NEURONAL ENCODING OF NATURAL IMAGERY IN DRAGONFLY MOTION PATHWAYS

BERNARD JOHN ESSEX EVANS

Discipline of Physiology
Adelaide Medical School
The University of Adelaide

August 2018

A thesis submitted for the degree of Doctor of Philosophy

TABLE OF CONTENTS

DEDICATIONS

I thank my supervisors, Steven Wiederman and David O'Carroll for their guidance, support and especially their patience. They have proven invaluable sources of knowledge and wisdom throughout my candidature.

I additionally thank my colleagues (of whom there were many) for their help, support and company throughout my candidature. Without them, it would have been not only harder but less fun. In an intentionally haphazard order (RNG):

- Benjamin Lancer
- Ben Parslow
- Benjamin Harvey
- Zahra Bagheri
- Sam Polacek
- Elisa Rigosi
- Joseph Fabian
- Matthew Schwartz
- Bo Bekkouche
- John James
- Plus, the 3rd years

More Importantly (which begs the question of being last, maybe larger font would help)

I would like to thank most of all, my wife Echo and my son Bertie. They have been nothing but supportive during my entire endeavour, have done their best to keep me sane (enough) and have shown a level of tolerance beyond what anyone could ask for.

Much love and too many hugs.

Bernard

1 Prelude

1.1 Abstract

Vision is the primary sense of humans and most other animals. While the act of seeing seems easy, the neuronal architectures that underlie this ability are some of the most complex of the brain. Insects represent an excellent model for investigating how vision operates as they often lead rich visual lives while possessing relatively simple brains. Among insects, aerial predators such as the dragonfly face additional survival tasks. Not only must aerial predators successfully navigate three-dimensional visual environments, they must also be able to identify and track their prey. This task is made even more difficult due to the complexity of visual scenes that contain detail on all scales of magnification, making the job of the predator particularly challenging.

Here I investigate the physiology of neurons accessible through tracts in the third neuropil of the optic lobe of the dragonfly. It is at this stage of processing that the first evidence of both wide-field motion and object detection emerges. My research extends the current understanding of two main pathways in the dragonfly visual system, the wide-field motion pathway and target-tracking pathway.

While wide-field motion pathways have been studied in numerous insects, until now the dragonfly wide-field motion pathway remains unstudied. Investigation of this pathway has revealed properties, novel among insects, specifically the purely optical adaptation to motion at both high and low velocities through motion adaptation. Here I characterise these newly described neurons and investigate their adaptation properties.

The dragonfly target-tracking pathway has been studied extensively, but most research has focussed on classical stimuli such as gratings and small black objects moving on white monitors. Here I extend previous research, which characterised the behaviour of target tracking neurons in cluttered environments, developing a paradigm to allow numerous properties of targets to be changed while still measuring tracking performance. I show that dragonfly neurons interact with clutter through the previously discovered selective attention system, treating cluttered scenes as collections of target-like features. I further show that this system uses the direction and speed of the target and background as one of the key parameters for tracking success. I also elucidate some additional properties of selective attention including the capacity to select for inhibitory targets or weakly salient features in preference to strongly excitatory ones. In collaboration with colleagues, I have also

performed some limited modelling to demonstrate that a selective attention model, which includes switching best explains experimental data.

Finally, I explore a mathematical model called divisive normalisation which may partially explain how neurons with large receptive fields can be used to re-establish target position information (lost in a position invariant system) through relatively simple integrations of multiple large receptive field neurons.

In summary, my thesis provides a broad investigation into several questions about how dragonflies can function in natural environments. More broadly, my thesis addresses general questions about vision and how complicated visual tasks can be solved via clever strategies employed in neuronal systems and their modelled equivalents.

1.2 Thesis Format

The body of this thesis represents a thesis by combination. Where published, works appear as originally published (with minor formatting changes). Other work has either been presented as manuscript papers (Chapters 3, 5 & 6) or as collections of data (Chapter 4). Other published work completed during the PhD but related more to a previous direction have been included as appendices (Chapters 9 & 10).

1.3 Thesis Declaration

I certify that this work contains no material which has been accepted for the award of any other degree or diploma in my name, in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission in my name, for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint-award of this degree.

I give consent to this copy of my thesis when deposited in the University Library, being made available for loan and photocopying, subject to the provisions of the Copyright Act 1968.

I acknowledge that copyright of published works contained within this thesis resides with the copyright holder(s) of those works.

I also give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library Search and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

I acknowledge the support I have received for my research through the provision of an Australian Government Research Training Program Scholarship.

1.4 Acronyms

BLC	Barlow-Levick Correlator
EMD	Elementary Motion Detector
GPCR	G-Protein Coupled Receptor
HRC	Hassenstein Reichardt Correlator
LPTC	Lobula Plate Tangential Cell
LTC	Lobula Tangential Cell
MLG	Male Lobula Giant
ND	Null Direction (Non-preferred direction of motion)
PD	Preferred Direction
STMD	Small Target Motion Detector
WFMS	Widefield Motion Sensitive

2 Introduction

2.1 Preliminary Remarks

Where to start? This is perhaps the apt question to begin any discussion on any topic. Of course, self-reference is rarely the correct answer. The difficulty in describing any body of work, is that the formation of the work and presentation of said work conform to frameworks which stand in fundamental opposition to one another. Information captured inside the brain of any creature is encapsulated in a vast array of interconnections, references and networks. Information captured within a text is subject to the constraint of linearity. How one turns a circle into a line or vice versa is matter of topological impossibility, but none-the-less, we must endeavour.

In this way, I sympathise with my subject, *Hemicordulia tau* a small hawking dragonfly of the family Odonata. Whilst it has but a small brain, it is by no means a simple creature. Even scratching the surface reveals vast numbers of neurons that interact with each other in surprising and often opaque ways. How it views the many phenomena that constitute this work is anyone's guess. Any attempt to elucidate the complexity of its rich internal world onto the page intrinsically diminishes its exquisiteness, but none-the-less, we must endeavour.

So, I return to 'where to start' (my first circle complete). What is more fundamental, anatomy, physiology or behaviour (what, how or why)? Surely, they must be explained in concert for what has no meaning without why, and why no explanation without how. Of course, there is also modelling which fortunately has its place subsequent to these (except of course where it underlies our understanding of these systems). I have therefore decided to segregate these topics as best I can, describing the particular physiology of specific anatomical units under their corresponding anatomy but leaving broader physiological discussions separate.

Broadly speaking, the theme that unifies this work is 1) the dragonfly, 2) vision, 3) natural environments & 4) explaining the previous three computationally. Within that rough framework the following topics were investigated in depth:

1. Wide field optic flow
2. Target discrimination and tracking
3. Recovering neuronal information using theoretical models

2.2 Behaviour of Insects

Invertebrates are, due to their vast phenotypical differences from humans, often treated as a different type of creature than their vertebrate cousins. Certainly, distance from humanity is often conflated with distance from intelligence. It is rare for anyone to properly identify with insects and this distance can lead to an intellectual disdain for 'vermin'. Yet, as recent studies have demonstrated, insects possess many of the cognitive faculties of 'higher-order' animals. While many insects exhibit only simplistic behaviour, dragonflies exhibit evidence for internal models of external objects (Mischiati & Lin 2014) and bees even have language (von Frisch 1967).

Thus, insects represent a resource for interrogating many kinds of behaviour and perhaps more importantly, the neuronal networks which underpin them. Each behaviour, whether it is a simple orienting reflex, or a more complex internal model-based target pursuit indicates the existence of a network of neurons designed to encode and facilitate that behaviour.

2.2.1 Orienting Reflexes

One of the simplest forms of behaviour is the capacity of an animal to attain its preferred location in the world. This can be as simple as a bacterium following a chemical gradient, to a wasp choosing a nesting site. In insects three of the simplest behavioural assays involve orienting reflexes, one a simple attempt to maintain the status quo (i.e. stability), a second a preference for facing and the third for avoidance.

2.2.1.1 Orienting for Stabilization

Many animals including insects perform rotational movements in response to wide-field visual motion. This motion 'corrects' for the wide-field motion to establish a non-moving surrounding (i.e. turning with the background motion). This 'optomotor reflex' is used to interrogate an insect's perception of wide-field motion (Gotz 1964, Haag & Borst 1997, Hassenstein & Reichardt 1956, Leonhardt et al., 2016). For example, using the torque generated from a fixed animal's attempts to move (using a free-moving ball for walking or wing-beat analysis for flight) it is possible to interrogate the parameters that influence the optomotor reflex. One example is that the size and contrast of a pattern generating optic flow directly affect the size of the response (Reichardt & Poggio 1983).

A freely moving insect (i.e. closed loop) it is also possible to interrogate these behaviours by using an enclosure that updates its visual stimuli dependent on the insect's behaviour potentially yielding differences between active and passive responses from the animal

(Reichardt & Wenking 1969). The optomotor reflex has been used extensively in tandem with recordings from neurons from the optic lobe. Correlations found between these two systems have been used to argue functional relevance for the neurons, specifically neurons that respond to large motion such as Lobula Plate Tangential Cells (LPTCs) in flies (Tuthill et al., 2011; Bahl et al., 2013; Haikala et al., 2013). These behavioural techniques can also be used to determine what kinds of motion, such as second-order motion (described below, Theobald et al., 2008) are visible to an insect.

2.2.1.2 Orienting to a Bar

In addition to the stabilization of optic flow, flies have also been shown to have a second reflex, which orients them towards individual moving bars (Reichardt & Wenking 1969). This reflex requires the bar to be in motion, as stationary bars elicit no response (Bahl et al., 2013). The reflex instead requires the luminance of the bar to change, which also occurs when a bar is in motion across the subject's retina (Bahl et al., 2013). In flies, this reflex has been shown to function entirely independently (though at reduced efficacy) from the wide-field motion-detection system and hence the optomotor reflex and likely is generated through a second pathway (Bahl et al., 2013).

2.2.1.3 Looming Reflexes

An object looming towards you can elicit two general responses, based on whether it is expected or not. In the case of flying insects an expected looming object is likely one upon which one wants to land. Thus such insects will attempt to slow their flight to tailor it to the rate of expansion so as to avoid impact (Borst & Bahde 1986).

Another cause of a looming stimulus is an object to be avoided, a predator or simply an object that may cause a fatal collision. Insects confronted with looming stimuli have been shown to exhibit avoidance behaviours (Robertson & Johnson 1993; Fotowat & Gabbiani 2011) reflected in their neuronal code (Rind & Simmons 1992). This generates a rather inconvenient ambiguity, namely, that looming stimuli and progressive stimuli (which contain several common features) can be confused. When both are present, both stimuli contribute to motion and summate in a non-linear manner (Reiser & Dickinson 2010).

2.2.2 Behavioural Strategies

In addition to basic reflexes, there are also behaviours that insects exhibit, which impacts their ability to see. The simplest example of this is attempting to remain motionless, reducing erroneous motion information induced by ego-motion. These strategies give some insight into

how motion systems work as a subject that performs a behaviour that minimizes or maximizes some aspect of a scene gives insight into what aspects of visual information are important to the subject. Here I discuss two flight strategies, the more general positioning strategies and strategies for prey selection.

