptors generated stronger responses. Hence, L4 seemed essential to reproduce the experimental data (as long as the corresponding cartridge is encapsulated between other cartridges).

3.3.13 The role of Amacrine cells is particularly well described with the multi-photoreceptor model

Finally, I asked whether the learnt weights varied between different photoreceptor model types. I used the same 3 different photoreceptor models as the ones used in **Figure 41**. I found that the weights learnt are very similar between the two single-photoreceptor models (**Figure 41**BD). However, the 7-photoreceptor model learnt slightly different weights, most of them being closer to the number of synapses we observe in the actual connectome (**Figure 41**FH). For example, Amacrine cells seem to have a lot of presynaptic connections, which the latter model seemed to reveal more than the other two models. The 7-photoreceptor model is also the only one that gives a higher value to the lateral L4 to L4 connections, as it happens to be in the real connectome. The single-photoreceptor models both showed a stronger lateral L2 to L4 connection, which is also true in real flies.
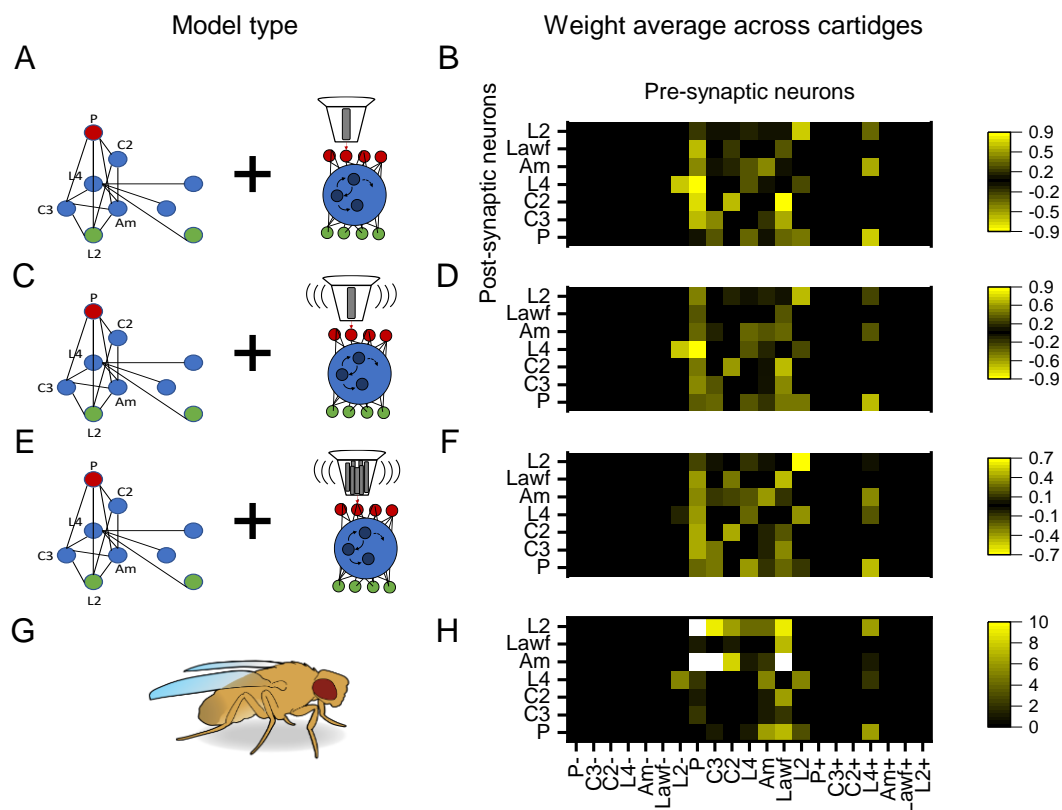
Figure 41: Comparison of the learnt weights between different photoreceptor models.
(A) Single-photoreceptor-saccade-less model. **(B)** Connection matrix of the weight values averaged across single cartridges, for which the colour shows the absolute value. **(C)** Single-photoreceptor-

81

saccadic model. **(D)** Same as in (B) but using the models in (C). **(E)** 7-photoreceptor-saccadic model. **(F)** Same as in (B) but using the models in (E). **(G)** Experimental data. **(H)** The number of synapses in the lamina. Data retrieved from the Material and Methods section of Rivera-Alba et al., 2011.

Overall, the BNN-like 7-photoreceptor model best reproduces the experimental acuity, orientation specificities and synaptic distribution when trained on the experimental L2 responses to dynamic gratings.

3.4 Discussion

Initially, I made a straightforward model of L2 neurons using a weighted sum of increasing and decreasing exponentials to try and reproduce the experimental responses of the L2 axon-terminals.

- $L2(t) = \sum_{i=1}^{t} 2f_i\left[e^{-(t-i)\lambda_d} - e^{-(t-i)\lambda_r}\right],$

- Where $f_i = \left(P(i) - P(i-1)\right)\frac{(1+s)d+1-s}{2},$

- With $s = sign(P(i) - P(i-1))$

I had several problems related to this model. Firstly, the responses seemed to contain a low-frequency sinusoid which I do not have in the real L2 responses. Secondly, the resolvability dropped much faster than for real L2 neurons. I thought that using photoreceptors' receptive field motion could improve the results. But our current findings invalidated this idea. These discrepancies were the primary reasons I used a BNN-inspired ANN to model the data.

Even though the final model follows mechanisms that are far more similar to the actual fly brain workings compared to the previous photoreceptor/lamina models, I have still simplified much of what is observed in the real world. First, the photoreceptor model I used is inspired by the photomechanical microsaccadic movement model (Kemppainen et al., 2022), which, for instance, involves 3D Gaussians receptive fields instead of 2D Gaussians used here. Second, I did not consider the 8 photoreceptors inside each ommatidium. The reason for this was that R7 and R8 have the same receptive fields and are mostly feeding the colour-involved channel. Third, the ANN I use to model the lamina layer also strongly simplifies the mechanisms occurring there. For example, the model did not consider the propagation time of information between two neurons, although we know that the fastest communication happens in gap junctions at ~120 m/s (Saint Marie & Carlson, 1985). Accordingly, R7 and R8 gap-junctions channel information to R1 and R6 photoreceptor axon terminals, accentuating their responses, just before the LMC synapses (Wardill et al., 2012). Finally, the model's hyper-parameters (such as activation functions, optimiser and weight initialisation) are the traditional ones used in current artificial intelligence (e.g. reLU and sigmoid, Ding et al., 2018). These simplifications mean that I must be careful about the conclusions I draw on the resulting weights when I compare them with the neurotransmitters and the number of synapses observed in the lamina.

However, those simplifications enabled faster model training and evaluations, making the model more usable and adaptable for neuroscience/AI-related purposes. Typically, we used 20 epochs (**Figure 31**) which took about 6 hours for the single-photoreceptor model, and 2 days for the 7-photoreceptor-model.

A simplification, which might have a big impact on the model outputs and overall conclusions, is that the junction between the photoreceptor model and the lamina model is made through the artificial neurons P (**Figure 30**D). In my model, one artificial neuron P gathers all the inputs from the photoreceptors. Hence, in the case of the 7-photoreceptor model, I lose some information (maybe crucial) by sending the aggregation of the 7 photoreceptors output passed through the P neuron's activation function instead of going directly to each of the connected lamina neurons. This simplification was done because of training constraints.

The models used in the first figures are much more simplified than the ones used in the last figures. This difference was because the latest models were not fully implemented at the beginning and were also very time-hungry when training adequately. Nevertheless, I kept the first figures as they brought about interesting insights to compare the different connectomes.

For the 7-photoreceptor model, I had to choose the initial photoreceptor receptive field positions based on a hexagonal lattice (see **Figure 24**B). Many issues arise from this constraint. Firstly, I chose a pattern that matched best the pattern on the measured positions of the photoreceptors inside every ommatidium of both eyes (**Figure 24**A). Because the recorded L2 axon terminals were next to one another, I assumed their connected photoreceptors would be in neighbouring ommatidia too. However, since the two-photon recording plane can only be a section of the L2 axon-terminal population, it is unclear whether the alignment is along the hexagon's long or short axis (see **Figure 42**). I went for the second option, as it seemed more intuitive to me, but also because it saved 50% of computation time for training and evaluation. Next, the experimental receptive field positions do not look like they are aligned (see **Figure 26**). Thus, I had to infer an approximation for the hexagonal lattice using a Principal Component Analysis (PCA) method. I understand that this technique is questionable as, for some flies, it creates a hexagonal lattice with a PCA direction far from expectations. An alternative could have been to use the average position of the experimental receptive fields and draw the lattice following the pseudo-pupil alignment mapping in **Figure 42**A. However, the receptive field

positions do not remain the same during training, so I considered those approximations to be fixed naturally during learning.
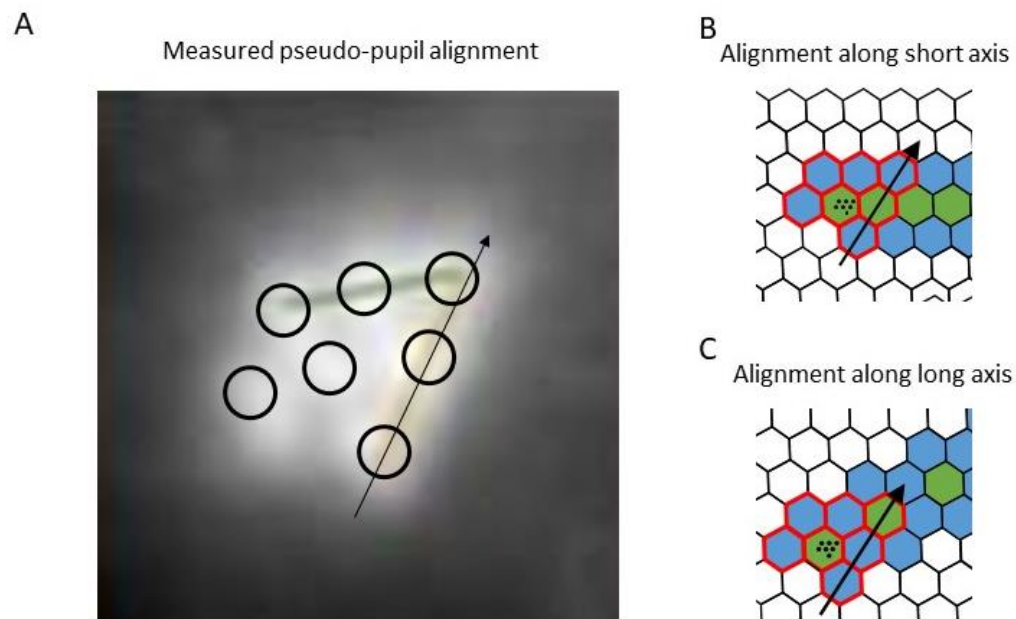
Figure 42: Multiple options for ommatidia-pattern choice.
 (A) Axis (black arrow) of the 7 neighbouring ommatidia feeding to the same cartridge, located near the recorded-L2 receptive field locations. Black circles: ommatidia pattern experimentally measured from the pseudo-pupil imaging (Kemppainen et al., 2022). **(B)** One option of ommatidia-pattern choice is that the alignment of the L2s' feeding ommatidia is along the short hexagon axis (option used in the model). **(C)** Another option is to align them along the long hexagon axis.

As an output for the model, I used the value of the artificial L2 neuron, normalised between 0 and a learnt positive value. Unfortunately, this procedure lacks the GCaMP Ca^{2+}-indicator component used when collecting the experimental L2 responses. Hence, it may create a gap that the model cannot fill when trying to reduce the loss. However, I showed in chapter 2 that the used GCaMP Ca^{2+}-indicator was fast enough to measure a fly's whole acuity range.

To model the lamina layer, I used an ANN, in which the connectome was similar to the real lamina (**Figure 28**) (Rivera-Alba et al., 2011). However, the lamina connectome is not an isolated closed system but should theoretically involve the whole fly brain. One has to draw the limit somewhere, and it could well be that more neurons in other optic lobes are strongly involved in the L2 activity (Wei et al., 2020). I used every neuron directly

connected to both the photoreceptors and the L2, as it is believed that indirectly connected neurons usually have only a weak influence on each other.

The ANN connectome follows the published description of the lamina layer (Rivera-Alba et al., 2011). However, one must remember that these connections have been inferred from a single cartridge. Hence, the model's connectome could be missing some contacts that, by chance, were missing on this single cartridge. In particular, lateral connections across distant cartridges are necessarily missed in this paper. Besides, the number of synapses (used as a final comparison in **Figure 40** and **Figure 41**) could vary from one cartridge to another (S. Y. Takemura et al., 2015).

When comparing the models' performances to reproduce the experimental data, one must remember that the number of parameters is not the same across all models. For instance, C_{dense} has a much higher amount of parameters than C_{real}. Typically, heavy models could struggle to converge to the global minima (e.g. higher chances to converge into local minima because of the higher dimensional parameter combination space, overfitting, etc.). Note that $C_{columns}$ also had more parameters than C_{real}. Also, the photoreceptor models had variations in their parameter numbers; for example, the microsaccade-less models lacked the trainable microsaccade vector, unlike the microsaccadic models. However, the number of parameters remained small (<400) compared to the dataset given to the model. The only times this issue can be significant is when I used the fully connected photoreceptor to the screen or when using C_{dense}, as the number of parameters is more than 1,800. To avoid any comparison issues between C_{real} and C_{random}, I attributed the same number of parameters C_{random} as there are in C_{real} (see Material and Methods section).

The dataset used for training the models was relatively small (only 32 recordings), and no repetitions were made for each stimulus (only one recording per stimulus). In contrast, the responses of the L2 terminals are subject to fast dynamic adaptation/processing (that can be [mis]interpreted as noise) and maybe some random processes, which make them imperfectly repeatable. Additionally, I used only two flies (but mainly one fly) for the training dataset. Hence, the model I created might be very specialised to the used dataset.

To avoid too strong data specialisation, I pre-processed the data before passing it on to the model. Two major modifications were the normalisation (in the sense of scaling between 0 and 1) and the flattening (see Material and Methods section). Firstly, the

variations in amplitude measured in the actual L2 responses depend strongly on the GCaMP fluorescence strength, which decreases over time. Because this phenomenon cannot be explained with the type of model I have, I normalised the data. Secondly, light adaptation and/or calcium indicator time-constant induce an initial substantial calcium rise in the L2 responses. Since our model was not designed to reproduce it, I artificially flattened the data. These two methods made the data more standardised and acuity-focused. However, we need to bear in mind that I might have lost some important L2-physiological properties through this process.

The batch size was chosen empirically, such that:

- the training process did not take too much time (which happens in the case of too large batch size)
- the model did not overfit (which happens in the case of too large batch size)
- the training loss value was not subject to strong fluctuations (which happens in the case of too small batch size)
- the training loss value did not diverge (which happens in the case of too small batch size)

(Y. A. Lecun et al., 2012; Pape et al., 2007)

Because the model is a recurrent neural network, its response depends on the inputs and a state (Jain & Medsker, 2000; see Material and Methods section). Since the state was empty at the beginning of the model evaluation, the initial response is incoherent. Hence, during the training process, the loss did not consider a warmup period corresponding to the first 0.04 s of the input time series.

All the models do not show OFF responses at the L2 level. Several factors might be involved in this particularity. One is that the receptive field locations are learnt during the training process. Thus, the model might sometimes learn to move the receptive field instead of changing the weights' polarities. Hence, I do not want to conclude with confidence that the sign of the resulting averaged weights of the models means that the model learnt neurotransmitters' types. Instead, I put more importance on the relative differences between the signs of the weights (see, for example, the explanation for **Figure 39**).

As seen in **Figure 39** and **Figure 40**, the model final weights (and probably output) are subject to side-flank effects caused by lateral interactions. Hence, it could have been more sensible to train the model on the central neurons only. To a larger extent, I can imagine that every model neuron is subject to side-flank effects along the axis perpendicular to the L2 alignment. Thus, one could think of adding surrounding neurons along both axes to avoid side-flank effects (the more, the fewer side-flank effects).

I aimed to have a realistic photoreceptor model to feed the lamina model because the lamina model's weight combination was likely to depend on it. Typically, when fed by an unrealistic photoreceptor model, the lamina model would not represent the actual lamina physiological structure. For example, the fully connected photoreceptor filters are very parameter-heavy, implying that, in reality, the photoreceptors do part of the jobs previously considered for the lamina only. For instance, the centre-surrounding mechanism in the lamina layer (Freifeld et al., 2013b) is explained with a 'Mexican hat' type of receptive field for fully connected photoreceptor filters (see **Figure 43**).
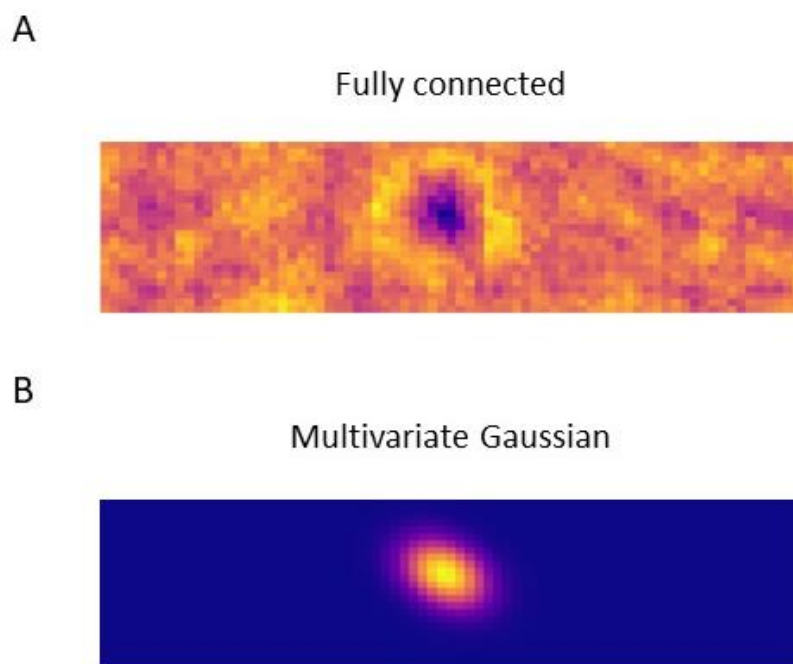
Figure 43: Fully connected photoreceptor layer learns centre-surrounding mechanism.
(A) An example of the learnt weights of a fully connected photoreceptor model that resembles a Mexican hat. **(B)** An example of a learnt multivariate-Gaussian photoreceptor model in which the centre-surrounding mechanism cannot be learnt.

More repetitions would need to be done (requiring maybe a more optimised algorithm for less time processing) to increase the confidence in the conclusions I made regarding the importance of the different static lamina models (**Figure 32**) and the significance of microsaccades for the final model (**Figure 36**). In particular, only one training was done for the 7-photoreceptor microsaccadic model.

The microsaccades of the two $C_{columns}$ models in **Figure 35** show very similar learnt directions, whereas the other models showed variations across trials. This unexpected result likely implies that only lateral connections enable multiple microsaccade combinations to reproduce the experimental data. In particular, the photoreceptor microsaccades influence the lateral cartridges as well.

A considerable proportion of the trained models (around 10%) showed high-frequency fluctuations in the responses (see **Figure 44**). I usually saw this phenomenon happening before divergence. However, some models stabilised whilst keeping these fluctuations in the responses. Since these strongly affected the resolvability analysis, I discarded these models.
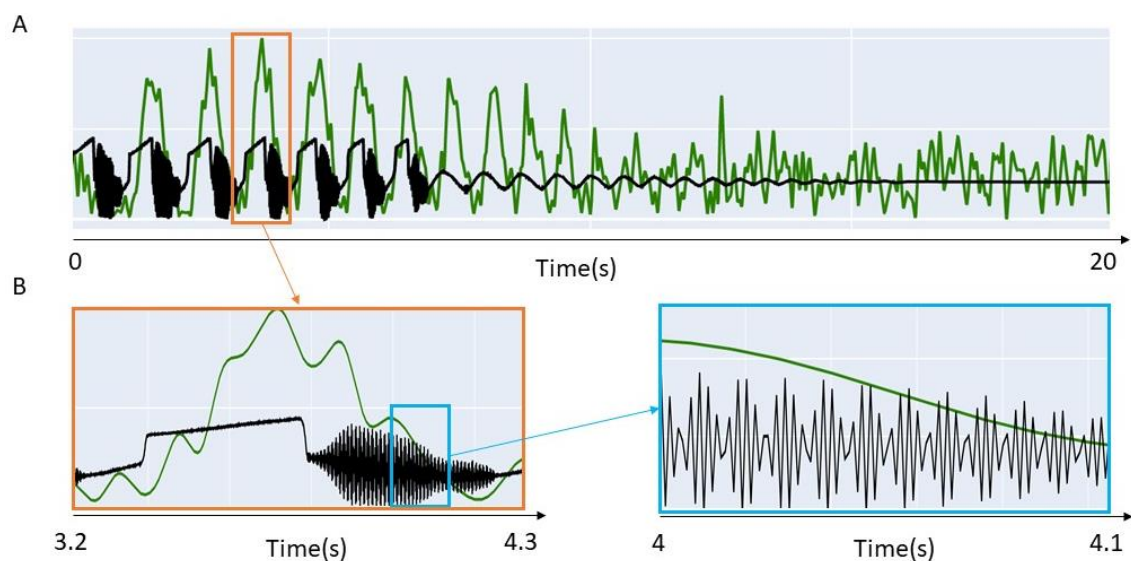
Figure 44: Example of fluctuations in a model's response.
(A) Response of a model (in black) to a dynamic-grating stimulus. Green traces: experimental data used for training. **(B)** close-ups of fluctuation periods. Y-axis is from 0 to 1 for A and B-left. Y-axis is from 0 to 0.5 for B-right.

The L4 asymmetrically connects lateral cartridges: it receives inputs from the L2 of a neighbouring cartridge (along an axis, say x) and from the L4 of the neighbouring cartridge (along a perpendicular axis, say y). It outputs to the photoreceptors (R1, R3, R4, R6), L2,

L4, and Am of the opposite neighbouring cartridge (along x) and to L4 of the opposite neighbouring cartridge (along y) (**Figure 45**). However, these are only relative axis. Because of difficulties in determining the orientation of the x and y axis corresponding to the recorded L2 cells, I arbitrarily chose only one axis (x) and two possible models: C_{real}^L for x, and C_{real}^R for opposite-x. Note that the Lamina connectome (Rivera-Alba et al., 2011) shows almost a 90° symmetry in the lateral L4 connections (particularly for the neurons used in our model). Thus, choosing the x or y axis should not matter.
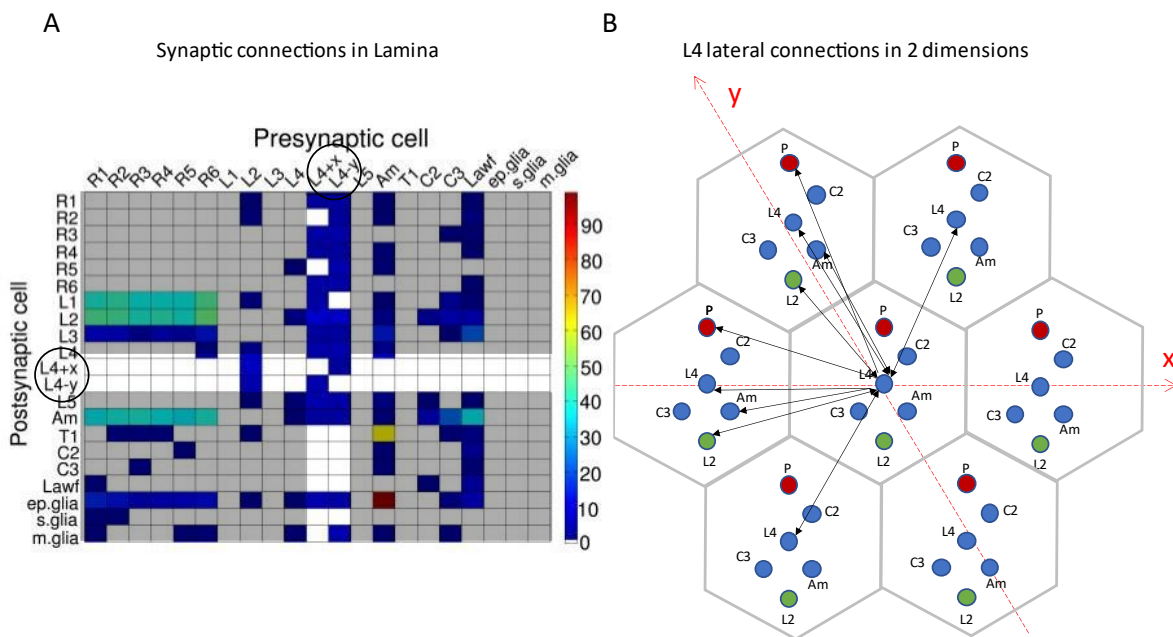
Figure 45: L4 lateral connections depend on axis orientations.
(A) Table of synaptic connections in the Lamina layer (Rivera-Alba et al., 2011), where L4 lateral connections are highlighted. **(B)** Representation in 2-dimension space of the L4 lateral connections. It displays only the connections coming from/to the L4 of the central cartridge. Red dotted arrows: arbitrarily chosen axis for x and y (used in table A).

I display below the weights learnt from different models (**Figure 46**).
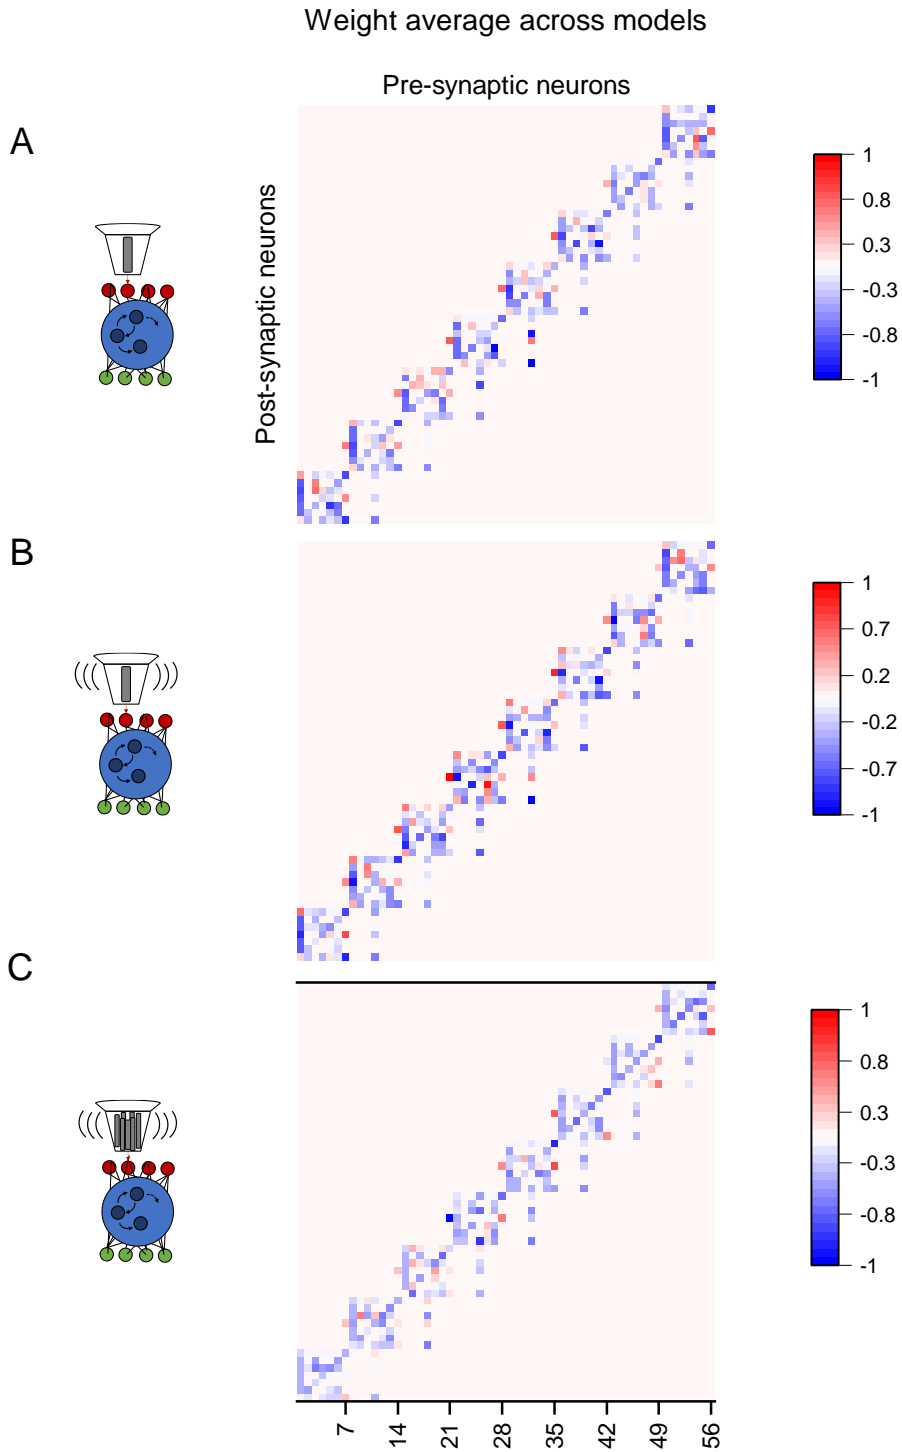
Figure 46: Learnt weights of the 3 different photoreceptor models.
 (A) Connection matrix of the weights averaged across all the single-photoreceptor-saccade-less models. **(B)** Same for the single-photoreceptor-saccadic models. **(C)** Same for the 7-photoreceptor-saccadic models. Each matrix's coordinates (modulo 7) represent the neuron name when assigned following this order (P, C3, C2, L4, Am, Lawf, L2).