2.2.2.1 Saccadic Turns

There are two basic forms of movements, rotational and translational. In the former, the subject simply rotates around a single axis resulting in a predictable motion where some points rotate and others exhibit large lateral motion. The second, translational is far more complex, resulting in parallax (i.e. differential motion of nearby and distant objects). The optic flow induced by translation, while more complex results in less overall motion energy. As such insects have established strategies to break these two motions apart.

Blowflies exhibit a saccadic separation of translational and rotational movements during flight (van Hateren & Schilstra 1999, Mongeau & Frye 2017). This is achieved during long flights by variations of thoracic and head movement. The thorax is kept largely still with the head used for gaze stabilization punctuated by periodic saccadic movements with the body following the head before re-establishing the stable flight mode. These saccadic movements are enabled via the production of efference copies, neuronal signals that suppress responses from gaze-stabilization pathways to allow for voluntary turns (Kim et al. 2015). This has also been shown to happen in bumblebees (Mertes et al 2014) and honeybees (Boeddeker et al., 2010).

The qualities of background information also contribute to this behaviour. For example, it was shown that the Wide-field Motion Sensitive (WFMS) system in bees responds to landmarks and may be used for landmark navigation to preferred sites (Mertes et al 2014). Additionally, flies exhibit distinct saccadic turns to bars but smoother motions to panoramas (Mongeau & Frye 2017) indicating a disjuncture between the object and ego-motion system as is seen in orienting reflexes mentioned above.

2.2.2.2 Proportional Navigation

Dragonflies are spectacular hunters, boasting successful capture rates from 90% (Combes et al., 2013) to 97% (Olberg et al., 2000, Combes et al., 2013). Much of this is down to the brute strength, which dragonflies bring to the insect world. However, pursuit strategy also plays an important role.

The simplest form of pursuit is to simply chase the target's current location (Figure 1 – left). However, this strategy is inefficient in multiple ways. Firstly, the path taken is more circuitous and thus longer. This strategy also involves more turning manoeuvres, which may add further energy inefficiencies.

Another computationally simple strategy is proportional navigation. This strategy works for two targets maintaining constant velocity vectors (Figure 1 - right). In these circumstances, maintaining a constant sightline angle (measured axially from the direction of motion of the chaser) results in optimal interception. Should the target change direction or speed, this will result in a drift in sightline, which can be corrected via either an increase or decrease in speed or a change in heading. While some species employ zero-degree proportional navigation (Land & Collett 1973) or precalculated trajectories (Collett & Land 1978), dragonflies have been shown to use a modified version of proportional navigation by maintaining central view of the target and using head-movements to both foveate and maintain constant sightline (Olberg et al., 2007, Mischiati & Lin et al., 2015). The speed of this foveation is surprisingly fast (30 ms).

Perhaps the most interesting benefit of a non-zero proportional navigation strategy is that it allows a slower object to intercept a faster one.

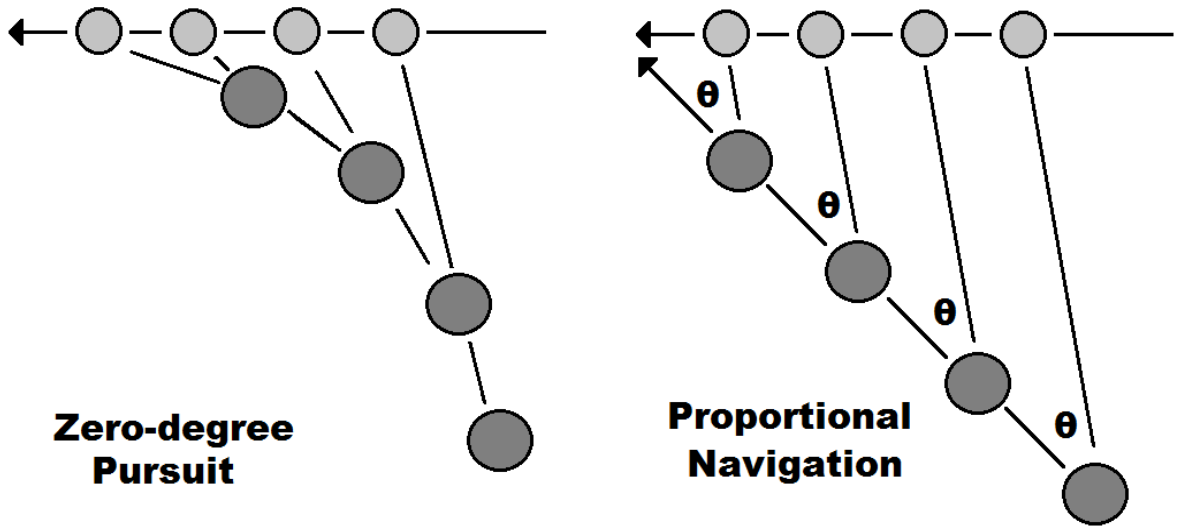

Figure 1: Zero-degree Pursuit vs. Proportional Navigation (dark grey predator, light grey prey): Left: Simply directing oneself at the current location of a target results in a progressive changing of heading and a more circuitous path. Right: For two objects moving at constant velocity, a simple strategy for determining the minimum time/distance to interception is to maintain a constant sightline bearing. Any slippage in bearing can be corrected via increase in speed or change in angle of approach. Such an approach requires no internal model of the target's location.

2.2.3 Positioning Strategies

Waiting in a favourable location has several benefits including preferential mate selection and food abundance. Beyond this, choosing a location also affects the visual input an insect observes. Dragonflies can be roughly split into two types based on behaviour, perching and hawking (Corbet & May 2008). In both these cases, the dragonfly minimizes the amount of motion information present in a scene by maintaining a stationary pose prior to initiating prey pursuits.

However, this doesn't answer the question of why this pond, why this frond, why this perch? It has been observed in dragonflies that they tend to pick identical perches more often, even over several days (Baird & May 1997). Some of the reasons are as mentioned above such as easy view of intruders (Switzer & Eason 2000) or the predictable availability of prey (Baird & May 1997). However, visual constraints also apply, with dragonflies tending to prefer locations with low shade and hence high-contrast (Remsburg et al., 2008) or locations where their dorsal eye is directed towards sky (and away from the sun) (Sauseng et al., 2003), thus, improving their catch rates based on purely vision-based grounds.

2.2.4 Selection Strategies

Predators by nature are required to be picky about their food. There are limitations on what constitutes prey (governed by size and speed). Moreover, any predation activity necessarily detracts from other behaviour such as territorial guarding and mate pursuit. To this end, one might imagine that predatory insects exhibit prey selection preferences such as only pursuing nearby prey (to save energy), slow prey (to maximize success) and larger prey (energy efficiency). However, it would appear dragonflies are largely bound by their detection capabilities rather than any predator/prey specific rules (Combes et al., 2013).

In general, the following principles apply. Perching dragonflies do not pursue objects that subtend too large a visual angle (3-4°). Beyond that, their pursuit initiations appear mostly to be governed by the visual size and velocity of the target (Combes et al., 2013). Thus, perching dragonflies have been shown to pursue larger targets which are further away and move faster rather than waiting for them to come closer (Combes et al., 2013). While this research is limited to perching dragonflies, it is likely that many of these same constraints apply to hawking dragonflies.

2.3 Anatomy of Insect Vision

The following section discusses the general anatomical features of visual pathways from basic eye structure to higher-order neurons involved in processing visual stimuli.

2.3.1 Eye Types

The first consideration of any visual system is the architecture of how various elements of the eye are brought together to form a cohesive unit. There are numerous structures of eyes ranging from simple concave pits to fully functional motile eyes like those found in humans. The design of the eye affects how its outputs can be further processed. While there are numerous architectures for eyes, only three are discussed here:

- Compound Eyes
- Superposition Eyes
- Lens Eyes

2.3.1.1 Compound Eyes

Compound eyes function by having individual capsules called ommatidia, each of which has light-sensitive cells (photoreceptors) at the base (Figure 2) surrounded by pigmented cells. These eyes each possess their own cornea and lens and can be conceived as independent units. Light is funnelled through the cornea (which focuses light) down through the rhabdom. The rhabdom acts as a transparent tube (or light guide) allowing maximum stimulation of the photoreceptors. Surrounding the ommatidia are pigmented cells. These cells are necessary to shield ommatidia and their photoreceptors from sources of light other than their cornea. In the event light could propagate tangentially through the ommatidia, its ability to resolve position would be compromised.

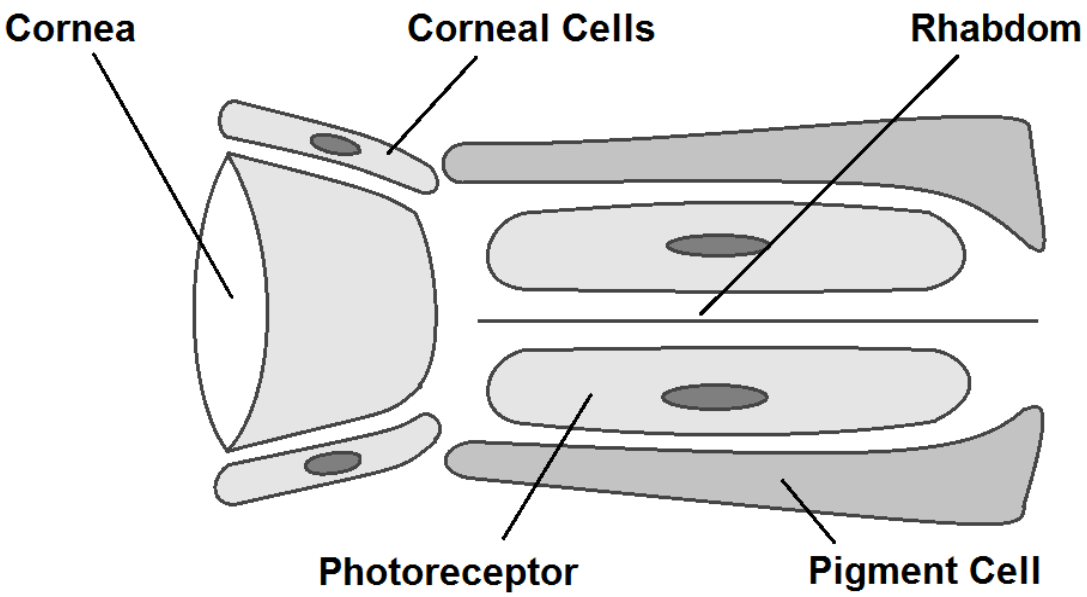

Figure 2: Example of a single ommatidia from a compound eye. Each ommatidia consists of its own cornea, corneal cells, rhabdom (light guide), photoreceptors (light detection) and pigment cells (shielding).

To provide a broad field of view, a compound eye requires hundreds (or thousands) of such ommatidia, each angled to regard a different point in space. The angular precision (i.e. resolution) of a compound eye is subject to two broad limitations. Firstly, the size of each individual ommatidia and its corresponding cornea limits the smallest area of space that the ommatidia can observe. To capture sufficient light from a small region of space requires the maximum number of photons to be absorbed. This tends towards a preference for larger corneas capable of aggregating more photons from a direction before focussing them into an ommatidium.