To push the model to learn stronger hyperacuity, I also trained it on a fly which hyperacuity was one of the strongest I have measured during this PhD (1.96±0.13deg). In this case, I see that the single-photoreceptor models learn stronger acuity than the one measured in **Figure 36** (**Figure 47**). Yet, surprisingly, the 7-photoreceptor model's acuity drops heavily. One possible explanation is that I used the same hexagonal-lattice pattern in Material and Methods (**Figure 24**B, **Figure 27**A) when training the model on data pooled from the right eye of a fly (e.g. **Figure 47**). In contrast, this fly's recorded neurons were located in the left hemisphere. This simplification was done assuming that the model could learn the symmetry. As another unexpected result, when flipping the hexagonal-lattice pattern for the left hemisphere, the model always converged to the local minimum that corresponds to the averaged training data.
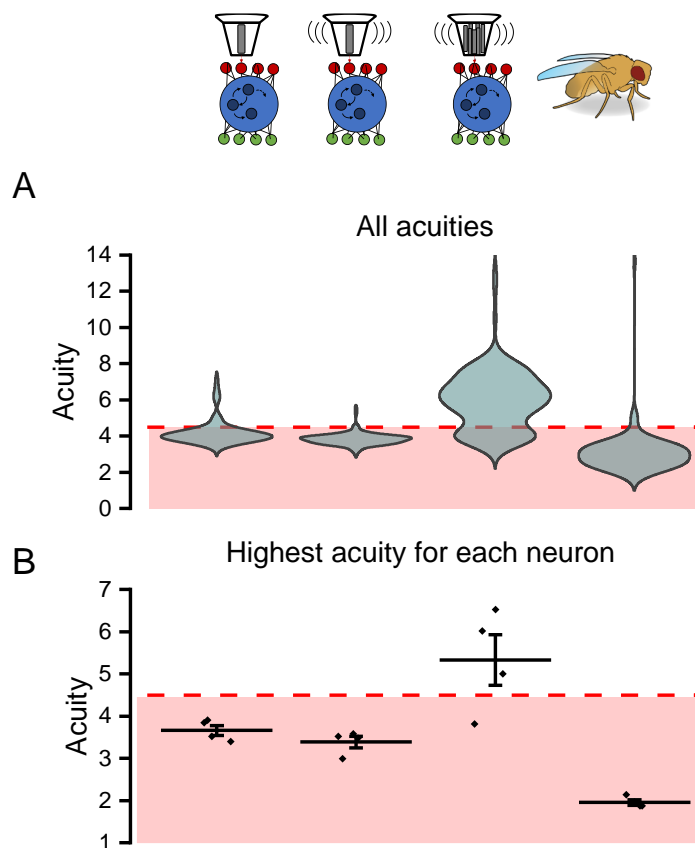
Figure 47: Visual acuity of the models trained on experimental L2 responses of a fly showing strong hyperacuity.
(A) The acuity distributions of all the responses for all the single-photoreceptor-saccade-less, single-photoreceptor-saccadic and 7-photoreceptor-saccadic models (ordered from left to right). Red dotted line: hyperacuity limit (4.5deg). Light-red region: Hyperacute region. **(B)** Same as above, but showing the highest acuity for each model's L2 neuron. Horizontal bar: mean. Error bar: standard deviation.

Notably, random connections impacted several aspects of model training (see **Figure 32** and). Firstly, for static models, trained acuities and orientations seem to be further from the experimental data when the connectome was random (**Figure 32**). Secondly, for microsaccadic random-connectome models, the learnt microsaccades were less similar to real microsaccades (**Figure 35**). Hence, one can ask if microsaccadic random-connectome models showed worse acuities and orientations. **Figure 48**B below shows one answer to this question, using *the reduced single photoreceptor and the multi-photoreceptor microsaccade model* as the test cases.

Interestingly, their acuity distributions did not differ from the real-connectome models. Additionally, the preferred orientations were not better with random-connectome models but slightly worse (**Figure 48**C). Overall, I did not see a significant effect in the microsaccadic model training when the connectome was randomised. Thus, this finding suggests that the high-precision visual encoding already happens at the level of photoreceptors.
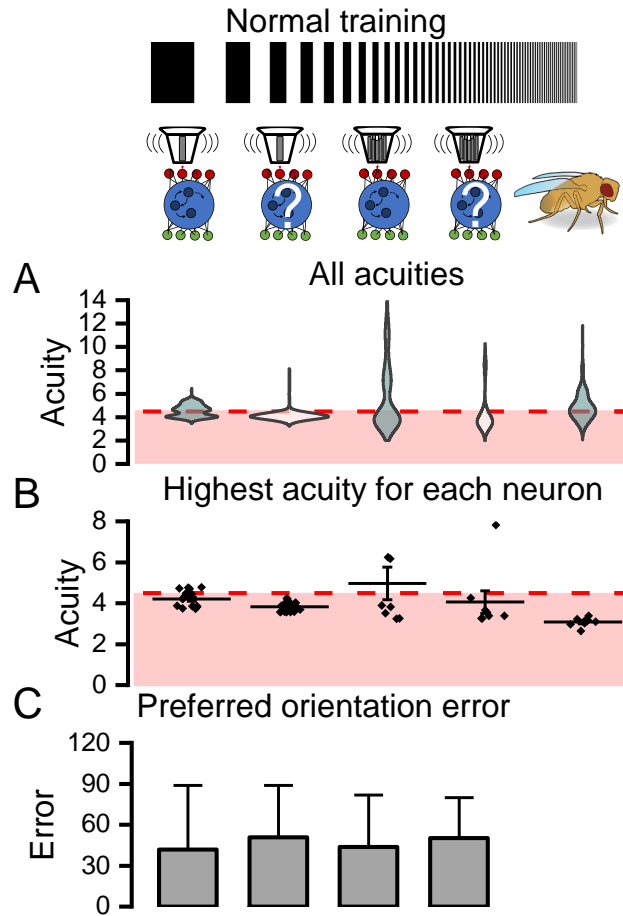
Figure 48: Random connections do not affect the model's acuities and preferred orientations.

(A) Acuity distributions of all the responses for the real-connectome and random-connectome single-photoreceptor and multi-photoreceptor saccadic models. Red dotted line: hyperacuity limit (4.5deg). Light-red region: Hyperacute region. **(B)** Same as above, but showing the highest acuity for each model's L2 neuron. Horizontal bar: mean. Error bar: standard deviation. **(C)** Absolute differences between the experimental highest-acuity orientations and the models'. Error bars: standard deviation.

Finally, the model built in this chapter can significantly improve the current AI used in computer vision. Indeed, this model makes a low-parameter (because constrained by the structure of the *Drosophila* visual pathway) representation of visual stimuli in a way that maintains the acuity at its best. Hence, this model can be seen as a low-parameter video compressor inspired by the fly's long-evolved mechanisms to encode its visual surroundings in space and time.

This thesis was funded by the EPSRC program grant 'Brains on Board' that aims to build artificial autonomous-low-energy drones using the new knowledge from insect vision research. This model can be integrated (after some tweaks) into this kind of prototype, especially because the compression is likely to be optimised for robots behaving like flies (https://brainsonboard.co.uk/).

Chapter 4: Hyperacuity in the visual pathway

4.1 Introduction

A handful of techniques exist that can improve understanding of how the brains work. In the mammalian model systems, for example, fMRI is widely used to record the activity of broad regions in the brain. However, for smaller brains, such as those of insects, recording techniques like electrophysiology enable more accurate activity estimates at the single neuron level. In the previous chapters, I showed how to use 2-photon Ca^{2+}-imaging to indirectly measure the activity of multiple individual neurons of the same type (L2 neurons), which were first genetically targeted to express a calcium indicator. This approach enabled me to infer some of their properties regarding visual encoding and orientation-sensitivity shifts between neighbouring neurons. But arguably, measuring the activity of a single type of neuron, even as an ensemble, restrains the amount of knowledge one can obtain. Indeed, a neuron cannot act independently in the network it is part of, and it is crucial to consider the whole network simultaneously as a multi-sensory combinatorial encoder to understand the underlying neural communication and computations.

In this chapter, I aim to unravel (in some detail) how the visual world is encoded throughout the whole visual pathway of the *Drosophila*, using a pan-neural expression of GCaMP. Specifically, I will trace how visual acuity information passes through different visual neuropils. I will first show that hyperacute communication occurs in every recorded layer in this fly strain. Secondly, I will present orientational and directional visual sensitivity results in the lamina and medulla circuits and compare these with the corresponding photoreceptor microsaccade directions, like in the previous chapter. I want to emphasise that this kind of pan-neural analysis is challenging and thus has been attempted here only in a limited sense. Issues such as neuron identification, neural activity mixing or functional alterations must be clarified in the future. These issues are discussed in the material and methods and discussion sections below.

4.2 Material and methods

I performed whole brain 2-photon Ca^{2+}-imaging in flies that express GCaMP6f across the entire brain (elav Gal4-line). Neural activity changes (fluorescence signals) to visual image motion were imaged using the same 2-photon microscope system and stimulus characteristics as for the L2 experiments above. The fly preparations were made as before, apart from the bigger cuticle opening, enabling the whole visual pathway imaging.

4.2.1 Stimulus

In the visual pathway, different neurons respond to visual stimuli differently. Hence, different stimulus/response analysis combinations will be relevant for different neurons to infer whether they show hyperacute responses to ultra-fine visual patterns. The dynamic-wavelength grating used for L2 acuity monitoring is well suited to monitor the acuity responses in many neurons in the visual pathway. It leads to local luminance changes (stimulating most of the lamina/medulla neurons) and local/global motion stimuli (stimulating most of the lobula/lobula plate) neurons. A coarse wavelength should trigger a significant response at each layer. In contrast, if the wavelength is small enough, it should trigger similar responses (noise and spontaneous activity) as a grey frame. Hence, visual acuity can be estimated by characterising the smallest wavelength triggering a response distinguishable from noise.

We chose accordingly a growing, dynamic wavelength, in which properties follow equation (11) of chapter 2, but with $\lambda_0 < \lambda_1$ (Figure 49). This way, the wavelength is multiplied by a constant factor frame after frame, enabling an accurate estimate of the point where the grating starts to induce a response in the visual pathway.
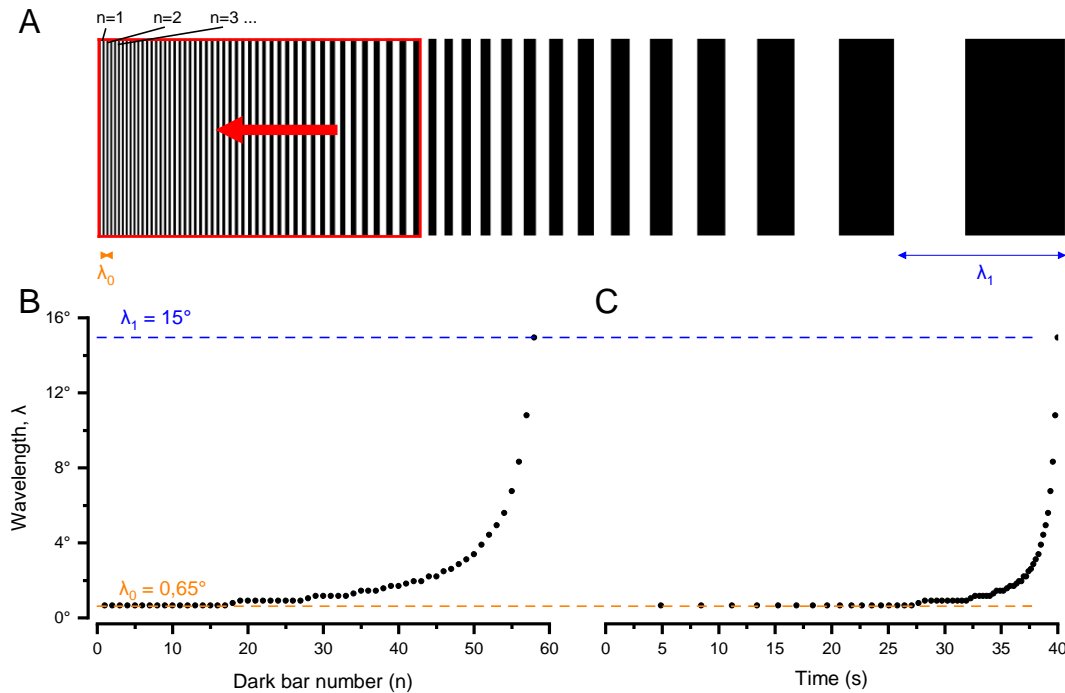
Figure 49: Description of the 4-parameters widening grating.
(A) Grating design. The wavelength widens from λ_0 to λ_1. Red rectangle: Screen (what is seen by the fly). Red arrow: Grating motion direction, speed is constant. **(B)** The wavelength at each dark bar. **(C)** Wavelength over time.

4.2.2 Neuron activity segmentation

Since the Elav flies express the calcium indicator in every brain cell, I built a Python algorithm to detect and classify neuronal activity across the recording plane. Like the L2 activity analysis, I tiled the recording plane with arbitrary-sized square ROIs (I chose 16x16 pixels squares, **Figure 50**EF).

For each ROI, the fluorescence intensity variations F were quantified after subtracting the background (mean intensity of a corner of the recording plane) from the mean intensity of the ROI. Ca^{2+} variations were then calculated according to the following equation:

$$\frac{\Delta F}{F} = \frac{F - F_0}{F},$$

eq. 4-1

where F_0 is the mean intensity before the visual stimulation.

I selected an arbitrary number of squares with the highest Ca^{2+}-variations standard deviations. I then grouped the ROI activities (**Figure 50C** and **E**) using a clustering

algorithm that utilises the function *TimeSeriesKMeans* of the Python package *tslearn.clustering*. The number of clusters chosen was 20. The maximum number of iterations was 50, and the metric used was Euclidean.

I monitored 250 ROIs' activities for each recording, clustered in 20 different groups. The stimuli were presented multiple times to the fly by varying the speed (usually $s =$ 20, 30, 60°/s) and the motion direction (usually every 15 or 30°).

4.2.3 Response analysis

For the sake of smoothness in the response, I constantly analysed the mean response for each cluster. The shape of the response differed regarding the type of cluster. As seen in the L2 experiments, this grating usually induced a succession of peaks in the lamina and medulla layer. For these responses, one can use the same method as for the L2 neurons to assess the time when the observed ROI starts resolving the grating. In the lobula and lobula plate layers, the responsive neurons raise calcium concentration as soon as the grating is resolved. However, for several reasons discussed at the end of this chapter, the ROI calcium variations in the lamina and medulla also followed an up-trend when the grating was resolved. Therefore, I built an algorithm in Python that detects the time of this rising. To do so, it first measures the response baseline b and its noise level n (standard deviation of the first 10 s). The algorithm outputs the last time the response is below $b + n$ (**Figure 50**C).
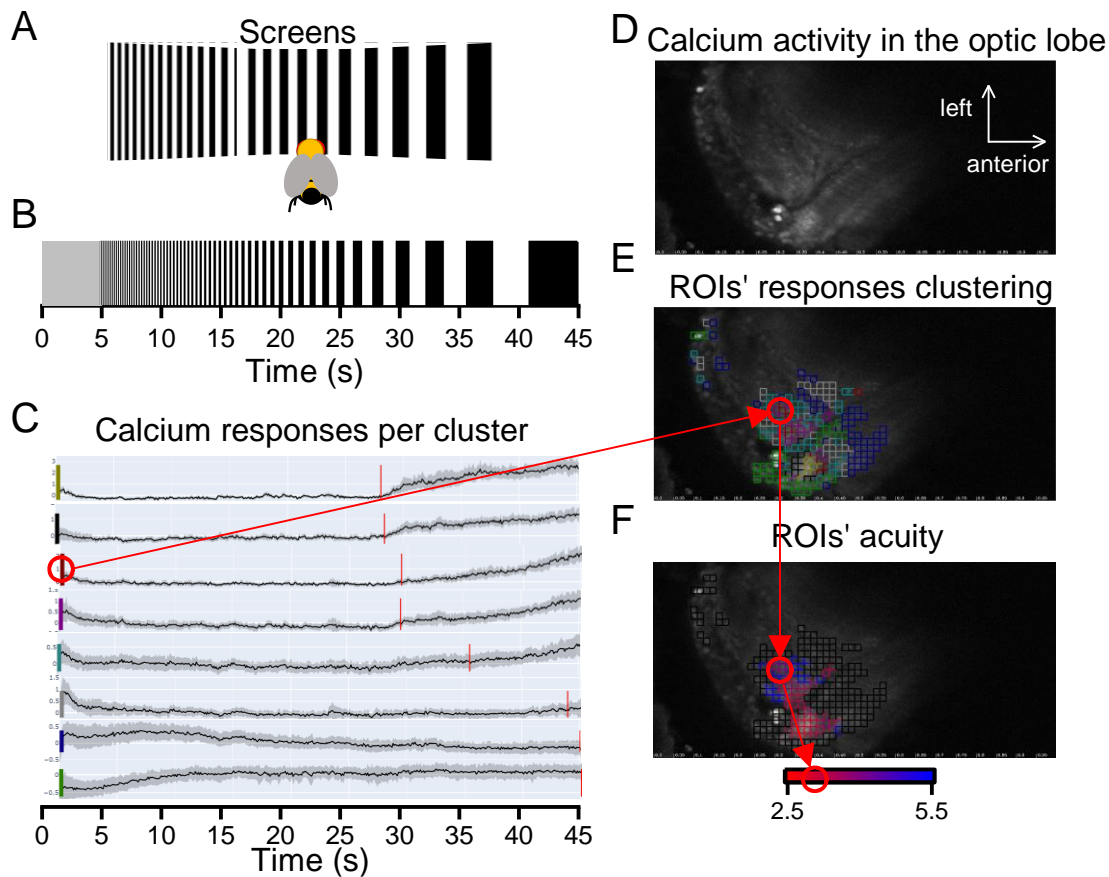
Figure 50: Acuity measurement across the optic lobe.
(A) Dynamic grating as seen by the fly through three screens spanning 150° x 50°. **(B)** Stimulus value over time. It starts with 5 s of the grey frame, followed by the widening grating explained in **Figure 49**. **(C)** ROIs' responses to the stimulus. These are clustered into eight groups according to the response shape. Colour bars: the colour of the cluster used in E. Red bars: the time the cluster starts responding to the widening grating. Grey margin: standard deviation over the cluster. **(D)** The standard deviation of each pixel over the 45 s of recording. **(E)** The 100 highest average ROIs from the tiled D. Each colour of an ROI is in the cluster of the same colour in C. **(F)** Acuity of each ROI. Red circles: tracking of an ROI in its cluster.

The grating is dynamic, meaning that at a given time, the wavelength is different from one point on the screen to another (**Figure 50**A). Hence, one needs the receptive field location to infer the visual acuity given the first time of grating resolving.

4.2.4 Receptive fields location

Similar to the L2 activity analysis, I calculated each recorded neuron's receptive field locations using two stimuli: a single light bar moving back and forth horizontally (e.g. **Figure 51**AB) and another vertically. As an approximation, I considered that the peak responses of each ROI (**Figure 51**C) correspond to the bar reaching the centre of its

receptive field. This approach allowed me to reconstruct a good approximation of the receptive field boundaries.
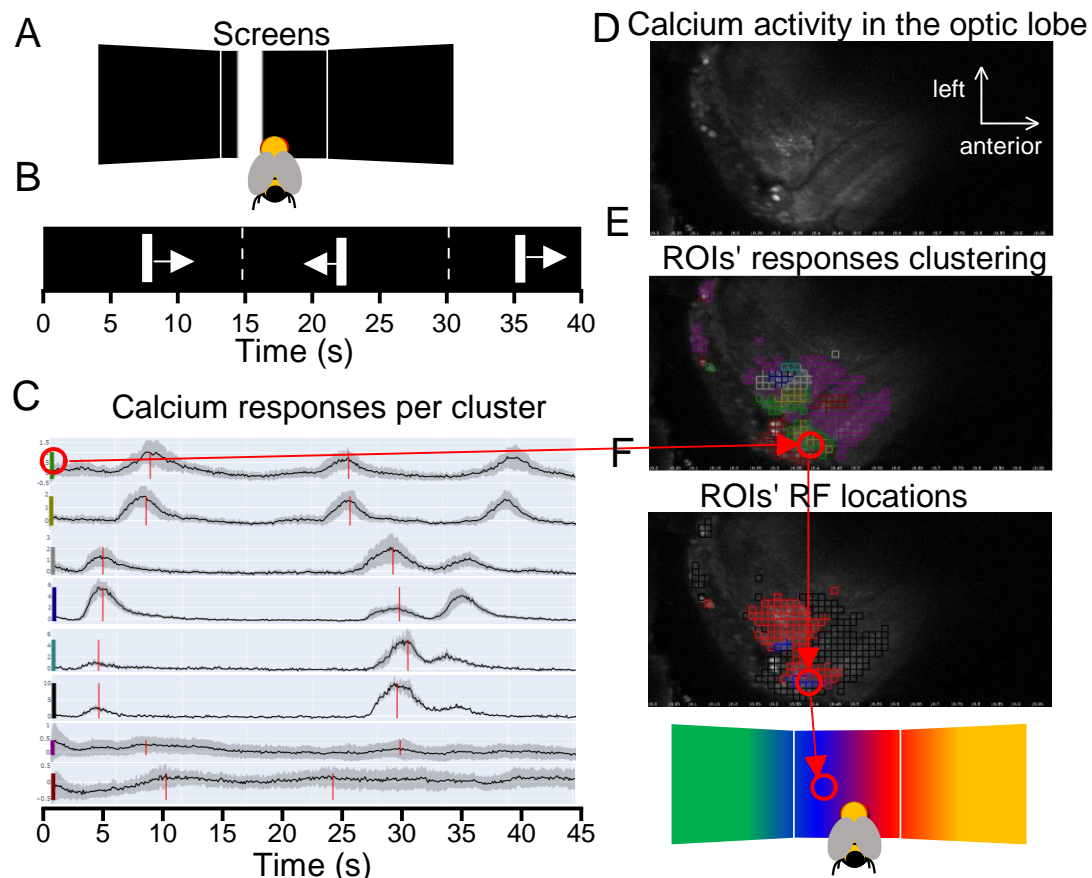
Figure 51: Receptive field (RF) measurement across the optic lobe (an example of using horizontal bars).
(A) A fly sees a single bar through three screens spanning 150° x 50°. (B) Bar motion direction over time. The bar takes 15 s to cross the 150° screens. (C) ROIs' responses to the stimulus as clustered in eight groups according to the response shape. Colour bars: the colour of the cluster used in E. Red bars: the two first times when the bar crosses the cluster's RF. Grey margin: standard deviation over the cluster. (D) The standard deviation of each pixel over the 40 s of recording. (E) The 100 highest average ROIs tiled over D. Each ROI colour matches the cluster in C. (F) Top: each ROI's RF location is represented by colour at the bottom. Note that the stimulus in this example only enables horizontal assessment of the RF locations. Red circles: tracking of an ROI in its cluster.

Hence, I have a set of visual acuity heatmaps corresponding to individual neurons in the visual pathway, and for these neurons, we know the receptive field location on the screen (**Figure 51**F). Note that this only applies to neurons responsive to the dynamic grating and the horizontal/vertical single bars. Many other neurons were only responsive to one of those two sets of stimuli and therefore had to be discarded.

4.2.5 Most sensitive directions / orientations

I calculated two values for each heatmap: the most sensitive direction and the most sensitive orientation. To do so, I used the same fitting algorithm as for the L2 analysis. Hence, for each fly, I have a map of the most sensitive orientations/directions positioned at the corresponding receptive field location.

4.2.6 Identifying the neurons

Since the Elav-flies express the calcium indicator GCaMP6f in every brain cell, it is impossible to identify the monitored neurons accurately. However, specific cues provided information about the layer of the visual pathway they are in and sometimes more of their properties.

Certain visual stimuli, such as a fixed wavelength alternating grating (**Figure 52**AB), enabled functional categorisation of the recorded neurons. I used a clustering algorithm on the time series to identify neurons' properties and locate them in the tissue. **Figure 52**C shows that some neurons respond to the luminance changes, signifying a low-order visual neuron whose dendrites branch in the lamina layer, like the L2 neurons, analysed in chapter 2. Others respond selectively to a direction of grating motion, most likely getting input from the medulla or lobula layers (these functional properties are reviewed in Borst et al., 2010). For example, the third and fourth clusters in **Figure 52** show (opposite) preferred directions, which suggests that their dendrites reside in the medulla or lobula. The second cluster responded to a moving grating, suggesting the medulla location. The first cluster responded to each bar crossing the receptive field, characteristic of the lamina neurons. These suggestive dynamics matched the neurons' physical tissue locations (**Figure 52**F and **Figure 53**).
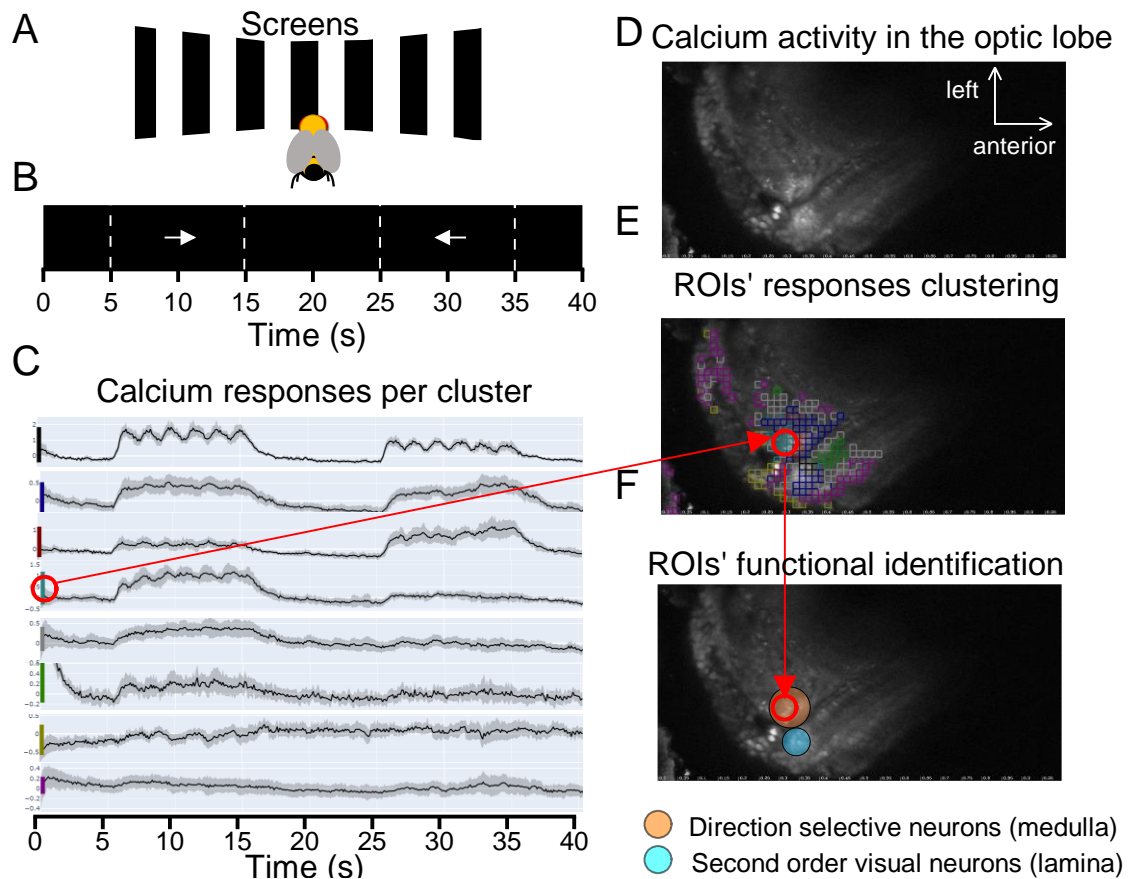
Figure 52: ROI functional identification with traditional gratings.
(A) A fly sees a single bar through three screens spanning 150° x 50°. (B) Grating motion direction over time. (C) ROIs' response to the stimulus was clustered in eight groups according to the response shape. Colour bars: the colour of the cluster used in E. Grey margin: standard deviation over the cluster. (D) The standard deviation of each pixel over the 40 s of recording. (E) The 100 highest average ROIs tiled over D. Each ROI colour matches the cluster colour in C. (F) Rough estimate of the location of in the lamina and medulla using the characteristics and locations of the clusters in C and E. Red circles: tracking of an ROI in a cluster likely to be in the medulla.

Figure 5 shows three different alignments of terminals, whose locations are consistent across the four flies. Therefore, in line with their response shapes in **Figure 52**, they likely indicate the separation between the retinal photoreceptors, lamina, medulla and lobula layers.
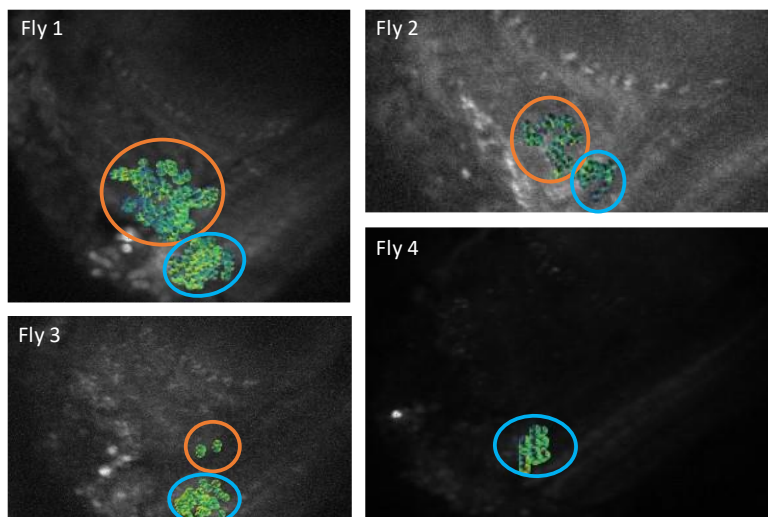
Figure 53: Spatial segregation between the lamina and medulla in 4 flies using tissue cues.
 Each of the four ROIs represents the calcium activity's standard deviation during 40 s recordings in the fly brain, expressing GCaMP pan-neuronally. Each heatmap is a visual acuity heatmap corresponding to roughly one neuron. Orange circle: cluster in the Medulla layer. Blue circle: cluster in the Lamina layer.

Note that one cannot conclude whether a neuron is in the OFF- or ON-pathways when using grating stimuli. This is because a neuron responsive to a positive edge shows the same response as a neuron responsive to a negative edge, shifted by the size of the grating bar. Hence, the only way to find this out is to use flashing stimuli.

4.3 Results
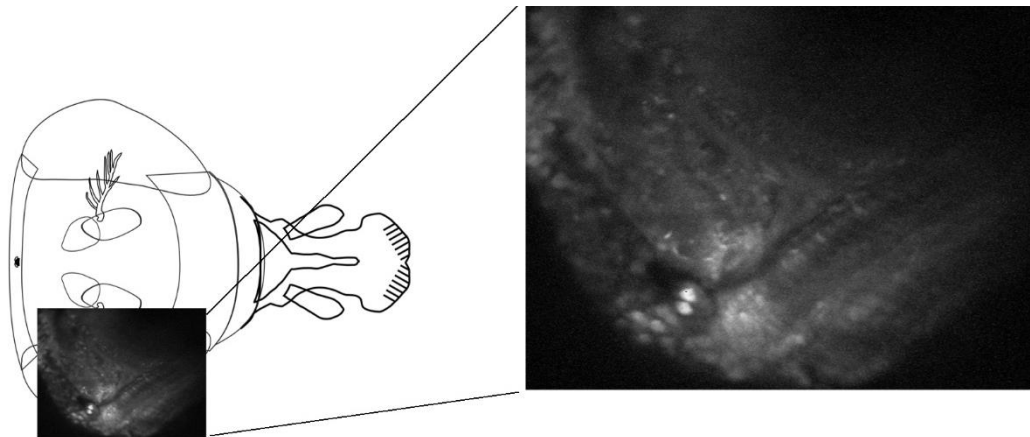
Figure 54: Pan-neuronal GCaMP6f expression visualised in the optic lobe.
A piece of tissue at the back of the fly head is removed to enable 2-photon-imaging of the calcium activity. The image shown here is the standard deviation of the calcium activity of a 40 s recording.