In addition, the number of distinguishable points in space is limited by the quantity and hence density of ommatidia. To have high resolution of a small region of space, the ommatidia must be arranged in such a manner that their angular differences are small (i.e. they are almost parallel). For example, dragonflies such as *Hemicordulia tau* can have an angular resolution less than 0.5° compared to *Drosophila Melanogaster*'s 5° representing a 10-fold improvement in visual resolution. Of course, dragonflies do not have uniquely acute vision among insects with even the humble honeybee drone boasting 0.6° (Rigosi et al., 2017).

These two constraints are in direct conflict with one another and lead to a few interesting outcomes. Firstly, as the desired resolution increases, both the quantity and size of ommatidia increase leading to an R^2 relationship (a doubling of resolution requires a quadrupling of eye

size, Land 1989). In essence, to achieve a very high resolution requires a very large eye making the compound eye an inherently limited design.

To compensate for this weakness, many insects have developed so-called ‘hot-spots’ (Land 1989). These regions of higher acuity result from a flattening of the eye (allowing more ommatidia to view a single region in space). These regions are quite analogous to the fovea in human vision. These regions often map to behaviourally relevant regions of an insect’s field of view such as for mate and prey identification.

While compound eye design is problematic from an evolutionary standpoint, it has a few incidental benefits for scientists, namely that individual ommatidia can be stimulated by point sources of light in a way that other kinds of eyes cannot. This serendipitously allows for a more detailed scrutiny of the behaviour of individual photoreceptor groups and their interactions with adjacent and downstream neurons.

2.3.1.2 Superposition Eyes

One structure designed to improve light sensitivity in compound eyes is the so-called neural superpositional eye. Such a structure takes light from several adjacent ommatidia and neurally combines them while maintaining their retinotopic pattern. Thus, each point in space is sampled through six independent ommatidia as found in Drosophila (Hardie 1986). This allows for more light to be collected by each photoreceptor, hence improving light sensitivity.

2.3.1.3 Lens Eyes

Lens eyes like those found in human and other vertebrates are an especially efficient design for eyes. The principle of such eye design is to pack as many light-sensitive cells into a small area and then using a refractive lens to focus light from the external environment onto the light-sensitive cells (Figure 3). The lack of a need for pigment cells and the additional acuity granted by the lens focusing the light allows these eye designs to be far smaller than their compound counterparts and are generally more efficient. However, both eyes structures conform to similar topological descriptions and thus are functionally very similar for the purposes of down-stream processing of light information.

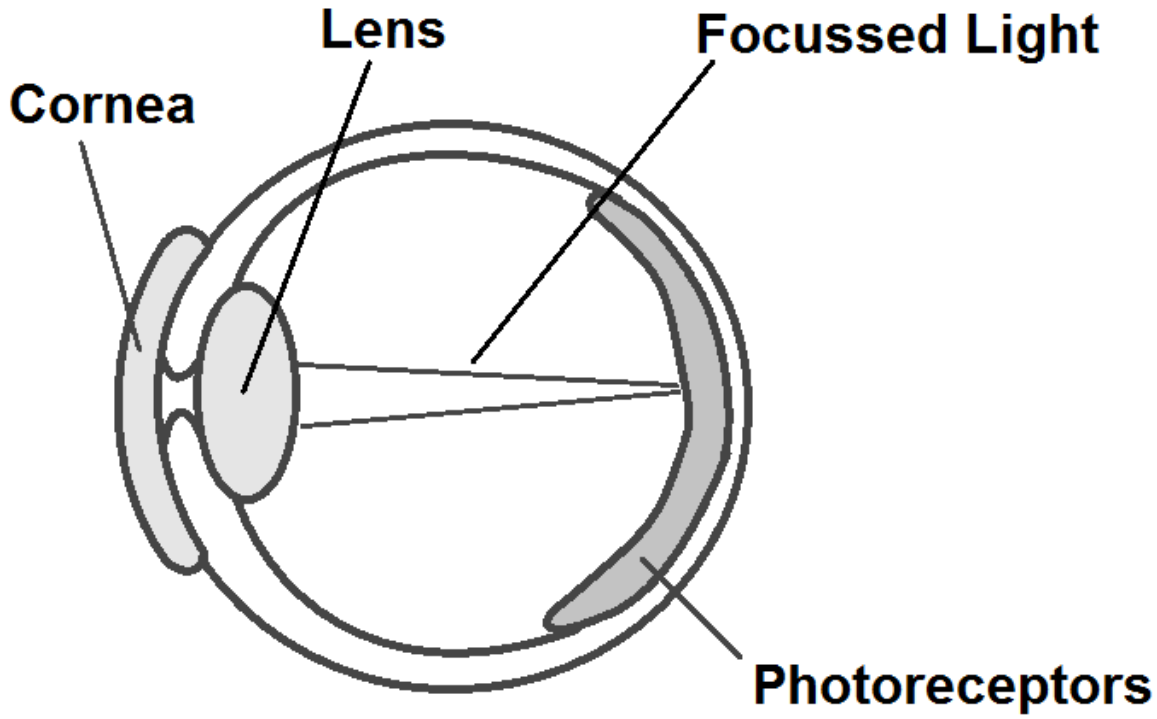

Figure 3: Example of a lens-based eye. Light is gathered through the cornea and focussed using the lens onto the photoreceptors at the back of the eye. Due to the focussed lens and lack of need for pigment cells, photoreceptors can be packed far more tightly.

2.3.2 Eye Structure in Dragonflies

The eye of *Hemicordulia tau* is divided into three regions, which apart from their visual dissimilarity also exhibit retinular differences. The dorsal region (red) has larger facets, while the ventral region (blue) has smaller facets and darker screening pigments (Laughlin & McGinness 1978). The third peripheral region (yellow) lies along the posterior margin of the eye.

In addition to these regional differences, *Hemicordulia Tau* boasts three separate acute regions in the eyes, dorsal, anterior and lateral (Horridge 1978). The dorsal acute region corresponds to *Hemicordulia's* hunting region (dragonflies typically hunt from below). The frontal region is presumed to be involved in inspecting objects. These acute zones boast far greater resolution through a combination of larger facets and smaller interommatidial angles. The effective resolution (as defined in minimal angle) ranges from 2° in the back of the eye to 0.3° in the acute region (Horridge 1978). Intermediate regions have approximately 0.5° resolution (Horridge 1978) which is very high for an insect and much better than flies and other commonly studied insects.

Hemicordulia also has a significant binocular overlap ~10° which differs from some other species of dragonfly (Horridge 1978). Finally, while the *Hemicordulia* has an acute-zone resolution of 0.3-0.5°, the acceptance angle of individual ommatidia are 0.7° (Horridge 1978) indicatinga fair degree of overlap between adjacent ommatidia.

2.3.3 Binocular Vision

One of the easiest methods for establishing depth in a scene is by using stereovision. By taking slightly different viewpoints of a scene, the three-dimensional nature of the scene can be readily computed. While this is very effective in larger creatures, the distance between the two eyes places a hard limit on the maximum range this approach can achieve. Dragonflies can estimate distance (depending on species) up to 70-100 cm (Olberg et al., 2005).

2.3.4 Visual Pathways

The visual pathway in insects involves a series of neuropils within the optic lobe. The general theory is that information flows through each layer sequentially, becoming more and more integrated and higher-level the more proximal to the insect's brain. The brain neuropils involved are as follows:

- Retina Light detection
- Lamina Light change detection
- Medulla Local movement detection
- Lobula Complex Global movement detection

2.3.4.1 Visual Pathways in Flies

In keeping with the above, the fly brain follows a similar pattern: Retina → Lamina → Medulla → Lobula. Figure 4 details the proposed circuitry underlying Lobula Plate Tangential Cells (LPTCs) in flies. Though this is not the only visual pathway which has been fully elucidated, it is the most detailed available. It should at the least, disavow any notion that insect vision is simple.

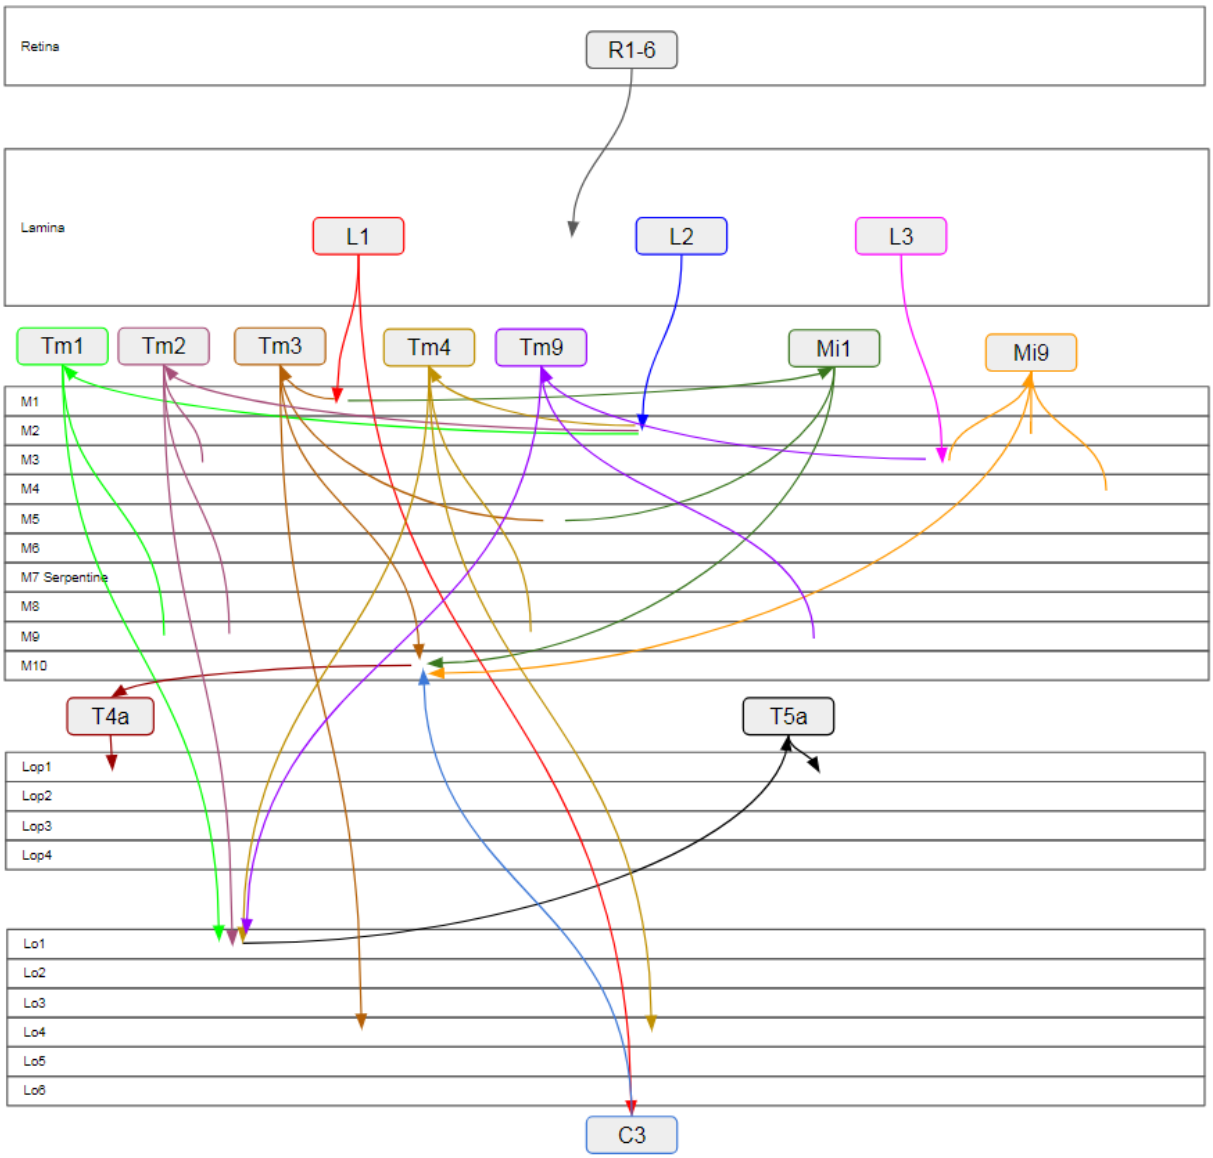

Figure 4: Schematic of the Fly Brain Layers and Neurons involved in motion. (Based on Fischbach & Dittrich 1989, Takemura et al., 2017 and Arenz et al., 2017)

2.3.5 Retina

Photoreceptors are the point of entry into all motion pathways. These cells transduce photons (largely via proteins called opsins) into electrical potentials. These electrical potential changes represent the first neuronal representation of light. Interestingly, the variations between photoreceptors among different species are almost more striking than the superficial differences of eye architecture.