4.3.1 Visual hyperacuity overall in the visual pathway

I presented to a fly dynamic gratings moving in 16 different directions (evenly spanning 360°) at either 20 or 30°/s. The wavelength gradually widened with time, starting at less than 1° and ending at more than 10° (Figure 49). Such stimuli induce a response in multiple low- and high-order neurons in the visual pathway (**Figure 50**). Indeed, for a long enough wavelength, some low-order neurons respond to the luminance change induced by the moving grating and some higher-order neurons to its local or global motion. For all the recorded neuron activities, I noticed that the calcium level either rose at a particular time (a sign of a responsive neuron) or did not show any obvious correlation. For the responsive neurons, I could infer the time when the neurons start to resolve the grating.
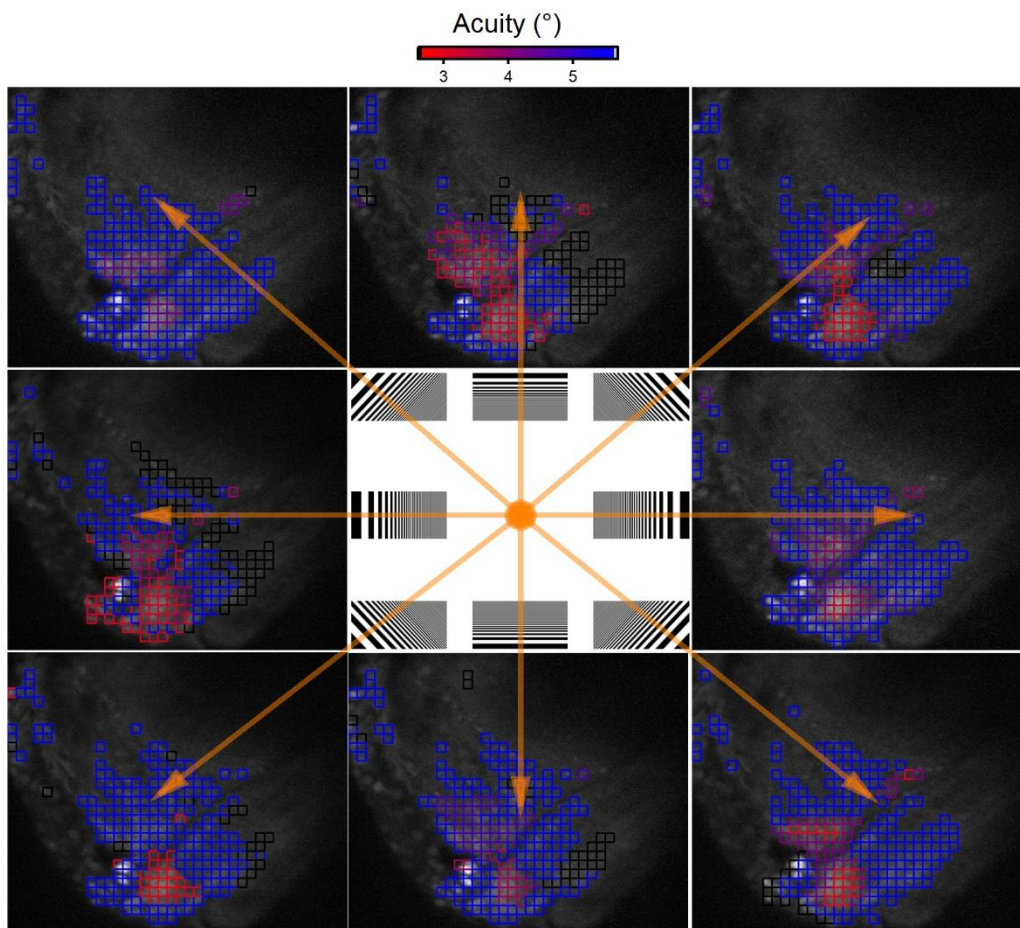
Figure 55: Visual acuity in the optic lobe for stripes moving in 8 directions at 20 °/s.
Each ROI shows the acuity of multiple pixelated regions to a specific stimulus. Arrows: direction of the stimulus.

I also presented horizontal and vertical bar-stimuli, crossing the screen at 5 °/s back and forth. These stimuli induced either two peaks in the calcium signal or no obvious correlation (**Figure 51**). I could infer the receptive field position for the responsive neurons by knowing the screen location to which they were responding.

Hence, for the neurons responding to both the dynamic gratings and single bar stimuli, I could accurately infer their visual acuity (the smallest inter-bar distance resolved) for all 32 directions/speeds used (**Figure 55**). Thus, as in Chapter 2, each neuron corresponds to a visual acuity heatmap (example of one fly in **Figure 56**A). I defined the acuity of a neuron as being the highest visual acuity for each ROI across responses to all stimuli.

I noticed that the monitored neurons typically gathered in two separate groups. Considering the visual cues and the types of responses (see **Figure 52** and **Figure 53**), I could conclude that one group was in the lamina and the other in the medulla.

I collected data from 4 flies, of which more than 100 ROIs were responsive to all stimuli.

The average acuity in the lamina and medulla across all flies was in the hyperacute range, with slightly finer acuities in the lamina. However, there was no significant difference between the two layers (**Figure 56Figure 56**B). Individually, I found that each fly showed <4.5° mean acuity (i.e., less than the average interommatidial angle) (**Figure 56**C), with some ROIs responding to <2° inter-bar-distances. Therefore, I can conclude that visual hyperacuity also readily occurs within the visual pathway's deeper layers.
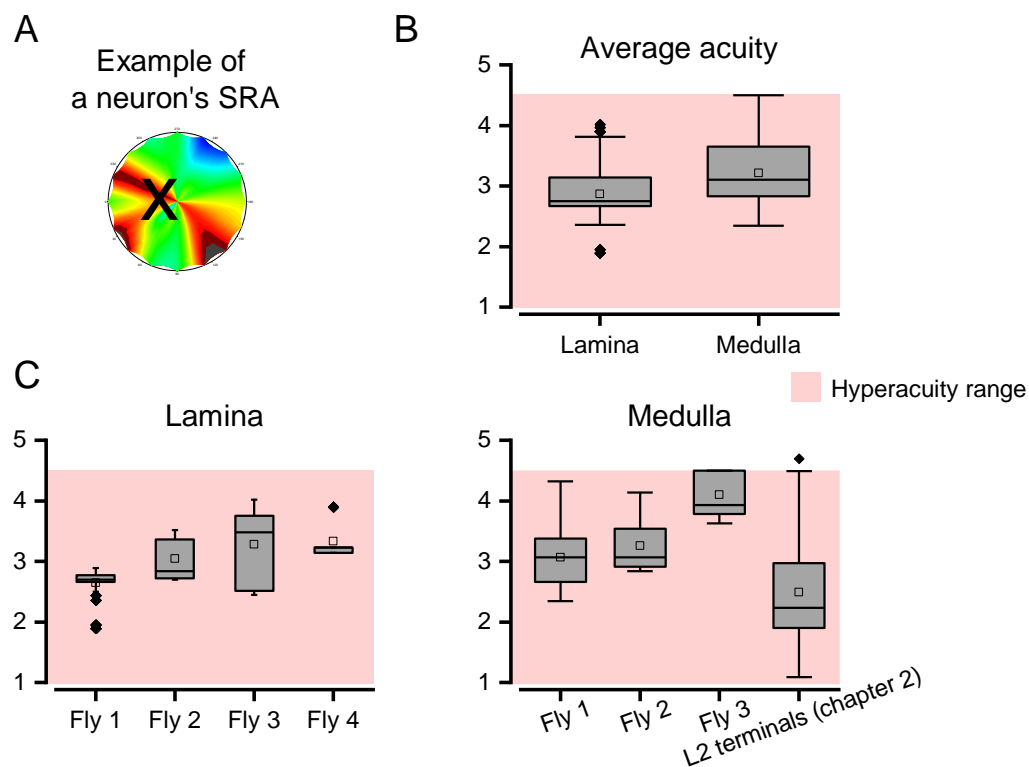
Figure 56: Visual acuity in the Lamina and Medulla.
A. Example of a visual acuity heatmap. The SRA is the minimum value of the heatmap (darkest red). **B.** Average of the SRAs across all the neurons of the four flies per layer. **C.** Average of the SRAs across all the neurons per fly, in the Lamina layer (left) and the Medulla layer (right). Note that Fly 4 only has recorded neurons in the lamina (see **Figure 53**).

4.3.2 Most sensitive directions/orientations for the whole visual pathway

Alike the earlier L2 neuron analyses, I then asked whether the recorded neurons showed a finer visual acuity for stimuli going in specific directions and if those directions aligned with the photoreceptors' saccades directions.

For each fly, I calculated visual sensitivity heatmaps for about 500 ROIs. By fitting the sensitivity with a sinusoidal function, I found that the finest visual acuity orientation (i.e.,

the smallest fitted visual acuity) was uncorrelated to the photoreceptor microsaccade directions (**Figure 57**A left) and the preferred orientation (**Figure 57**B left). When considering only the fits for which the error was < 12°, I noticed that only about a 5[th] of the ROIs showed a clear preferred orientation (**Figure 57**A right). Interestingly, these orientations primarily aligned with the corresponding photoreceptor microsaccade orientations (i.e., their back-and-forth moving directions). However, half of the ROIs showed orientations misaligned more than 30° with the microsaccade orientation. I obtained similar results when considering the directions instead of orientations, with still half of the ROIs showing preferred directions pointing more than 60° away from the microsaccade directions (**Figure 57**B right). Note that, for directions, I considered fitting errors < 24° (**Figure 57**B, right) because I measured the absolute direction differences, which span twice as much space (180°) compared to orientations' orthogonality (90°).
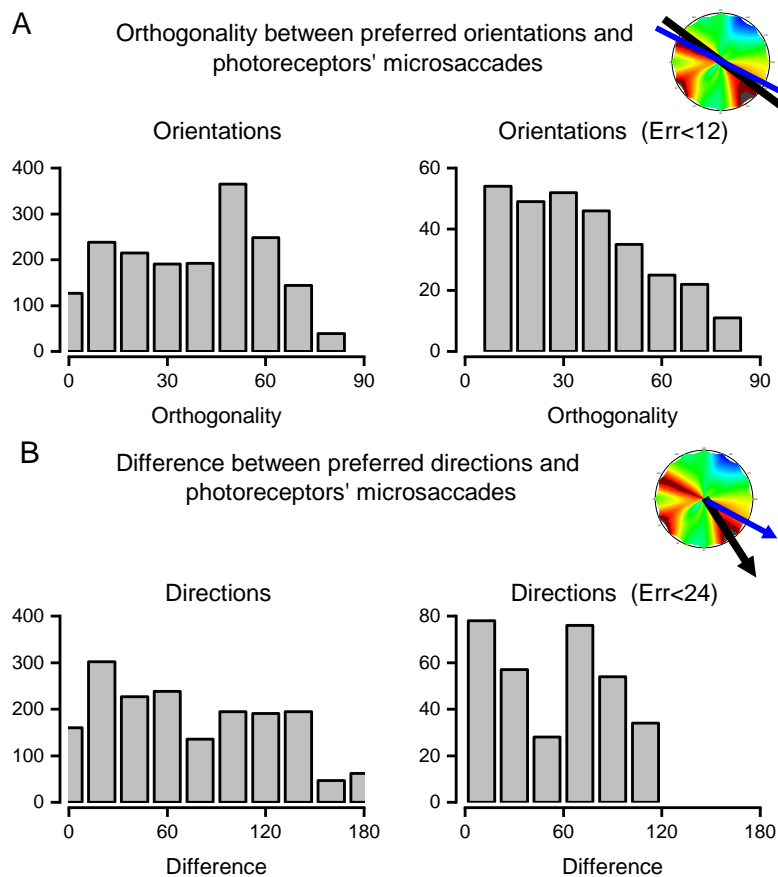
Figure 57: Preferred directions and orientations of all recorded neurons compared to photoreceptor microsaccades.
A. Distribution of the orthogonality between the preferred orientations and the corresponding microsaccade orientations. Left: all orientations. Right: orientations for which the fitting error is smaller than 12°. **B.** Distribution of the absolute angular differences between the preferred and corresponding microsaccade directions. Left: all directions. Right: directions for which the fitting error is smaller than 24°.

4.3.3 Most sensitive directions/orientations per layer

When only considering the preferred directions and orientations (for which the fitting error < 12°), I noticed that the lamina layer showed preferred orientations close to the photoreceptor microsaccades orientations. In contrast, the directions seemed uncorrelated (**Figure 58**). Conversely, in the medulla layer, the preferred orientations appeared uncorrelated with the photoreceptor microsaccades orientations, whereas the directions seemed similar.
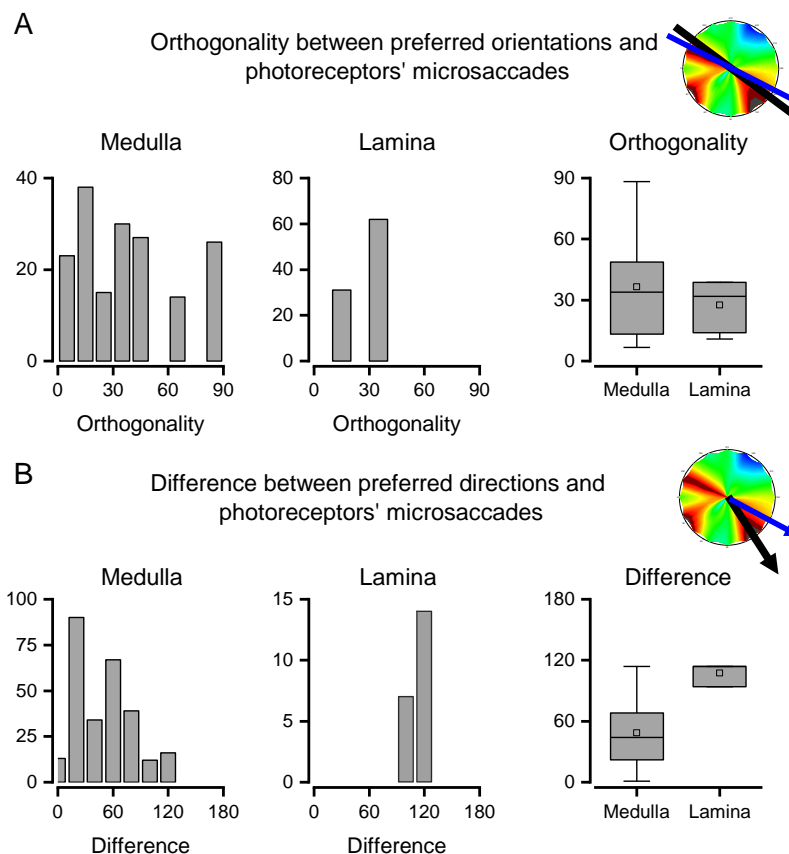
Figure 58: Comparing the preferred orientations/directions to the photoreceptor microsaccades in the lamina and medulla.
A and B (left and middle). Similar distributions as in **Figure 57**, but within the lamina and medulla layers, for which the fitting error <12° (orientations) or 24° (directions). **A and B (right).** The distributions are shown in box plots.