Invertebrate photoreceptors use entirely different photo-transduction pathways to vertebrate eyes. The most striking difference is that while vertebrate photoreceptors close ion channels (thus hyperpolarising) in response to light, invertebrate photoreceptors do the opposite,

depolarising to light (Fain et al., 2010, Hardie & Raghu 2001). In addition, the physical substructure of photoreceptors differs between species.

Opsins are membrane-bound proteins and as such are limited to the outer membrane (or membrane like structures) of cells. To maximize the absorption of light (and hence light sensitivity) it is necessary to maximize the total available membrane (i.e. surface area). Vertebrate and invertebrates achieve this goal in different ways. Vertebrate rod cells contain numerous internal disks, connected to the external membrane via cilium. Each disk possesses numerous opsin proteins (Figure 5A). Invertebrate vision instead relies on microvilli (Figure 5B) which protrude from the membrane and form the internal rhabdomere (Fain et al., 2010, Hardie & Raghu 2001, Rister & Desplan 2011).

There are advantages and disadvantages to both designs. First and foremost, microvilli out-perform cilium-based structures in very dim conditions whilst maintaining good performance in bright conditions enabling better dynamic range. Vertebrates, which employ ciliary photoreceptors account for this weakness by augmenting rod-cells with cone cells. Rod cells are slow and specialised for low-light conditions, while cone cells are better suited to bright light conditions. This specialisation represents a space-saving benefit with cones and rods more efficiently counting photons in a space-sense (Fain et al., 2010).

Most research in the invertebrate retina has been performed in flies and so specific details of insects will here be limited to flies and dragonflies (the research animal examined in this thesis).

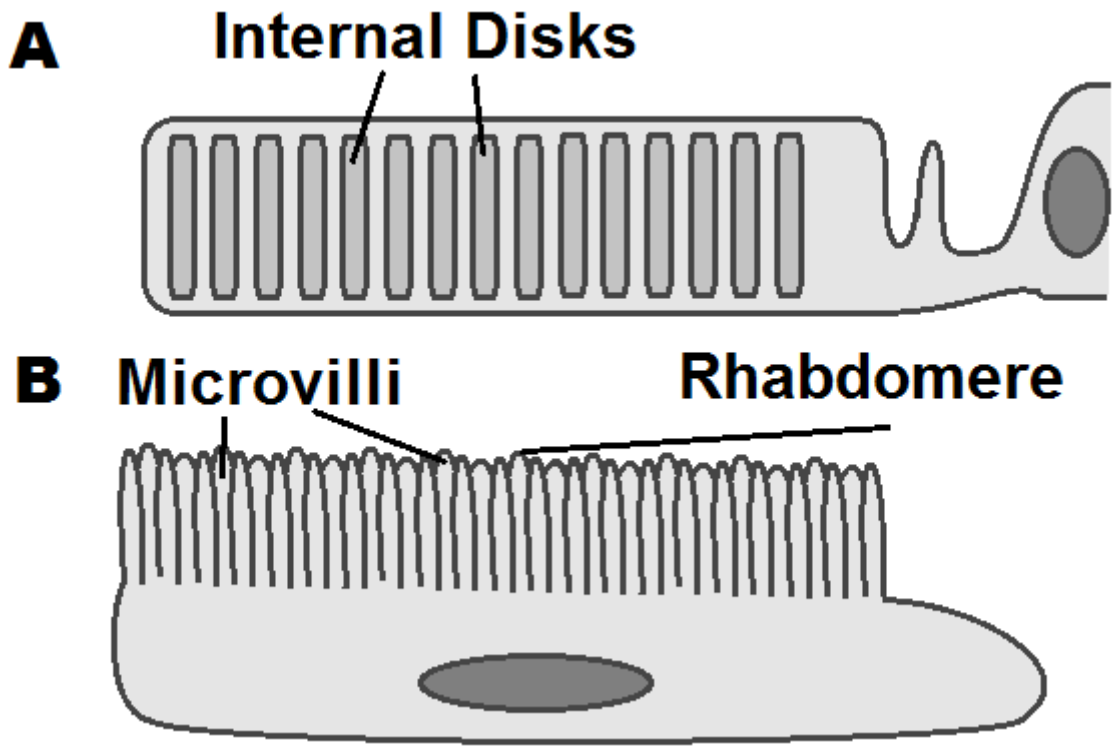

Figure 5: Differences between vertebrate rods and invertebrate photoreceptors. A) Vertebrate rods possess internal disks which carry membrane-bound rhodopsins. B) Invertebrate photoreceptors have many microvilli processes which form the rhabdomere in invertebrates. These microvilli also possess numerous membrane-bound opsin proteins.

2.3.5.1 Phototransduction

As previously mentioned, photo-transduction is performed via the use of opsins. These G-protein coupled receptors (GPCRs) contain two major components, one of which is retinal, the chemical responsible for light transduction. Light transduction occurs when a photon is absorbed by retinal causing a conformation change from cis-form to a trans-form. This change occurs in less than 200 fs (Schoenlein 1991). This change is sufficient to kick-off a secondary messenger cascade (via cGMP) which results in the change in membrane potential.

The differences between vertebrates and invertebrates also exist at the level of opsin proteins. Firstly, the different structures (cilium vs microvilli) place constraints on the opsins present. Cilium (i.e. rods) express c-opsins, while microvilli (i.e. invertebrates) express r-opsins (Fain et al., 2010) indicating that even at the level of the protein, there are significant differences between vertebrate and invertebrate vision.

The chemical cascades which transduce light in invertebrates are significantly faster (up to 10 times so in Drosophila, Hardie & Raghu 2001) in part because of the microvilli structure which promotes diffusion (Fain et al., 2010) and in part as r-opsins allow for photon-

absorption in their trans-form state (Stavenga 1995) allowing light detection speeds of up to 300 Hz in flies (Tatler et al., 2000).

These photoreceptors also exhibit a much higher dynamic range than vertebrates (which separate dim and bright conditions using rods and cones respectively) via large changes in intracellular calcium (Fain et al., 2010). Drosophila photoreceptors have been shown to have a dynamic range up to 10^6 while locusts have been shown capable of detecting single photons (Laughlin & Lillywhite 1982).

2.3.5.2 Colour Sensitivity

Most invertebrates possess photoreceptors which are sensitive to different spectra of light. The specific wavelengths of greatest interest vary from species to species (Laughlin 1976, Yang & Osorio 1991, Arikawa et al., 1999) however two general bands which appear to be held commonly among many insect species are sensitivity to green and UV light. It is proposed (Pichaud et al., 1999) that this general distinction represents a simplistic behavioural advantage. UV light does not reflect off objects in a scene and thus UV light is representative of open areas, while green light is more representative of regions rich in food (Pichaud et al., 1999).

Many invertebrates expand on these two conserved wavelengths, including various blue-light opsins, however red-sensitivity is rare in insects (Pichaud et al., 1999) and not present in flies (Rister & Desplan 2011) or dragonflies (Yang & Osorio 1991) the two insects most relevant to this study.

The detection of colour vision requires at least two different receptors centred at different wavelengths. In many species (including flies and mice, Rister & Desplan 2011) photoreceptors of different frequencies arise from common gene sets but exhibit stochastic randomness on expression resulting in random mosaics (Rister & Desplan 2011).

Finally, although it has been demonstrated that flies can discriminate coloured and polarized light, it has not been formally demonstrated that they have colour vision (Pichaud et al., 1999).

2.3.5.3 Polarization Cues

Due to the wave-like nature of light, it is possible for light to be polarized (Figure 6). Numerous insects can detect these polarization cues and use them for tasks as diverse as navigation (Homberg 2004) and habitat selection. Dragonflies have been shown to use polarization cues to determine their rendezvous and oviposition sites (Wildermuth 1998).

These systems can also malfunction. For example, dragonflies were shown to prefer polished black gravestones and behave as if near water. The gravestones were shown to reflect polarized light similarly to water indicating this may represent a visual cue used for territory discrimination in dragonflies (Horvarth et al., 2007).

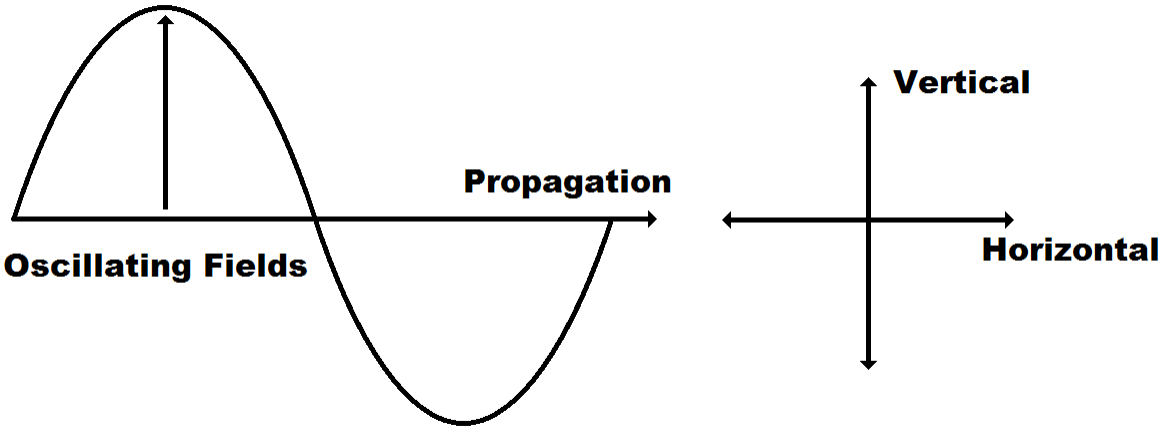

Figure 6: *Polarization of Light. Light can be oriented such that the main axis of oscillation is limited to a single dimension. The polarization of light in natural environments can indicate the presence of certain surfaces.*

2.3.5.4 Retina Cells of the Dragonfly

In *Hemicordulia tau*, there are eight photoreceptors per ommatidium, though organised differently to their fly counterparts. Dragonfly photoreceptors can be subdivided into two types: typical and vestigial. Vestigial photoreceptors are characterised by their trivial contribution to rhabdomere formation (Laughlin & McGinness 1978). They layout of these photoreceptors depends on eye region (see Laughlin & McGinness for a detailed explanation).

Hemicordulia photoreceptors can also be subcategorised based on their spectral sensitivities (Laughlin 1976). The five frequency peaks are as follows (Yang & Osorio 1991):

- UV: 330nm
- Blue: 410nm
- Blue: 460nm
- Green: 525nm
- Orange: 630nm

These spectral sensitivities are not dissimilar to other odonata (Autrum & Kolb 1968, Chappell and Devoe 1975) indicating that these properties may be common amongst most dragonflies.

Interestingly, the UV channels of *Hemicordulia tau* also exhibit polarisation sensitivity. This is in line with other insects that can detect polarisation and may play a part in migration seen in other dragonfly species. *Hemicordulia* also exhibit linked photoreceptors, which respond to an aggregate of frequencies indicating inputs from all three rhodopsin types (Laughlin 1976).

2.3.5.5 Retina Cells in Flies

In *Drosophila*, there are eight light-sensitive photoreceptors (R1-R8) which divide into two categories (Arnett 1971, Franceschini et al., 1981) exhibiting five separate spectral classes (Hardie 1986). Of the eight, only photoreceptors R1-R6 are responsible for motion detection. R7 and R8 are instead generally used for spectral wavelength discrimination (Heisenberg & Buchner 1977) though may also play a role in orienting reflexes and perhaps navigation. There are exceptions, as in the male *Musca* acute zone, R7 neurons are repurposed for motion detection (Franceschini et al., 1981).

R1-R6

R1-R6 form the start of the major motion pathways in flies being involved in the optomotor response (Heisenberg & Buchner 1977). They exhibit two absorption peaks (green & ultraviolet; Hardie 1986). Photoreceptors R1-R6 have many gap junctions connecting themselves (indicating non-independent behaviour) and synapse onto numerous neuron classes in the Lamina (L1-L4, Meinerhertzhagen & O'Neil 1991; Clark et al 2011). They also receive feedback from Amacrine cells in the Lamina.

R7-R8

R7-R8 photoreceptors are chromatic photoreceptors. In some fly species they have been shown to have several different opsins covering UV, Blue and Green wavelengths (Hardie 1986). However, they are not involved in motion pathways (Heisenberg & Buchner 1977). These photoreceptors also directly synapse into the medulla, bypassing the lamina (Hardie 1986; Stausfeld & Lee 1991).

2.3.6 Lamina

The lamina represents the first neuropil within the optic lobe and receives its inputs directly from the retina. Like the retina, it remains largely retinotopic, separated into neuro-ommatidia or cartridges (Strausfeld 1971; Rister et al., 2007). It is at the level of the lamina that first efferent neurons appear indicating that feedback is an inherent part of the visual pathway.

The lamina contains many different cells and their processes. These can be roughly cut into three groups, the Lamina Monopolar Cells (LMC) L1-L3, L5 which have terminating processes in the medulla, L4 and amacrine cells (with the associated T1) which cross between adjacent cartridges and C2-C3 which are centrifugal neurons originating in the medulla. Various genetic knockouts in Drosophila indicate that the tuning properties of the optomotor response remain unaffected by removal of various pathways (L1, L2) and that L1 & L2 are not direction selective (Reiff et al., 2010) indicating that while the Lamina may provide inputs to motion detection it does not perform the motion calculation (Rister et al., 2007).

The Lamina cells also exhibit inhibitory surround (Zettler & Järvilehto 1971; Laughlin & Osorio 1989). This consists of high-gain synapse at the Photoreceptor-LMC junction combined with inhibitory inputs from adjacent ommatidia.

From the Lamina onwards the synaptic organisation of neurons becomes increasingly complex. Unlike the idealised models found in many textbooks, optic synapses are anything but simple. It is more common than not, that neurons will form multi-synaptic groups (referred to as triads, tetrads etc.) making detailed analysis of these neurons increasingly complex.

2.3.6.1 Lamina Cells of Dragonflies

Lamina cells of the dragonfly are dominated by Large Monopolar Cells (LMCs). LMCs depolarize to dark stimuli while hyperpolarizing to light stimuli (Laughlin 1973). The strength of the signal is amplified, but also contains significant transient components (for example a light stimulus generates an initial hyperpolarization followed by a slower depolarization). These neurons receive their inputs from their retinotopically equivalent retinular cells (Laughlin 1973). However, it is not a one-to-one mapping. The increased noise measured in these neurons indicates the presence of multiple inputs to each LMC taken from surrounding photoreceptors, enabling the so-called ‘inhibitory surround’ found in other insects.

Inhibitory surround is a process whereby neighbouring units provide negative input to a central unit (Figure 7, left). This negative input selects against large objects, instead being maximally stimulated by objects which subtend the size of the central unit surrounded by a stimulus of opposing polarity (Figure 7, middle, right). For example, if a cell subject to inhibitory surround were maximally excited by darkness, then the optimal stimuli would consist of a dark point surrounded by a light ring (Figure 7, middle). In this case the excitation would be maximized and surround inhibition minimized. Likewise, the anti-

optimal stimuli would be a light point, surrounded by a dark ring (Figure 7, right), maximizing the surround inhibition while minimizing the excitation.

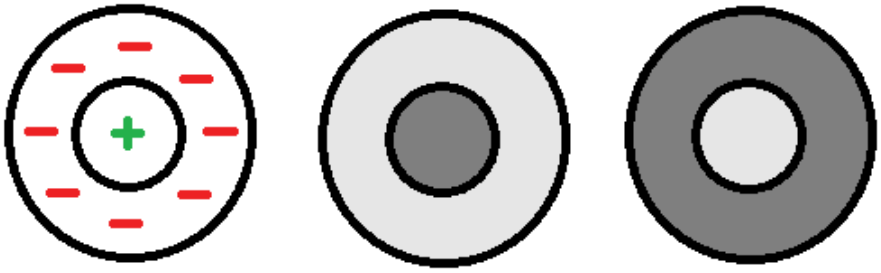

Figure 7: Example of inhibitory surround. The area in the centre is positively stimulated by a stimulus (the exact stimulus type is unimportant) while the surrounding region negatively stimulates the centre (left). This leads to maximal stimulation (if we assume dark is excitatory) to a dark point surrounded by a light ring (middle) or minimal stimulation when stimulated by a light point surrounded by a dark ring (right).

2.3.6.2 Lamina Cells of Flies

L1 & L2

In *Drosophila*, L1 & L2 are lamina output cells that receive their inputs from R1-R6 (Strausfeld 1971; Meinertzhagen & O'Neil 1991; Clark et al 2011) but also efferent inputs from C2 & C3, centrifugal neurons originating in the medulla (Meinertzhagen & O'Neil 1991). The transfer of information between photoreceptors and LMCs is dependent on the background adaptation state. In dim conditions, LMCs simply invert the photoreceptor signal, but in bright conditions encode changes in brightness (Juusola et al., 1995). Given insect phototransduction, this means that L1 & L2 hyperpolarize in response to light increments and depolarize in response to dark increments (Tuthill et al., 2013). In flies these synapses are facilitated by histaminergic chloride channels, which can be rendered defective in the so-called *ort*-mutation (Borst et al., 2010).

Both L1 & L2 play important roles within the motion pathway as evidenced by the optomotor response (Rister et al., 2007, Tuthill et al., 2013). Degeneration of L1 & L2 has been shown to suppress the optomotor response while preserving fixation (Coombe & Heisenberg 1986). Blocking activity within both neurons completely suppresses the optomotor response (Rister et al., 2007, Tuthill et al., 2013). Blocking L1 & L2 individually also introduces a change in the effective contrast sensitivity to standard motion signals (gratings) (Rister et al., 2007). However, the picture becomes more complicated when the density (the total number of

features in a scene) is modulated. Dense scenes produce standard responses, whereas sparse (very few objects) induce a counter-motion behaviour. When L2 is suppressed translational movements of flies were shown to be entirely suppressed, while still allowing some rotational responses to remain (Katsov & Clandinin 2008) indicating an early separation of different aspects of motion.

L1 & L2 have also been implicated in the optomotor response when the components of the response are decoupled, for example direction and polarity. Silencing L1 suppressed ON edge responses while silencing L2 suppressed OFF edge responses (Joesch & Schnell et al., 2010). This can be explained by L2's role as a brightness decrement detector (Reiff et al., 2010). L2 can detect brightness decrements with all other Lamina cells silenced, however its kinetics become slower indicating that the remaining cells modulate L2's kinetics (Reiff et al., 2010). L1 also appears to detect brightness decrements (Clark et al., 2011). Interestingly it is in the reverse-phi domain (see below for a detailed explanation of reverse-phi motion) that the real differences between L1 & L2 emerge with L1-only flies insensitive to two-point bright-dark and L2-only flies insensitive to two-point dark-bright which induce reverse-phi behaviour. L1 also receives indirect input from L1 via L5, though L5 is not implicated in motion detection (Takemura et al., 2013).

L2 also synapses onto L4. This circuit is thought to be involved in responses to progressive motion with silencing of either neuron resulting in optomotor impairments (Tuthill et al., 2013). Efferent processes from L2 (as well as L4 and Amacrine cells) also synapse back onto R1-R6 creating a negative feedback loop that may increase information flow (Zheng et al., 2006). On the forward path, L1 and L2 synapse onto multiple layers of the medulla (M1, M5 and M2 respectively) where medullary neurons pass motion signals to deeper layers of the optic lobe.

L3

L3 is another LMC found in the lamina that receives inputs from all six of the preceding photoreceptors (Strausfeld & Campos-Ortega 1973) but does not receive inputs from the surrounding L4 neurons indicating it may be insulated from surround inhibition effects. L3 does have a downstream role in vision synapsing onto M3 (Rister et al., 2007), Mi9 (Takemura et al., 2017) and Tm9 (Borst & Helmstaedter 2011) and is also implicated in optomotor turning behaviour (Silies et al., 2013). Moreover, as motion behaviour is abolished via removal of L1 & L2, but orientating behaviour remains (fixation on large objects) it is

hypothesized that this reflex occurs through L3 (though not exclusively) (Rister et al., 2007). Silencing L3 & L1 can remove dark edge responses (Silies 2013) indicating that L2 requires input from L1 or L3 to detect dark edges.

Unlike L1 & L2, L3 is also associated with connections from R7-R8 (Strausfeld & Lee 1991), the photoreceptors which provide chromatic information.

L4

Unlike the other LMC's L4 does not have any distal (with respect to the retina) neurites forming all its synapses within the lamina layer itself in both the x and y direction (Meinerhertzhagen & O'Neil 1991). It synapsed both presynaptically and postsynaptically with L2 in some adjacent cartridges but receives input from all six neighbouring L4 neurons (the retina has a hexagonal pattern).

L4 connects to M2 and M4, shows no direction sensitivity to bars and can function independently of L2 silencing (Silies et al., 2013) though it can respond to dark moving bars.

Despite its small reach, L4 is integral to motion detection (Tuthill et al., 2013). Drosophila with L4 selectively blocked were reported to be motion blind (Reiff et al., 2010).