4.4 Discussion

Unlike the L2 analysis, I selected ROIs after tiling the standard deviation plane with 16 x 16-pixel squares and taking the top 100 squares. For the L2 analysis, I used a more accurate algorithm that continuously finds the highest standard-deviation squares in the standard-deviation plane. However, in this chapter, the field of view was too large, making the latter computationally too costly. Hence, the method I used had some drawbacks. Since the used fly line expressed the GCaMP6f-calcium indicator in every neuron, the response generated in an ROI is likely to be a mixture of different sections of neurons responding differently to the stimulus. However, this was a minor issue as the macroscopical responses of nearby neurons were easily analysable using the first raise point method (instead of a peak-to-peak resolvability), as explained in detail below. Besides the conclusions I make out of this, the macroscopical analysis remained conservative for three reasons:

- The signal-to-noise ratio can only be underestimated, as we might have only portions of neurons in each ROI.
- ON and OFF pathway neurons can cancel each other out when averaged together.
- Two non-hyperacute neurons averaged together cannot result in a hyperacute response.

Lamina neurons are low-order visual neurons that do not respond to the dynamic-wavelength grating with an increase in calcium concentration but rather a succession of peaks. As shown in **Figure 59**C, I still measured these neurons' acuity using the first-raise point, unlike the L2 analysis done in chapter 2, where I calculated the resolvability of each pair of peaks. This is because the analysed ROIs contain mixed activities of several antagonist neurons at once, with smaller raise time constants than decay time constants. In practice, this gives the succession of peaks an up-trend (see **Figure 59**C) and enables using the first-raise point algorithm.

In this chapter, we decided to infer the visual acuity of neurons using a widening dynamic-wavelength grating instead of a narrowing one as previously used in Chapter 2. In the previous chapter, using a widening or narrowing stimulus would not have made a significant difference: L2 neurons respond symmetrically to widening or narrowing gratings since they respond to light decrement, which induces a succession of peaks in the calcium response (Joesch et al., 2010b). However, most ROIs chosen in the Elav-flies showed a raising point whenever the grating is resolved, where the raise time constant is shorter

than the decay time constant (**Figure 59**A). Some exceptions exist, especially for the ROIs located in the medulla layer (see **Figure 50**C). Overall, it is still easier and more accurate to measure the point when the grating starts to be seen than when the grating is not seen anymore.
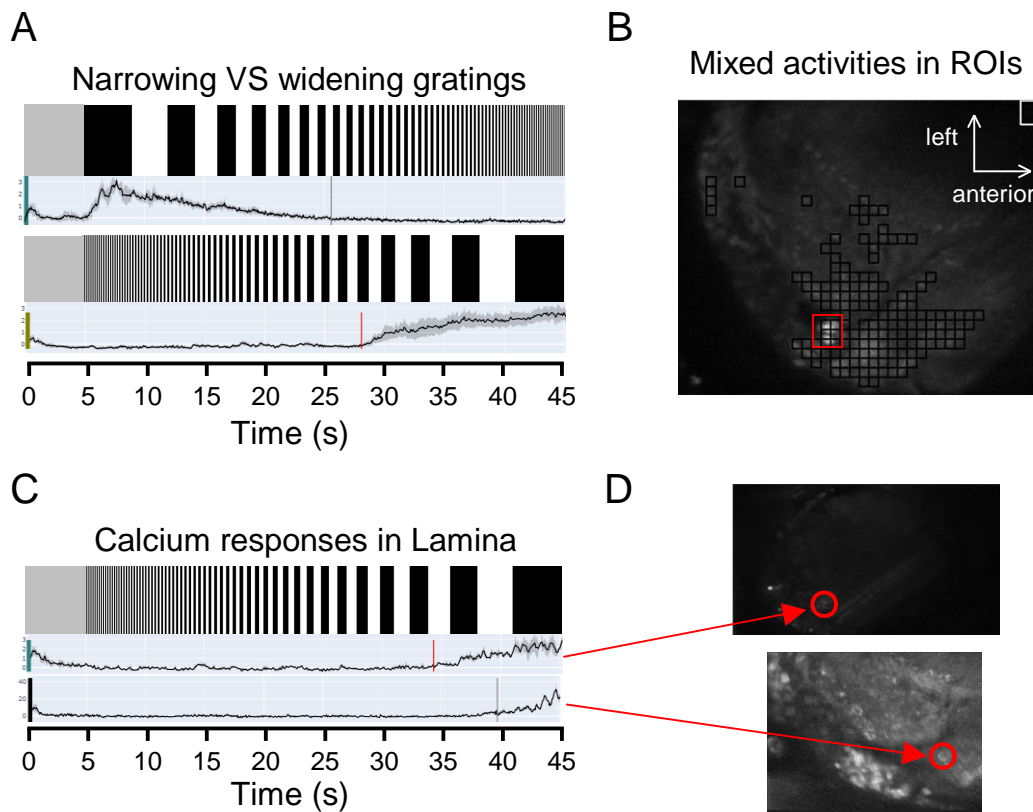
Figure 59: Widening gratings advantages and ROI limits.
 (A) Two examples of cluster calcium responses to a narrowing grating (top) and widening grating (bottom). Gray vertical bar: the point when the grating becomes unresolved. Red bar: the point when the grating becomes resolved. The grey bar calculation is less accurate than the red bar calculation because the raising time constant is higher than the decay time constant. **(B)** Example of an ROI selection where some ROIs (e.g., those in the red rectangle) share the calcium activities from parts of two or more distinct neurons. **(C)** Two examples of calcium responses of clusters showing low-order visual stimulus responses (presence of oscillations). The oscillations follow an increasing trend, which enables a valid calculation of the first-time raise point (red and grey vertical bars). **(D)** Location (red circles) in the brain of the two clusters used in C. The images confirm that the clusters are in the layer following the photoreceptors, probably the lamina. Note that the gratings shown in A and C do not represent the exact pixel value over time but rather give the type of grating used.

As seen in **Figure 57**, more than 75% of the recorded ROIs did not show a clear preferred orientation/direction (i.e., the fitting error > 13% of the space size). This phenomenon also occurred in (Kemppainen et al., 2022) but to a lesser extent, where the most sensitive L2 orientations were calculated. I separately measured the number of clear preferred

orientations and directions within the lamina and medulla. **Figure 60** shows that the medulla neurons usually show a stronger direction and orientation sensitivity than the lamina ones. This finding is consistent with the idea that the medulla is the first layer housing direction-selective neurons (such as T4 and T5 cells) (Strother et al., 2017; Tuthill et al., 2013b). Within the lamina layer, I recorded from ROIs that showed a stronger orientation than a direction sensitivity, concurring with the findings of Kemppainen et al., 2022. These hotspots likely included the photoreceptor, L2, L4, Lawf and amacrine cell terminals, with the largest fluorescence changes demarcating Ca^{2+}-dependent synaptic transmission. However, the reasons why most neurons did not show a clear directional/orientational sensitivity remain unclear. A high signal-to-noise ratio might be necessary to measure such directional or orientational sensitivity. Hence, since the signal-to-noise ratio of the ROIs was lower than in Kemppainen et al., 2022 (as discussed above), the hyperacute orientational or directional sensitivity, if there is one, might be hidden in the noise. The fact that some ROIs might share the activity of different neurons (**Figure 59**B) can also be a factor in mixed directional sensitivities.
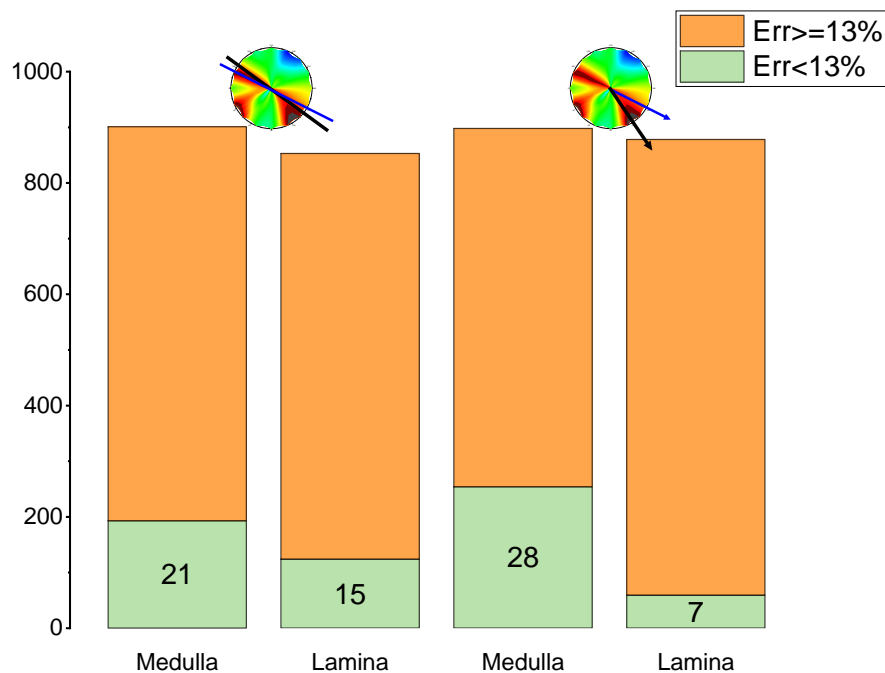
Figure 60: Proportions of good direction/orientation sensitivity fits.
 Each bar shows the number of ROIs analysed across the four flies per neuropile. In green is the amount of ROIs that show a clear (fitting error <13% of the angular space) directional (right) or orientational (left) sensitivity. The percentage of green is labelled in the green regions.

One needs to measure the RF locations of the recorded neurons to infer visual acuity or to compare the direction/orientation sensitivities with the photoreceptors of the same cartridges. The stimulus I used to calculate it is limited to the neurons in question. Indeed, a single bar does induce RF luminance changes (stimulating lamina neurons) and local motion (stimulating medulla), but not global movement (stimulating the lobula). Hence, the analysis used in this chapter cannot test the lobula neurons. To be more rigorous, we can also ask if the moving bar is relevant for local motion neurons, which are supposed to be direction-selective and not respond to the bar in both directions of motion. One can see that, even if the clustering algorithm finds some groups of ROIs that show slight direction preferences (**Figure 51**), the average activities of antagonist neurons induce responses to both directions. Once one knows the peak response times, an approximation can be made regarding the location. Indeed, for L2 neurons, a peak response corresponds to the bar leaving the neuron's RF, whereas, for L1, it corresponds to the bar entering the RF. Again, since I average multiple antagonist neurons, I considered the peak response as the bar reaching the centre of the cluster's RF. Note that inferring the visual acuity of lobula neurons is more challenging with the dynamic wavelength grating I used as they do not have a local receptive field but respond to the whole grating at once. Since this grating is always different spatially at a specific time, it is not known how to measure the acuity in such cases.

Conclusion

The work done during this PhD resulted in novel discoveries regarding the visual acuity of some specific interneurons in the *Drosophila* visual pathway, the L2 neurons. I confirmed the recent findings of Juusola's lab on fly hyperacuity, by showing that these neurons respond to hyperacute patterns, respond surprisingly better to specific orientation, and that these saccades are mostly aligned with their connected photoreceptor's microsaccade directions. I extended this work to the whole visual pathway using a pan-neuronal calcium indicator expression, which confirmed again that hyperacuity is transmitted and encoded throughout the whole fly brain. I hope this latter work is going to be beneficial for the brave scientists who will dare studying pan-neuronal expressions, which unavoidably are a great challenge.

To make a more concrete use of my findings, I made an Artificial Neural Network, which structure is based on the bio-physical properties of the photoreceptor and the lamina layer in the *Drosophila*. This model, trained on the data that I recorded myself during my PhD, can predict neural activity at the initial layers of the fly visual pathway and encode ultra-fine details with low energy requirements. Furthermore, it proved that photoreceptor microsaccades are involved in direction specificity and hyperacuity, and that neural superposition is likely to be involved in hyperacuity by inducing sharp and fast responses to moving edges. Eventually, I hope this model can inspire for computer vision, robotics and neuroscience.

A Drosophila *- drawn by Léïla Roebben*

Reference

Ache, J. M., Namiki, S., Lee, A., Branson, K., & Card, G. M. (2019). State-dependent decoupling of sensory and motor circuits underlies behavioral flexibility in Drosophila. *Nature Neuroscience*. https://doi.org/10.1038/s41593-019-0413-4

Amari, S. ichi. (1993). Backpropagation and stochastic gradient descent method. *Neurocomputing*, *5*(4–5), 185–196. https://doi.org/10.1016/0925-2312(93)90006-O

Arakaki, T., Barello, G., & Ahmadian, Y. (2019). Inferring neural circuit structure from datasets of heterogeneous tuning curves. *PLoS Computational Biology*, *15*(4), e1006816. https://doi.org/10.1371/journal.pcbi.1006816

Barnhart, E. L., Wang, I. E., Wei, H., Desplan, C., & Clandinin, T. R. (2018). Sequential Nonlinear Filtering of Local Motion Cues by Global Motion Circuits. *Neuron*. https://doi.org/10.1016/j.neuron.2018.08.022

Behnia, R., Clark, D. A., Carter, A. G., Clandinin, T. R., & Desplan, C. (2014). Processing properties of on and off pathways for Drosophila motion detection. *Nature*, *512*(7515), 427–430. https://doi.org/10.1038/nature13427

Borst, A., & Egelhaaf, M. (1990). Direction selectivity of blowfly motion-sensitive neurons is computed in a two-stage process. *Proceedings of the National Academy of Sciences of the United States of America*, *87*(23), 9363–9367. https://doi.org/10.1073/pnas.87.23.9363

Borst, A., & Haag, J. (2002). Neural networks in the cockpit of the fly. In *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology* (Vol. 188, Issue 6, pp. 419–437). J Comp Physiol A Neuroethol Sens Neural Behav Physiol. https://doi.org/10.1007/s00359-002-0316-8

Borst, Alexander, Haag, J., & Mauss, A. S. (2020). How fly neurons compute the direction of visual motion. In *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology* (Vol. 206, Issue 2, pp. 109–124). Springer. https://doi.org/10.1007/s00359-019-01375-9

Borst, Alexander, Haag, J., & Reiff, D. F. (2010). *Fly Motion Vision*. https://doi.org/10.1146/annurev-neuro-060909-153155

Borst, Alexander, & Helmstaedter, M. (2015). Common circuit design in fly and mammalian motion vision. In *Nature Neuroscience* (Vol. 18, Issue 8, pp. 1067–1076). Nature Publishing Group. https://doi.org/10.1038/nn.4050

Braitenberg, V. (1967). Patterns of projection in the visual system of the fly. I. Retina-lamina projections. *Experimental Brain Research*, *3*(3), 271–298. https://doi.org/10.1007/BF00235589

Brand, A. H., & Perrimon, N. (1993). Targeted gene expression as a means of altering cell fates and generating dominant phenotypes. *Development*.