L5

Unlike other LMCs, L5 has been shown to have no role in motion vision (Takemura et al., 2013) and has no significant synaptic engagement in the lamina (Meinerhertzhagen & O'Neil 1991).

C2 & C3

C2 & C3 are centrifugal neurons (Douglass & Strausfeld 1995), which project from the medulla into the lamina (Meinerhertzhagen & O'Neil 1991). As per Figure 4, C3 is connected to motion detection in T4 cells (see below) and thus motion detection. C2 also is involved in the motion detection pathway (Douglass & Strausfeld 1995) synapsing on both L1 & L2 (Meinertzhagen & O'Neil 1991).

Perhaps most interesting, is that their contribution appears to be most significant when considering reverse-phi motion (an optical illusion where a contrast inversion falsely signals a reverse motion stimulus). Both C2 & C3 appear to have complementary roles when interpreting reverse-phi motion, causing an inhibition and enhancement of the reverse-phi optomotor inversion respectively (Tuthill et al., 2013).

Lawf1 & Lawf2

The first widefield cells (cells which respond to aggregate motion over a large region) in the motion pathway are Lawf (Lamina Widefield) cells, which receive their inputs in the medulla and synapse back in the Lamina (Tuthill et al., 2014). These cells respond robustly to flicker of low frequency (~1Hz) but are not direction selective. The responses of Lawf2 cells are modulated by behaviour (flight) (Tuthill et al., 2014).

Amacrine Cells

Amacrine cells connect adjacent lamina cartridges. They receive their inputs from R1-R6 and synapse onto T1, which connects amacrine cells to the medulla (Rister et al., 2007). It has been proposed (Strausfeld & Campos-Ortega 1977) that amacrine cells are involved in lateral inhibition and photoreceptor adaptation. Unlike other LMCs, Amacrine cells are sign preserving (Douglass & Strausfeld 2005). Amacrine cells are motion-sensitive and are possibly orientation selective (Douglass & Strausfeld 2005).

T1 (Medullary Cell)

Cell T1 is postsynaptic to amacrine cells (Rister et al., 2007) but like L5, there is no evidence connecting T1 to motion of any kind (Tuthill et al., 2013).

2.3.7 Medulla

The medulla is presumed to contain the cellular machinery for motion correlation (Tuthill et al., 2013). In *Drosophila* it is divided into ten separate layers (Fischbach & Dittrich 1989) each defined by the termination of various lamina neurons. What follows will largely consider neurons found in the fly visual pathways. While no equivalent experiments have been performed in dragonflies, it is likely that many of the cells and mechanisms found in flies will have analogues in dragonflies. However, it should also be noted that the dragonfly medulla is significantly more complex than the fly having many more layers of neurons (Fabian et al., in submission).

2.3.7.1 Motion Detection

While the retina and lamina largely influence the temporal properties of light, it is in the medulla that spatial correlations (and hence motion detection) really originates (Strother et al., 2014). The fly medulla contains a vast number of neurons, many of which have been implicated in motion pathways. These simple motion pathways are thought to be based upon

correlation-based motion detection (Figure 8). These detectors take two spatially separated inputs and combine the output after one of the inputs has been delayed. In effect, one input is compared to the delayed version of the second input. If the detector is subject to motion in the preferred direction of the detector these two signals become correlated spatiotemporally, in effect a form of autocorrelation. These units, called Elementary Movement Detectors (EMDs) are thought to reside in the medulla.

Numerous neurons within the medulla have been implicated in motion detection to greater or lesser extents (Behnia et al., 2014, Stother et al., 2014). While direct evidence for their role in motion detection is difficult to show, numerous studies have found neurons with temporal and anatomical properties that make them sensible candidates.

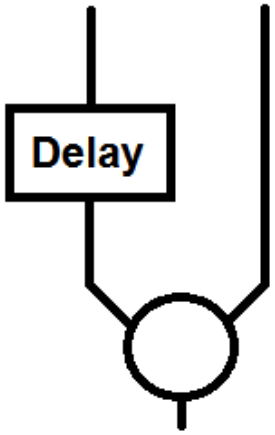

Figure 8: Example of a correlation detector. Two spatially separated units (one temporally delayed) are combined (this combination must be non-linear). The resulting output predicts the correlation between the two inputs over time, an analogue of motion.

2.3.7.2 Transmedullary Neurons (Tm)

Transmedullary neurons exhibit a wide variety of morphology but are essentially the same kind of neuron structurally (Strausfeld & Lee 1991). These neurons have their cell body distal to the medulla and connect the distal medulla with the lobula with a naming convention which prioritizes the layer at which transmedullary neurons arborise (Fischbach & Dittrich 1989). Of the transmedullary neurons, Tm1, Tm2, Tm3, Tm4 & Tm9 have been implicated in motion pathways (Behnia et al., 2004; Strother et al., 2014; Arenz et al., 2017). As each of these neurons functionally act to filter out high frequency or low frequency information (or both), it is postulated that they may form arms of EMDs in *Drosophila*.

2.3.7.3 Medulla Intrinsic Neurons (Mi)

Medulla Instrinsic neurons (Mi) differ morphologically from Tm neurons but appear to fulfil similar roles in motion detection. Mi1, Mi4 and Mi9 are all implicated in motion detection (Behnia et al., 2004; Strother et al., 2014; Arenz et al., 2017) forming connections between LMC and T4/T5 neurons (Takemura & Meinertzhagen 2008). As each of these neurons functionally act as either bandpass or lowpass filters, it is postulated that they may be present on the inputs of any two-point correlation-based motion detection used in the ON and OFF motion pathways in *Drosophila.*

2.3.8 Lobula Complex

The lobula complex is divided into many subregions of specialisation, showing significant divergence between species (Fabian et al., in submission). Within the lobula complex, one region is particularly well studied, the lobula plate. It is within this subregion that many of the widefield-motion sensitive neurons of flies and other insects are found, however this is not universally the case such as in bees where the equivalent neurons exist in a separate neuropil (Devoe et al., 1982). This is also where target-tracking neurons are found. If one were to summarize the lobula (at least in insect vision), it would appear to house many small and large receptive-field neurons that integrate motion information over multiple ommatidia of the insect. It is these higher-order neurons that are the main study of this thesis.

2.3.8.1 Separation of ON and OFF channels

It has been shown in *Drosophila*, mice and humans (Borst & Helmstaedter 2015) that the Elementary Motion Detectors (EMDs) separate ON (light increments) and OFF (light decrements) into separate pathways (Riehle & Franceschini, 1984). This allows for the detection of both light and dark edges within a visual scene (Fisher et al., 2015a). Both ON and OFF edges contribute to behavioural responses to motion (Fisher et al., 2015a). This also helps explain the reverse-phi optomotor response (moving against the signal motion) without the need for a sign-correct multiplier (Clark et al., 2011).

The separation of ON and OFF channels occurs early in some species (as early as the Lamina – Reiff et al., 2010) indicating that contrast polarity is an important input to motion vision). One significant benefit of early separation is the more efficient encoding of information (Borst & Euler 2011, Gjorgjieva et al., 2014). More importantly, the statistics of natural images is a driving factor for this separation. ON and OFF channels tuned to different frequency optima (as seen in *Drosophila*) results in better encoding of velocity in natural scenes (Leonhardt et al., 2016).

Despite the existence of ON-ON and OFF-OFF correlators, their counterparts (ON-OFF & OFF-ON) are unlikely to exist in the wide-field motion pathway (Eichner et al., 2011).

2.3.9 Wide-field Motion Sensitive Neuron Properties

Wide-field Motion Sensitive (WFMS) neurons have been found in most invertebrate and vertebrate species that have vision. These cells respond to motion occurring over a large region of an animal's field of view usually due to translation or rotation. Rotational movement is simpler to analyse as the effective image on the retina changes very little with rotation (only its relative location) whereas translational movement introduces motion at many different speeds related to the distance of features in an image. These motion differences may even play a role in encoding such as respond strongly to near contrast edges (Schwegmann et al., 2014)

WFMS neurons have been found in many invertebrates, including beetles cockroaches (Kathman et al 2014), flies (Hausen 1982a, b), bees, butterflies, moths (Theobald et al 2010) as well as vertebrates including cats (Rauschecker et al., 1987) and monkeys (Dubner & Zeki 1971, Tanaka & Saito 1989, Duffy & Wurtz 1991, Duffy 1998) and pigeons (Wang & Frost 1992) but remain as yet unexamined in dragonflies.

2.3.9.1 Sinusoidal Tuning

One of the classic methods for examining the responses of WFMS cells is to stimulate them with sinusoidal gratings of varying spatial (sinusoidal narrowness/broadness) and temporal frequencies (number of light-dark-light cycles at a single point measured in Hz). In each of these two domains, the neurons in question will respond maximally to specific spatial and temporal frequencies and attenuating their response on either side of this maximum (Figure 9). This is explicable for each separate axis (space and time) due to entirely different effects.

Space: When considering spatial frequency optima, there are two competing factors generating tuning. Firstly, every eye is subject to a limited resolution and beyond a certain spatial frequency (for example 10^6 cycles/°) there is no way to tell black from white and the two blur together into grey. On the low spatial frequency side, it is more complex. This is thought to be a result of inhibitory surround generated by LMCs in the lamina. This provides a fundamental limit on contrast sensitivity. In addition to inhibitory surround, the contrast resolution (i.e. which contrasts can be easily distinguished) also plays a role. In large wavelength sinusoids, adjacent samples do not differ from one another by much requiring a mental comparison of further and further distant points.

Time: The preferred temporal tuning found in WFMS cells is due in part to two predominant factors. Firstly, the necessary dependence on correlation thought to underlie motion detection (i.e. EMDs, Hassenstein & Reichardt 1956) and the effects of temporal adaptation of the preceding neurons. In effect, if the preceding neurons are too sluggish to encode changes in input light, it is impossible for an EMD to sensibly correlate signals. Likewise, if too much adaptation occurs, slow-signals (i.e. low temporal frequencies) will also be eliminated. Thus, the spatiotemporal tuning curves observed are a combination of the underlying time-constants thought to enable EMDs and the temporal kinetics of the preceding neurons (such as photoreceptors and LMCs). In addition to these, the WFMS neurons themselves are subject to their own adaptation constraints (see Motion Adaptation below for details).

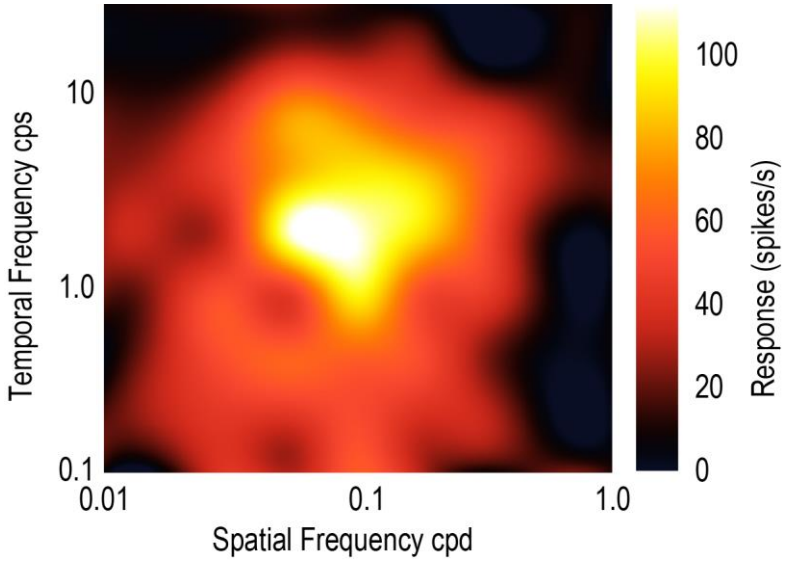

Figure 9: Example of a spatiotemporal tuning curve. The neuron is maximally stimulated by a grating with a spatial frequency of 0.1 cycles/° (i.e. one sinusoid per 10 degrees) and a temporal frequency ≈ 2Hz. At both higher spatial and temporal frequency, the response drops away.