Burtt, E. T., & Catton, W. T. (1954). Visual perception of movement in the locust. *The Journal of Physiology*, *125*(3), 566–580. https://doi.org/10.1113/jphysiol.1954.sp005181

Campos-Ortega, J. A., & Strausfeld, N. J. (1972). Columns and Layers in the Second Synaptic Region of the Fly's Visual System: The Case for Two Superimposed Neuronal Architectures. In *Information Processing in the Visual Systems of Anthropods* (pp. 31–36). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-65477-0_5

Chiappe, M. E., & Jayaraman, V. (2012). Performing electrophysiology and two-photon calcium imaging in the adult drosophila central brain during walking behavior. *Neuromethods*. https://doi.org/10.1007/978-1-62703-014-4-5

Chiappe, M. E., Seelig, J. D., Reiser, M. B., & Jayaraman, V. (2010). Walking modulates speed sensitivity in drosophila motion vision. *Current Biology*, *20*(16), 1470–1475. https://doi.org/10.1016/j.cub.2010.06.072

Clandinin, T. R., & Zipursky, S. L. (2002). Making connections in the fly visual system. In *Neuron*. https://doi.org/10.1016/S0896-6273(02)00876-0

Clark, D. A., Bursztyn, L., Horowitz, M. A., Schnitzer, M. J., & Clandinin, T. R. (2011). Defining the Computational Structure of the Motion Detector in Drosophila. *Neuron*. https://doi.org/10.1016/j.neuron.2011.05.023

Courgeon, M., & Desplan, C. (2019). Coordination between stochastic and deterministic specification in the Drosophila visual system. *Science*, *366*(6463). https://doi.org/10.1126/science.aay6727

Creswell, A., White, T., Dumoulin, V., Arulkumaran, K., Sengupta, B., & Bharath, A. A. (2018). Generative Adversarial Networks: An Overview. In *IEEE Signal Processing Magazine* (Vol. 35, Issue 1, pp. 53–65). Institute of Electrical and Electronics Engineers Inc. https://doi.org/10.1109/MSP.2017.2765202

Dau, A., Friederich, U., Dongre, S., Li, X., Bollepalli, M. K., Hardie, R. C., & Juusola, M. (2016). Evidence for dynamic network regulation of Drosophila photoreceptor function from mutants lacking the neurotransmitter histamine. *Frontiers in Neural Circuits*, *10*(MAR). https://doi.org/10.3389/fncir.2016.00019

Deng, B., Li, Q., Liu, X., Cao, Y., Li, B., Qian, Y., Xu, R., Mao, R., Zhou, E., Zhang, W., Huang, J., & Rao, Y. (2019). Chemoconnectomics: Mapping Chemical Transmission in Drosophila. *Neuron*, *101*(5), 876-893.e4. https://doi.org/10.1016/j.neuron.2019.01.045

Denk, W., Strickler, J. H., & Webb, W. W. (1990). Two-photon laser scanning fluorescence microscopy. *Science*. https://doi.org/10.1126/science.2321027

Ding, B., Qian, H., & Zhou, J. (2018). Activation functions and their characteristics in deep neural networks. *Proceedings of the 30th Chinese Control and Decision Conference, CCDC 2018*, 1836–1841. https://doi.org/10.1109/CCDC.2018.8407425

Dombeck, D. A., Harvey, C. D., Tian, L., Looger, L. L., & Tank, D. W. (2010). Functional imaging of hippocampal place cells at cellular resolution during virtual navigation. *Nature Neuroscience*.

https://doi.org/10.1038/nn.2648

Dubs, A. (1982). The spatial integration of signals in the retina and lamina of the fly compound eye under different conditions of luminance. *Journal of Comparative Physiology □ A*, *146*(3), 321–343. https://doi.org/10.1007/BF00612703

E. T. Burtt and W. T. Catton. (1962). *A Diffraction Theory of Insect Vision I. An Experimental Investigation of Visual Acuity and Image Formation in the Compound Eyes of Three Species of Insects*. Proceedings of the Royal Society B: Biological Sciences. https://www.jstor.org/stable/90271

Elyada, Y. M. (2009). *Intracellular processing of motion information in a network of blowfly visual interneurons*.

Fischbach, K. F., & Dittrich, A. P. M. (1989). The optic lobe of Drosophila melanogaster. I. A Golgi analysis of wild-type structure. *Cell and Tissue Research*, *258*(3), 441–475. https://doi.org/10.1007/BF00218858

Fischbach, Karl Friedrich, & Hiesinger, P. R. (2008). Optic lobe development. In *Advances in Experimental Medicine and Biology* (Vol. 628, pp. 115–136). https://doi.org/10.1007/978-0-387-78261-4_8

Fisher, Y. E., Silies, M., & Clandinin, T. R. (2015). Orientation Selectivity Sharpens Motion Detection in Drosophila. *Neuron*, *88*(2), 390–402. https://doi.org/10.1016/J.NEURON.2015.09.033

Freifeld, L., Clark, D. A., Schnitzer, M. J., Horowitz, M. A., & Clandinin, T. R. (2013a). GABAergic Lateral Interactions Tune the Early Stages of Visual Processing in Drosophila. *Neuron*. https://doi.org/10.1016/j.neuron.2013.04.024

Freifeld, L., Clark, D. A., Schnitzer, M. J., Horowitz, M. A., & Clandinin, T. R. (2013b). GABAergic Lateral Interactions Tune the Early Stages of Visual Processing in Drosophila. *Neuron*, *78*(6), 1075–1089. https://doi.org/10.1016/j.neuron.2013.04.024

Fu, Q., & Yue, S. (2020). Modelling Drosophila Motion Vision Pathways for Decoding the Direction of Translating Objects Against Cluttered Moving Backgrounds. In *arXiv*.

Fujiwara, T., Cruz, T. L., Bohnslav, J. P., & Chiappe, M. E. (2017). A faithful internal representation of walking movements in the Drosophila visual system. *Nature Neuroscience*. https://doi.org/10.1038/nn.4435

Gonzalez-Bellido, P. T., Wardill, T. J., & Juusola, M. (2011). Compound eyes and retinal information processing in miniature dipteran species match their specific ecological demands. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(10), 4224–4229. https://doi.org/10.1073/PNAS.1014438108

Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M., & Cohen, J. D. (2001). An fMRI investigation of emotional engagement in moral judgment. *Science*, *293*(5537), 2105–2108. https://doi.org/10.1126/science.1062872

Grienberger, C., & Konnerth, A. (2012). Imaging Calcium in Neurons. In *Neuron*.

https://doi.org/10.1016/j.neuron.2012.02.011

Gschweng, K. M. M., Sampson, M. M., Hardcastle, B. J., Sizemore, T. R., Dacks, A. M., Frye, M. A., & Krantz, D. E. (2019). Serotonergic Modulation of a Visual Microcircuit in Drosophila melanogaster. *BioRxiv*. https://doi.org/10.1101/619759

Hardcastle, B. J., & Krapp, H. G. (2016). Evolution of Biological Image Stabilization. In *Current Biology* (Vol. 26, Issue 20, pp. R1010–R1021). Cell Press. https://doi.org/10.1016/j.cub.2016.08.059

Hardie, R. C. (1987). Is histamine a neurotransmitter in insect photoreceptors? *Journal of Comparative Physiology A*, *161*(2), 201–213. https://doi.org/10.1007/BF00615241

Hardie, R. C., & Postma, M. (2008). Phototransduction in Microvillar Photoreceptors ofDrosophila and Other Invertebrates. In *The Senses: A Comprehensive Reference* (Vol. 1, pp. 77–130). Elsevier Inc. https://doi.org/10.1016/B978-012370880-9.00402-3

Hardie, R. C., & Raghu, P. (2001). Visual transduction in Drosophila. In *Nature* (Vol. 413, Issue 6852, pp. 186–193). https://doi.org/10.1038/35093002

Hardie, Roger C., & Franze, K. (2012). Photomechanical responses in Drosophila photoreceptors. *Science*. https://doi.org/10.1126/science.1222376

Hardie, Roger C., & Juusola, M. (2015). Phototransduction in Drosophila. In *Current Opinion in Neurobiology* (Vol. 34, pp. 37–45). Elsevier Ltd. https://doi.org/10.1016/j.conb.2015.01.008

Heisenberg, M., & Buchner, E. (1977). The rôle of retinula cell types in visual behavior of Drosophila melanogaster. *Journal of Comparative Physiology ▪ A*, *117*(2), 127–162. https://doi.org/10.1007/BF00612784

Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*. https://doi.org/10.1113/jphysiol.1962.sp006837

Jain, L., & Medsker, L. (2000). *Recurrent Neural Networks: Design and Appication*. https://www.taylorfrancis.com/books/9781420049176

Järvilehto, M., & Zettler, F. (1973). Electrophysiological-histological studies on some functional properties of visual cells and second order neurons of an insect retina. *Zeitschrift Für Zellforschung Und Mikroskopische Anatomie*, *136*(2), 291–306. https://doi.org/10.1007/BF00307446

Joesch, M., Plett, J., Borst, A., & Reiff, D. F. (2008). Response Properties of Motion-Sensitive Visual Interneurons in the Lobula Plate of Drosophila melanogaster. *Current Biology*, *18*(5), 368–374. https://doi.org/10.1016/j.cub.2008.02.022

Joesch, M., Schnell, B., Raghu, S. V., Reiff, D. F., & Borst, A. (2010a). ON and off pathways in Drosophila motion vision. *Nature*. https://doi.org/10.1038/nature09545

Joesch, M., Schnell, B., Raghu, S. V., Reiff, D. F., & Borst, A. (2010b). ON and OFF pathways in

Drosophila motion vision. *Nature 2010 468:7321*, *468*(7321), 300–304. https://doi.org/10.1038/nature09545

Juusola, M. (1993). Linear and non-linear contrast coding in light-adapted blowfly photoreceptors. *Journal of Comparative Physiology A*, *172*(4), 511–521. https://doi.org/10.1007/BF00213533

Juusola, M., Kouvalainen, E., Järvilehto, M., & Weckström, M. (1994). Contrast gain, signal-to-noise ratio, and linearity in light-adapted blowfly photoreceptors. *Journal of General Physiology*, *104*(3), 593–621. https://doi.org/10.1085/jgp.104.3.593

Juusola, Mikko, Dau, A., Song, Z., Solanki, N., Rien, D., Jaciuch, D., Dongre, S. A., Blanchard, F., De Polavieja, G. G., Hardie, R. C., & Takalo, J. (2017). Microsaccadic sampling of moving image information provides Drosophila hyperacute vision. *ELife*. https://doi.org/10.7554/eLife.26117

Juusola, Mikko, & French, A. S. (1997). Visual acuity for moving objects in first- and second-order neurons of the fly compound eye. *Journal of Neurophysiology*. https://doi.org/10.1152/jn.1997.77.3.1487

Juusola, Mikko, & Hardie, R. C. (2001). Light adaptation in Drosophila photoreceptors: II. Rising temperature increases the bandwidth of reliable signaling. *Journal of General Physiology*, *117*(1), 27–41. https://doi.org/10.1085/jgp.117.1.27

Kemppainen, J., Scales, B., Haghighi, K. R., Takalo, J., Mansour, N., McManus, J., Leko, G., Saari, P., Hurcomb, J., Antohi, A., Suuronen, J. P., Blanchard, F., Hardie, R. C., Song, Z., Hampton, M., Eckermann, M., Westermeier, F., Frohn, J., Hoekstra, H., … Juusola, M. (2022). Binocular mirror–symmetric microsaccadic sampling enables Drosophila hyperacute 3D vision. *Proceedings of the National Academy of Sciences of the United States of America*, *119*(12). https://doi.org/10.1073/pnas.2109717119

Kirschfeld, K., & Franceschini, N. (1968). Optische Eigenschaften der Ommatidien im Komplexauge von Musca. *Kybernetik*, *5*(2), 47–52. https://doi.org/10.1007/BF00272694

Kirschfeld, Kuno. (1967). Die projektion der optischen umwelt auf das raster der rhabdomere im komplexauge von Musca. *Experimental Brain Research*, *3*(3), 248–270. https://doi.org/10.1007/BF00235588

Klapoetke, N. C., Nern, A., Rogers, E. M., Rubin, G. M., Reiser, M. B., & Card, G. M. (2022). A functionally ordered visual feature map in the Drosophila brain. *Neuron*, *110*(10), 1700-1711.e6. https://doi.org/10.1016/j.neuron.2022.02.013

Kolodziejczyk, A., Sun, X., Meinertzhagen, I. A., & Nässel, D. R. (2008). Glutamate, GABA and acetylcholine signaling components in the lamina of the Drosophila visual system. *PLoS ONE*, *3*(5). https://doi.org/10.1371/journal.pone.0002110

Krapp, H. G., & Hengstenberg, R. (1996). Estimation of self-motion by optic flow processing in single visual interneurons. *Nature*, *384*(6608), 463–466. https://doi.org/10.1038/384463a0

Kumar, J. P., & Ready, D. F. (1995). Rhodopsin plays an essential structural role in Drosophila photoreceptor development. *Development (Cambridge, England)*, *121*(12), 4359–4370.

https://doi.org/10.1242/DEV.121.12.4359

Land, M. F. (1997). Visual acuity in insects. *Annual Review of Entomology. Vol. 42*. https://doi.org/10.1146/annurev.ento.42.1.147

Laughlin, S. B., & Hardie, R. C. (1978). Common strategies for light adaptation in the peripheral visual systems of fly and dragonfly. *Journal of Comparative Physiology □ A*, *128*(4), 319–340. https://doi.org/10.1007/BF00657606

Lawson, K. K. K., & Srinivasan, M. V. (2020). Contrast sensitivity and visual acuity of Queensland fruit flies (Bactrocera tryoni). *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*. https://doi.org/10.1007/s00359-020-01404-y

Lecun, Y. A., Bottou, L., Orr, G. B., & Müller, K. (2012). *Efficient BackProp*. *0*(1998), 9–48.

Lecun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. In *Nature*. https://doi.org/10.1038/nature14539

Logothetis, N. K. (2008). What we can do and what we cannot do with fMRI. In *Nature* (Vol. 453, Issue 7197, pp. 869–878). Nature Publishing Group. https://doi.org/10.1038/nature06976

Lott, G. K., Rosen, M. J., & Hoy, R. R. (2007). An inexpensive sub-millisecond system for walking measurements of small animals based on optical computer mouse technology. *Journal of Neuroscience Methods*. https://doi.org/10.1016/j.jneumeth.2006.10.007

MaBouDi, H., Marshall, J. A. R., & Barron, A. B. (2020). Honeybees solve a multi-comparison ranking task by probability matching. *Proceedings of the Royal Society B: Biological Sciences*, *287*(1934), 20201525. https://doi.org/10.1098/rspb.2020.1525

Maimon, G., Straw, A. D., & Dickinson, M. H. (2010). Active flight increases the gain of visual motion processing in Drosophila. *Nature Neuroscience*. https://doi.org/10.1038/nn.2492

Mano, O., Creamer, M. S., Badwan, B. A., & Clark, D. A. (2021). Predicting individual neuron responses with anatomically constrained task optimization. *Current Biology*, *31*(18), 4062-4075.e4. https://doi.org/10.1016/j.cub.2021.06.090

Meinertzhagen, I. A., & O'Neil, S. D. (1991). Synaptic organization of columnar elements in the lamina of the wild type in Drosophila melanogaster. *The Journal of Comparative Neurology*, *305*(2), 232–263. https://doi.org/10.1002/CNE.903050206

Miyawaki, A., Llopis, J., Heim, R., Michael McCaffery, J., Adams, J. A., Ikura, M., & Tsien, R. Y. (1997). Fluorescent indicators for Ca2+ based on green fluorescent proteins and calmodulin. *Nature*. https://doi.org/10.1038/42264

Morante, J., & Desplan, C. (2005). Photoreceptor axons play hide and seek. In *Nature Neuroscience* (Vol. 8, Issue 4, pp. 401–402). Nature Publishing Group. https://doi.org/10.1038/nn0405-401

Nikolaev, A., Zheng, L., Wardill, T. J., O'Kane, C. J., De Polavieja, G. G., & Juusola, M. (2009). Network Adaptation Improves Temporal Representation of Naturalistic Stimuli in Drosophila Eye: II Mechanisms. *PLoS ONE*, *4*(1), 4306. https://doi.org/10.1371/JOURNAL.PONE.0004306

Nityananda, V., & Read, J. C. A. (2017). Stereopsis in animals: evolution, function and mechanisms. *The Journal of Experimental Biology*, *220*(14), 2502. https://doi.org/10.1242/JEB.143883

Pape, L., Ruessink, B. G., Wiering, M. A., & Turner, I. L. (2007). Recurrent neural network modeling of nearshore sandbar behavior. *Neural Networks*, *20*(4), 509–518. https://doi.org/10.1016/j.neunet.2007.04.007

Paukert, M., & Bergles, D. E. (2012). Reduction of motion artifacts during in vivo two-photon imaging of brain through heartbeat triggered scanning. *Journal of Physiology*. https://doi.org/10.1113/jphysiol.2012.228114