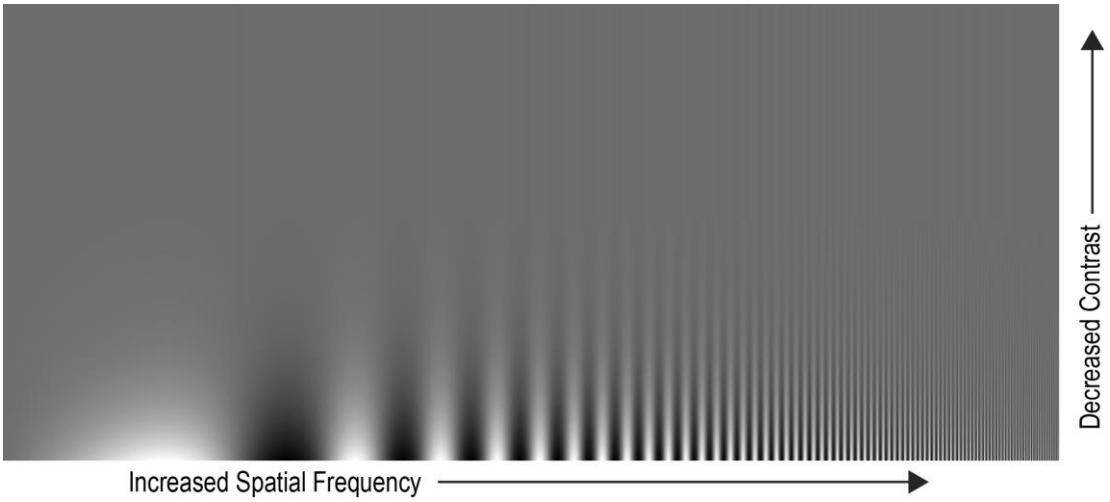

Figure 10: *Example of the effects of optical blurring and inhibitory surround resulting in visual tuning for spatial frequency. Contrast decreases higher in the image while frequency increases to the right.*

2.3.9.2 Motion Adaptation

Neurons sensitive to widefield motion often exhibit a form of motion-adaptation (Maddess & Laughlin 1985; Maddess 1986). This is an effect similar to photoreceptor adaptation to prolonged stimulation. When exposed to a motion signal for an extended period, the response will tend to start with a strong transient component before decaying to a lower steady-state (which is not necessarily zero). This motion-adaptation effect is local (Maddess & Laughlin 1985) and highly dependent on the temporal frequency of the motion stimulus. It is poorly recruited by flicker stimulus (Maddess & Laughlin 1985, presumably it only occurs because flicker stimuli recruit weak motion signals). In H1 & HS cells in flies, motion adaptation has been demonstrated through stronger coherence to pseudorandom velocity patterns rather than constant velocity patterns (Haag & Borst 1997). As motion-adaptation is differentially recruited depending on temporal frequency, it can affect the measurement of temporal frequency optima, causing a leftward shift as higher frequencies are suppressed more strongly than lower frequencies (Maddess & Laughlin 1985).

It was hypothesized that motion adaptation might be accounted for by a change in the temporal delay of EMDs (Clifford & Langley 1996), however this was contradicted by change of experimental paradigm (Harris et al., 1999). In these experiments, the delay optimum was determined using single-frame sinusoidal gratings separated by a delay interval. This paradigm removed the effects of adapting to a stationary grating and thus inducing the apparent changes in temporal delay. Additionally, some elements of motion adaptation are

spatially localized (O'Carroll et al., 2012) but are most strongly recruited by elongated orthogonal features.

Disentangling the components of motion adaptation is complex. For example, given that motion-sensitive units (such as EMDs) draw their inputs from photoreceptors (which adapt), some of that adaptation may pass through the system to WFMS neurons.

Ignoring the precise causes of motion adaptation, the effect itself can be measured in several ways. A classic approach is the test-adapt-test methodology, which involves a test, followed by an adaptor and a repeat of the original test to measure the effect of the adaptor. In flies, such experiments demonstrated that multiple forms of motion-adaptation exist (Harris et al., 2000). A brief overview is shown in Figure 11. The forms of adaptation are:

- Contrast gain reduction (a rightward shift, reducing overall sensitivity)
- Output range reduction (a vertical compression limiting the maximum response)
- Vertical shifts (representing changes in the underlying membrane potential)

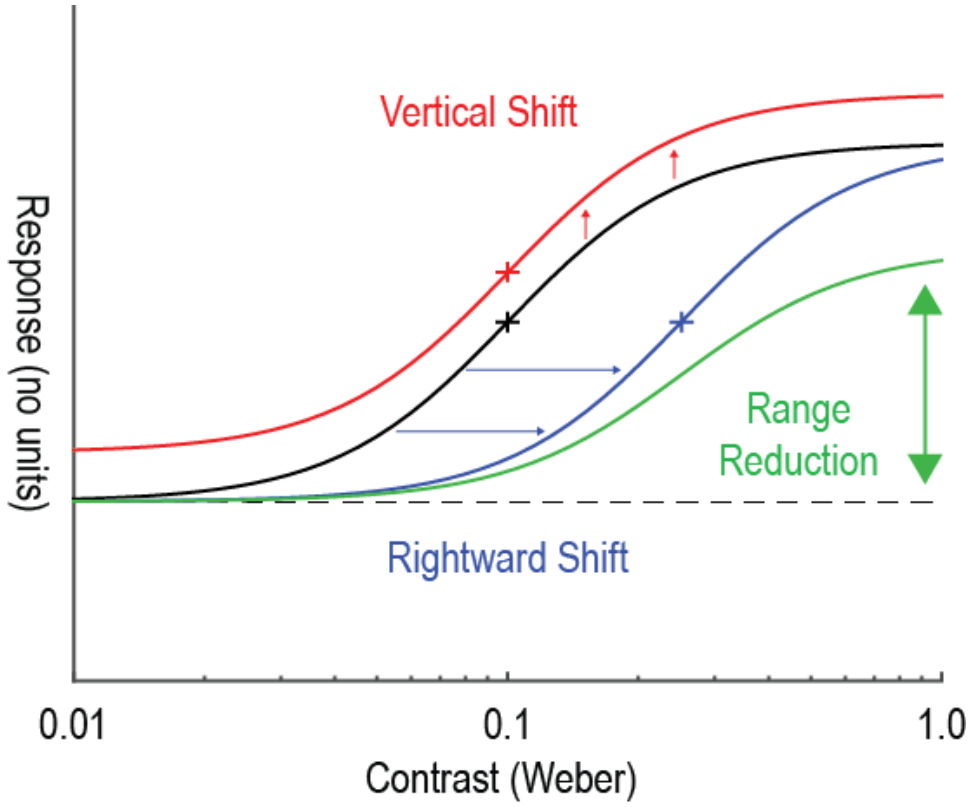

Figure 11: Example of motion adaptation and its components. The contrast curves here represent the pre and post-adapt responses of a WFMS neuron subjected to a test-adapt-test protocol. The black represents the pre-adaptation test and the coloured curves the theoretical post-adaptation responses. Vertical shifts (up or down, red) can occur as the result of changes in the underlying membrane potential changing due to prolonged excitation or inhibition. These changes oppose the adapting response. Rightward shifts (blue, also known

as contrast gain reduction) occur as the neuron becomes effectively less sensitive. Output range reduction (green) represent a kind of range squashing, limiting the maximum response a cell can generate (even with a powerful stimulus).

Motion adaptation has many useful functions, including improving velocity contrast (Maddess & Laughlin 1985) via relief from saturation (Barnett et al., 2010) on a timescale similar to the response (Nordström et al., 2011). It also improves velocity encoding of natural images (Shoemaker et al., 2005, Straw et al., 2008, Barnett et al., 2010) and enhance differentiation between foreground and background features (Li et al., 2017).

2.3.10 Wide-field Motion Sensitive Neurons

2.3.10.1 T4 & T5 (Flies)

T4 & T5 are lobula-based cells which taken their inputs from the medulla and output in the lobula plate (Borst & Helmstaedter 2015). T4 & T5 represent the first direction selective neurons in the wide-field motion pathway of the fly. They are tuned to fairly low temporal frequencies in flies (1Hz) though this can be altered by the application of CDM (an octopaminergic agonist, Arenz et al., 2017).

Both classes are direction selective, being both excitatory in their preferred direction and inhibitory to null-direction motion (Haag et al., 2017). Interestingly, this direction selectivity is distributed over their receptive fields with the inhibitory input on one side and the excitatory on the other (Haag et al., 2017). T4 & T5 can be further subdivided into four subtypes based on their direction selectivity, corresponding to the cardinal directions; up, down, left and right (Buchner 1984, Strausfeld & Lee 1991). Both T4 & T5 play a critical role in motion detection. When both are suppressed, all LPTC activity is abolished (Schnell et al., 2012, Leonhardt et al., 2016). This extends to behaviour where optomotor responses are also abolished, though fixation behaviour remains (Bahl et al., 2013).

T4 and T5 neurons differ from one another both anatomically and physiologically (although they are quite analogous in their function). T4 are ON-edge selective (Maisak et al., 2013) and receive inputs from L1 via Mi1 (Takemura & Meinertzhagen 2008) and Tm3 (Takemura et al., 2013). They also provide inhibitory feedback via LPi (Lobula Plate interneurons) onto equivalent cells of the reverse motion resulting in noise reduction in their motion calculation (Arenz et al., 2017).

T5 neurons instead detect OFF edges (Maisak et al., 2013) and receive their inputs from L2 via Tm1, Tm2, Tm4 & Tm9 (Arenz et al., 2017).

2.3.10.2 LPTCs aka DSMD (Flies)

Flies possess ~60 individually identifiable lobula-plate neurons exhibiting motion sensitivity (Spalthoff et al., 2010). Found in numerous species (including *Bombylius, Calliphora, Drosophila* & *Eristalis*) Lobula Plate Tangential Cells (LPTCs) are unified by their directional encoding of widefield motion. LPTCs are temporally tuned, usually peaking ~1Hz in fly species (Borst et al., 2010, Schnell et al., 2010) though some exhibit double peaking (O'Carroll et al., 1996). LPTCs integrate motion from numerous smaller motion-sensitive neurons (T4 & T5) with their dendritic inputs responsive to the locally relevant direction for the global function (Spalthoff et al., 2010).