Paulk, A. C., Kirszenblat, L., Zhou, Y., & van Swinderen, B. (2015). Closed-loop behavioral control increases coherence in the fly brain. *Journal of Neuroscience*. https://doi.org/10.1523/JNEUROSCI.0691-15.2015

Pichaud, F., Briscoe, A., & Desplan, C. (1999). Evolution of color visione. In *Current Opinion in Neurobiology* (Vol. 9, Issue 5, pp. 622–627). Current Biology Ltd. https://doi.org/10.1016/S0959-4388(99)00014-8

Raghu, S. V., Claussen, J., & Borst, A. (2013). Neurons with GABAergic phenotype in the visual system of Drosophila. *Journal of Comparative Neurology*, *521*(1), 252–265. https://doi.org/10.1002/cne.23208

Ramos-Traslosheros, G., Henning, M., & Silies, M. (2018). Motion detection: Cells, circuits and algorithms. *Neuroforum*. https://doi.org/10.1515/nf-2017-A028

Redcay, E., Kennedy, D. P., & Courchesne, E. (2007). fMRI during natural sleep as a method to study brain function during early childhood. *NeuroImage*, *38*(4), 696–707. https://doi.org/10.1016/j.neuroimage.2007.08.005

Reichardt, W. (1987). Evaluation of optical motion information by movement detectors. *Journal of Comparative Physiology A*, *161*(4), 533–547. https://doi.org/10.1007/BF00603660

Reiff, D. F., Plett, J., Mank, M., Griesbeck, O., & Borst, A. (2010). Visualizing retinotopic half-wave rectified input to the motion detection circuitry of Drosophila. *Nature Neuroscience*. https://doi.org/10.1038/nn.2595

Rein, K., Zöckler, M., Mader, M. T., Grübel, C., & Heisenberg, M. (2002). The Drosophila standard brain. *Current Biology*, *12*(3), 227–231. https://doi.org/10.1016/S0960-9822(02)00656-5

Rister, J., Pauls, D., Schnell, B., Ting, C. Y., Lee, C. H., Sinakevitch, I., Morante, J., Strausfeld, N. J., Ito, K., & Heisenberg, M. (2007). Dissection of the Peripheral Motion Channel in the Visual System of Drosophila melanogaster. *Neuron*. https://doi.org/10.1016/j.neuron.2007.09.014

Rivera-Alba, M., Vitaladevuni, S. N., Mischenko, Y., Lu, Z., Takemura, S. Y., Scheffer, L., Meinertzhagen, I. A., Chklovskii, D. B., & De Polavieja, G. G. (2011). Wiring economy and volume exclusion determine neuronal placement in the Drosophila brain. *Current Biology*. https://doi.org/10.1016/j.cub.2011.10.022

Saint Marie, R. L., & Carlson, S. D. (1985). Interneuronal and glial-neuronal gap junctions in the

lamina ganglionaris of the compound eye of the housefly, Musca domestics. *Cell and Tissue Research*, *241*(1), 43–52. https://doi.org/10.1007/BF00214624

Scheffer, L. K. (2020). Graph Properties of the Adult Drosophila Central Brain. *BioRxiv*, 2020.05.18.102061. https://doi.org/10.1101/2020.05.18.102061

Schilling, T., Ali, A. H., Leonhardt, A., Borst, A., & Pujol-Martí, J. (2019). Transcriptional control of morphological properties of direction-selective T4/T5 neurons in Drosophila. *Development (Cambridge)*. https://doi.org/10.1242/dev.169763

Schneider, J., Murali, N., Taylor, G. W., & Levine, J. D. (2018). Can drosophila melanogaster tell who's who? *PLoS ONE*. https://doi.org/10.1371/journal.pone.0205043

Schnell, B., Joesch, M., Forstner, F., Raghu, S. V., Otsuna, H., Ito, K., Borst, A., & Reiff, D. F. (2010). Processing of horizontal optic flow in three visual interneurons of the Drosophila brain. *Journal of Neurophysiology*, *103*(3), 1646–1657. https://doi.org/10.1152/jn.00950.2009

Seelig, J. D., Chiappe, M. E., Lott, G. K., Dutta, A., Osborne, J. E., Reiser, M. B., & Jayaraman, V. (2010). Two-photon calcium imaging from head-fixed Drosophila during optomotor walking behavior. *Nature Methods*. https://doi.org/10.1038/nmeth.1468

Seelig, J. D., & Jayaraman, V. (2013). Feature detection and orientation tuning in the Drosophila central complex. *Nature*. https://doi.org/10.1038/nature12601

Shi, J., & Tomasi, C. (1994). Good features to track. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. https://doi.org/10.1109/cvpr.1994.323794

Shinomiya, K., Huang, G., Lu, Z., Parag, T., Xu, C. S., Aniceto, R., Ansari, N., Cheatham, N., Lauchie, S., Neace, E., Ogundeyi, O., Ordish, C., Peel, D., Shinomiya, A., Smith, C., Takemura, S., Talebi, I., Rivlin, P. K., Nern, A., … Meinertzhagen, I. A. (2019). Comparisons between the ON- and OFF-edge motion pathways in the Drosophila brain. *ELife*. https://doi.org/10.7554/eLife.40025

Silies, M., Gohl, D. M., Fisher, Y. E., Freifeld, L., Clark, D. A., & Clandinin, T. R. (2013). Modular Use of Peripheral Input Channels Tunes Motion-Detecting Circuitry. *Neuron*, *79*(1), 111–127. https://doi.org/10.1016/J.NEURON.2013.04.029

Snyder, A. W., Stavenga, D. G., & Laughlin, S. B. (1977). Spatial information capacity of compound eyes. *Journal of Comparative Physiology ■ A*, *116*(2), 183–207. https://doi.org/10.1007/BF00605402

Song, Z., & Juusola, M. (2014). Refractory sampling links efficiency and costs of sensory encoding to stimulus statistics. *Journal of Neuroscience*, *34*(21), 7216–7237. https://doi.org/10.1523/JNEUROSCI.4463-13.2014

Song, Z., Postma, M., Billings, S. A., Coca, D., Hardie, R. C., & Juusola, M. (2012). Stochastic, adaptive sampling of information by microvilli in fly photoreceptors. *Current Biology*, *22*(15), 1371–1380. https://doi.org/10.1016/j.cub.2012.05.047

Srinivasan, M. V., & Bernard, G. D. (1975). The effect of motion on visual acuity of the compound eye: A theoretical analysis. *Vision Research*. https://doi.org/10.1016/0042-6989(75)90029-2

Stavenga, D. G. (2003). Angular and spectral sensitivity of fly photoreceptors. II. Dependence on facet lens F-number and rhabdomere type in Drosophila. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, *189*(3), 189–202. https://doi.org/10.1007/s00359-003-0390-6

Strother, J. A., Wu, S. T., Rogers, E. M., Eliason, J. L. M., Wong, A. M., Nern, A., & Reiser, M. B. (2018). Behavioral state modulates the on visual motion pathway of drosophila. *Proceedings of the National Academy of Sciences of the United States of America*. https://doi.org/10.1073/pnas.1703090115

Strother, J. A., Wu, S. T., Wong, A. M., Nern, A., Rogers, E. M., Le, J. Q., Rubin, G. M., & Reiser, M. B. (2017). The Emergence of Directional Selectivity in the Visual Motion Pathway of Drosophila. *Neuron*. https://doi.org/10.1016/j.neuron.2017.03.010

Sun, Y., Nern, A., Franconville, R., Dana, H., Schreiter, E. R., Looger, L. L., Svoboda, K., Kim, D. S., Hermundstad, A. M., & Jayaraman, V. (2017a). Neural signatures of dynamic stimulus selection in Drosophila. *Nature Neuroscience*. https://doi.org/10.1038/nn.4581

Sun, Y., Nern, A., Franconville, R., Dana, H., Schreiter, E. R., Looger, L. L., Svoboda, K., Kim, D. S., Hermundstad, A. M., & Jayaraman, V. (2017b). Neural signatures of dynamic stimulus selection in Drosophila. *Nature Neuroscience*. https://doi.org/10.1038/nn.4581

Takalo, J., Piironen, A., Honkanen, A., Lempea, M., Aikio, M., Tuukkanen, T., & Vähäsöyrinki, M. (2012). A fast and flexible panoramic virtual reality system for behavioural and electrophysiological experiments. *Scientific Reports*. https://doi.org/10.1038/srep00324

Takemura, S.-Y., Lu, Z., & Meinertzhagen, I. A. (2008). Synaptic circuits of the *Drosophila* optic lobe: The input terminals to the medulla. *The Journal of Comparative Neurology*, *509*(5), 493–513. https://doi.org/10.1002/cne.21757

Takemura, S. Y. (2015). Connectome of the fly visual circuitry. In *Microscopy*. https://doi.org/10.1093/jmicro/dfu102

Takemura, S. Y., Bharioke, A., Lu, Z., Nern, A., Vitaladevuni, S., Rivlin, P. K., Katz, W. T., Olbris, D. J., Plaza, S. M., Winston, P., Zhao, T., Horne, J. A., Fetter, R. D., Takemura, S., Blazek, K., Chang, L. A., Ogundeyi, O., Saunders, M. A., Shapiro, V., … Chklovskii, D. B. (2013). A visual motion detection circuit suggested by Drosophila connectomics. *Nature*. https://doi.org/10.1038/nature12450

Takemura, S. Y., Karuppudurai, T., Ting, C. Y., Lu, Z., Lee, C. H., & Meinertzhagen, I. A. (2011). Cholinergic circuits integrate neighboring visual signals in a drosophila motion detection pathway. *Current Biology*, *21*(24), 2077–2084. https://doi.org/10.1016/j.cub.2011.10.053

Takemura, S. Y., Xu, C. S., Lu, Z., Rivlin, P. K., Parag, T., Olbris, D. J., Plaza, S., Zhao, T., Katz, W. T., Umayam, L., Weaver, C., Hess, H. F., Horne, J. A., Nunez-Iglesias, J., Aniceto, R., Chang, L. A., Lauchie, S., Nasca, A., Ogundeyi, O., … Scheffer, L. K. (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), 13711–13716. https://doi.org/10.1073/pnas.1509820112

Tang, S., & Juusola, M. (2010). Intrinsic Activity in the Fly Brain Gates Visual Information during Behavioral Choices. *PLoS ONE*, *5*(12), e14455. https://doi.org/10.1371/journal.pone.0014455

Taylor, G. J., Paulk, A. C., Pearson, T. W. J., Moore, R. J. D., Stacey, J. A., Ball, D., Van Swinderen, B., & Srinivasan, M. V. (2015). Insects modify their behaviour depending on the feedback sensor used when walking on a trackball in virtual reality. *Journal of Experimental Biology*. https://doi.org/10.1242/jeb.125617

Troje, N. (1993). Spectral categories in the learning behaviour of blowflies. *Zeitschrift Fur Naturforschung - Section C Journal of Biosciences*, *48*(1–2), 96–104. https://doi.org/10.1515/znc-1993-1-218

Tschopp, F. D., Reiser, M. B., & Turaga, S. C. (2018). A connectome based hexagonal lattice convolutional network model of the drosophila visual system. In *arXiv*.

Tseng, E., Colburn, S., Whitehead, J., Huang, L., Baek, S. H., Majumdar, A., & Heide, F. (2021). Neural nano-optics for high-quality thin lens imaging. *Nature Communications*, *12*(1), 1–7. https://doi.org/10.1038/s41467-021-26443-0

Tuthill, J. C., Nern, A., Holtz, S. L., Rubin, G. M., & Reiser, M. B. (2013a). Contributions of the 12 neuron classes in the fly lamina to motion vision. *Neuron*, *79*(1), 128–140. https://doi.org/10.1016/J.NEURON.2013.05.024

Tuthill, J. C., Nern, A., Holtz, S. L., Rubin, G. M., & Reiser, M. B. (2013b). Contributions of the 12 Neuron Classes in the Fly Lamina to Motion Vision. *Neuron*, *79*(1), 128–140. https://doi.org/10.1016/j.neuron.2013.05.024

van Hateren, J. H. (1984). Waveguide theory applied to optically measured angular sensitivities of fly photoreceptors. *Journal of Comparative Physiology A*, *154*(6), 761–771. https://doi.org/10.1007/BF00610676

van Hateren, J. H. (1987). Neural superposition and oscillations in the eye of the blowfly. *Journal of Comparative Physiology A*. https://doi.org/10.1007/BF00610226

van Hateren, J. H. (1992). Theoretical predictions of spatiotemporal receptive fields of fly LMCs, and experimental validation. *Journal of Comparative Physiology A*, *171*(2), 157–170. https://doi.org/10.1007/BF00188924

Wardill, T. J., List, O., Li, X., Dongre, S., McCulloch, M., Ting, C. Y., O'Kane, C. J., Tang, S., Lee, C. H., Hardie, R. C., & Juusola, M. (2012). Multiple spectral inputs improve motion discrimination in the drosophila visual system. *Science*, *336*(6083), 925–931. https://doi.org/10.1126/science.1215317

Warrant, E. J., & McIntyre, P. D. (1993). Arthropod eye design and the physical limits to spatial resolving power. In *Progress in Neurobiology* (Vol. 40, Issue 4, pp. 413–461). Prog Neurobiol. https://doi.org/10.1016/0301-0082(93)90017-M

Washburn, T., & O'Tousa, J. E. (1989). Molecular defects in Drosophila rhodopsin mutants. *The Journal of Biological Chemistry*, *264*(26), 15464–15466. https://doi.org/10.1016/S0021-9258(19)84852-4

Weckstrom, M., Juusola, M., & Laughlin, S. B. (1992). Presynaptic enhancement of signal transients in photoreceptor terminals in the compound eye. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, *250*(1327), 83–89. https://doi.org/10.1098/rspb.1992.0134

Wei, H., Kyung, H. Y., Kim, P. J., & Desplan, C. (2020). The diversity of lobula plate tangential cells (LPTCs) in the Drosophila motion vision system. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, *206*(2), 139–148. https://doi.org/10.1007/s00359-019-01380-y

Wernet, M. F., Labhart, T., Baumann, F., Mazzoni, E. O., Pichaud, F., & Desplan, C. (2003). Homothorax switches function of Drosophila photoreceptors from color to polarized light sensors. *Cell*, *115*(3), 267–279. https://doi.org/10.1016/S0092-8674(03)00848-1

Wijngaard, W., & Stavenga, D. G. (1975). On optical crosstalk between fly rhabdomeres. *Biological Cybernetics*, *18*(2), 61–67. https://doi.org/10.1007/BF00337126

Wilson, T. (1985). Scanning optical microscopy. *Scanning*, *7*(2), 79–87. https://doi.org/10.1002/sca.4950070203

Yamaguchi, S., Wolf, R., Desplan, C., & Heisenberg, M. (2008). Motion vision is independent of color in Drosophila. *Proceedings of the National Academy of Sciences of the United States of America*, *105*(12), 4910–4915. https://doi.org/10.1073/pnas.0711484105

Yang, H. H. H., St-Pierre, F., Sun, X., Ding, X., Lin, M. Z. Z., & Clandinin, T. R. R. (2016). Subcellular Imaging of Voltage and Calcium Signals Reveals Neural Processing In Vivo. *Cell*. https://doi.org/10.1016/j.cell.2016.05.031

Zheng, L., De Polavieja, G. G., Wolfram, V., Asyali, M. H., Hardie, R. C., & Juusola, M. (2006). Feedback Network Controls Photoreceptor Output at the Layer of First Visual Synapses in Drosophila. *The Journal of General Physiology*, *127*(5), 495. https://doi.org/10.1085/JGP.200509470