The temporal tuning of LPTCs has been shown to be amenable to pharmacological intervention (Longden & Krapp 2009). Flight behaviour and octopamine agonists (including CDM) bind to receptors causing a change in the temporal tuning properties of LPTCs (Suver et al., 2012) mimicking behavioural states and raising the temporal optima. This has also been shown to occur in walking flies (Chiappe et al., 2010) though this is dependent on the nutritional state of the animal (Longden et al. 2014). There are numerous subtypes of LPTCs, with many having been well characterised as detailed below.

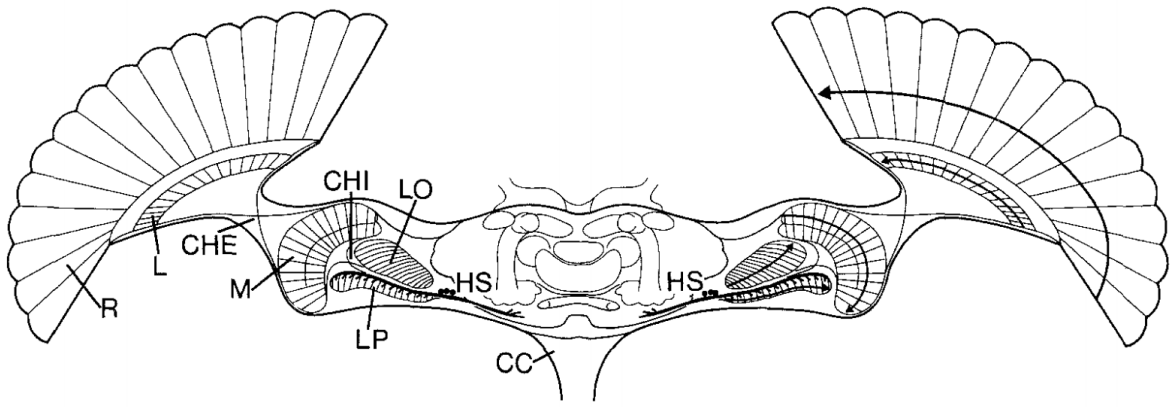

Figure 12: Reproduced from Hausen 1982a. Example schematic of the fly brain including brain regions and an example HS neuron (see below). Labelling is as follows (R – Retina, L – Lamina, CHE – External chiasm, M – Medulla, CHI – Internal chiasm, LO – Lobula, LP – Lobula plate, HS – Horizontal System Neuron, CC – Cervical connective)

Horizontal System (HS)

Found in the lobula plate (Figure 12), HS neurons encode horizontal motion (i.e. yaw) in *Calliphora* (Hausen 1982a, b) and *Drosophila* (Schnell et al., 2010). There are several types of HS neurons whose nomenclature is derived from their relative locations and orientation of

their dendrites (HSN – North/Dorsal, HSE – Equatorial, HSS – South/Ventral, Hausen 1982b) with each neuron covering ~40% of the lobula plate.

HS neurons encode progressive (front-to-back motion) and receive inputs from contralateral H1 neurons, improving the distinction between rotational and translational movement (Horstman et al., 2000). HS neurons are ipsilateral (Hausen 1982a) with their inputs terminating in the first layer of the Lobula (Fischbach & Dittrich 1989).

Vertical System (VS)

Like their HS counterparts, VS neurons encode wide-field motion but prefer motion oriented vertically (Krapp & Hengstenberg, 1996; Krapp et al. 1998). They have slightly different tuning (2Hz) and respond to both ON and OFF Flicker (Hengstenberg 1982).

H1 & H2 – Flies

H1 & H2, though sensitive to horizontal motion like HS neurons are distinguished by their interhemispheric arborizations, signal structure and receptive field shape (Maddess & Laughlin 1985; Krapp et al. 2001). H1 has been shown to detect regressive motion (Maddess & Laughlin 1985), which coupled with HS neurons helps distinguish progressive and rotational movements.

Like H1, H2 neurons provide post-synaptic output to HS neurons (Horstmann et al., 2000). H2 responses to binocular stimuli are not simple linear combinations. However, H2 neurons do not respond to contralateral stimulation in absence of ipsilateral stimulation (Haag & Borst, 2001).

Centrifugal Horizontal (CH) Cells

The Centrifugal Horizontal (CH) cells are a set of efferent Lobula Plate neurons found in blowflies, sensitive to rotational yaws (Eckert & Dvorak, 1983; Krapp et al. 2001). These neurons have been shown to be both post-synaptically connected to H1 and H2 (Gauck et al., 1998) cells (see above) and provide inhibitory inputs to FD neurons (see below), enabling their size selectivity (Warzecha et al., 1993).

Anatomically, CH cells have their cell body in the contralateral midbrain (compared to recording site) and have two main arborisations, one in the lobula plate (where they inhibit FD) and one in the ipsilateral ventrolateral brain (where they are post-synaptic to H1 and H2). The lobula plate arborisation is involved in both ipsilateral and contralateral stimulation, whereas the ventrolateral brain is involved in only contralateral stimulation (Egelhaaf et al., 1993). They have two main distributions (Eckert & Dvorak, 1983) and hence subtypes: DCH

(dorsal) and VCH (ventral). It is the latter that is implicated in inhibition of FD1 (Warzecha et al., 1993).

2.3.10.3 WFMS Neurons in Hawkmoths

Hawkmoths, which operate in very light-poor environments exhibit hovering feeding behaviours necessitating detection of very slow velocities. Comparisons between crepuscular and nocturnal variants have indicated that the velocity tuning is maintained at the expense of introducing longer visual delays (and hence greater temporal integration) in a nocturnal species (Theobald et al., 2010).

Like flies, Hawkmoths have both a horizontal and vertical system of WFMS neurons (Wicklein & Varju 1999) which occupy similar locations in the brain. They also exhibit tuning, though higher than flies at 2Hz (O'Carroll et al., 1996).

2.3.10.4 WFMS Neurons in Bees

Unlike flies or hawkmoths, hovering is not represented in the honeybee behavioural repertoire. More important are the long-distance flights seen in activities such as pollen collection. It has been suggested (Ibbotson 2001) that the honeybee has two separate visual systems designed to cope with this. The first, mimic the behaviour of LPTC cells in flies (Ibbotson & Goodman 1990), whereas the second are velocity-tuned neurons, geared to encode motion on long flights.

Descending Neurons (Bees)

Direction-selective WFMS neurons are not restricted to optic lobes. In honeybees, a series of descending neurons (DN) have been found to be responsive to wide-field motion exhibiting both direction selectivity and classical tuning (8-11 Hz, Ibbotson & Goodman 1990).

In addition to the more LPTC-like DN neurons, another class of DN neurons, dubbed Velocity-Tuned (VT) exhibit direction selectivity but are tuned to velocity, rather than temporal frequency (Ibbotson 2001).

Lobula Widefield Cells (Bees)

Bees also possess neurons within the lobula (dubbed lobula widefield cells – LWCs), which are sensitive to motion (Paulk et al., 2008, Mertes et al., 2014). These neurons exist in multiple layers, with some being sensitive to motion, some sensitive to colour cues and some responsive to both cues. The motion sensitive neurons are direction selective (Paulk et al., 2008). These neurons also encode information about landmarks from the environment (Mertes et al., 2014).

One of the more curious aspects of these neurons is their presence in a different brain region to those of LPTCs found in flies. The bee does not have a lobula plate but does possess a neuropil which contains neurons performing similar functions to the lobula plate, named the sublobula.

2.3.11 Feature Detection

Feature detection is a broad term encompassing almost all visual stimuli (they are, after all features). None-the-less in vision, this term tends to have a narrower definition, referring to small objects within a scene (rather than the scene itself). Here we run into a curious issue of natural languages where terms like ‘small’, ‘large’, ‘fast’ and ‘slow’ which are effectively relative terms begin to simply confuse things. This section is roughly divided into two subsections, target detection and feature detection. Features in this context refer to large objects whereas targets here refer to the small (< 5 degree) objects which are likely objects relating to other insects.

2.3.11.1 Figure Detecting Neurons of the Fly (FD)

In concert with the well described LPTC cells, flies also possess figure-detecting (FD) neurons (Egelhaaf 1985a, b). Found in the Lobula Plate (Egelhaaf 1985b), these neurons are sensitive to progressive or regressive sinusoidal input subtending a subsection of the field of view (i.e the receptive field of the neuron). The features detecting by FD neurons are still quite large, their name originating in contrast to LPTC (whose optimal stimuli, large gratings, inhibit FD neurons).

FD1 & FD2

FD1 has an ipsilateral receptive field tuned for small figures (12°) moving progressively (Egelhaaf 1985b). It is direction opponent (inhibited by regressive motion). Its receptive field is large, subtending at least 40° laterally and extending vertically in the ventral part of the fly’s visual field (Warzecha et al., 1993). It is also inhibited by ‘widefield’ motion (i.e. very large background motion), even when the motion is only restricted to the ipsilateral side (i.e. same side) and regardless of its direction (Egelhaaf 1985b).

Anatomically, FD1 takes its inputs from the frontal/caudal lobula plate, covering the entire dorsal-ventral extent (Egelhaaf 1985b), though the dorsal inputs appear thinner and thus generate greater attenuation (Warzecha et al., 1993). There are two variants, distinguished by contralateral axonal extent (though with similar physiological properties) and cell body location. Both produces outputs in the midbrain.

The size selectivity observed in FD1 is caused by inhibition by surrounding GABA-ergic VCH (Ventral CH) neurons. This was demonstrated via injection of picrotoxin (effectively silencing the inhibition) resulting in FD1 losing both its size selectivity and direction opponency, though it remained weakly direction selective (Warzecha et al., 1993).

FD2 is similar to FD1 except that it prefers regressive motion. It has broadly similar receptive field shape and anatomical structure.

FD3 & FD4

Unlike FD1 & FD2, FD3 has a more lateral receptive field, centred ~50° from the mid-axis. It is a regressive-motion sensitive neuron exhibiting direction-opponency. Interestingly, the widefield inhibition in this neuron is maximized by ipsilateral-only stimulation, a surprise given that contralateral stimulation also inhibits it (Egelhaaf 1985b). FD4 likewise has a lateral receptive field, but much more diffuse, covering the entire ipsilateral extent.

2.3.11.2 Loom-Sensitive Neurons

Perhaps one of the simplest feature detections which has immediate behavioural relevance is the detection of looming objects. Many insects exhibit aversive behaviour to such stimuli (Robertson & Johnson 1993). Conceptually, such neurons are somewhat ambiguous in nature, as distinguishing looming stimuli from progressive or other translator motion is non-trivial.

Loom-Sensitive Neurons in Hawkmoths

Hawkmoths have been shown to have at least two classes of cells that respond selectively to looming and receding stimuli (Wicklein & Strausfeld 2000). One class responds selectively to expanding discs while the other to more progressive flow-field stimuli. Interestingly such cells are contrast-polarity independent, responding to both light and dark expanding disks.

Loom-Sensitive Neurons in Locusts

Locusts have also been the subject of looming neuron research, with two neurons, the DCMD (Descending Contralateral Movement Detector) and LGMD (Lobula Giant Movement Detector) both being sensitive to the motion of small (< 30°) objects, specifically those that are moving towards the subject (Rind & Simmons 1992; Judge & Rind 1997; Fotowat & Gabbiani 2011).

DCMDs are monocular neurons that are most sensitive to approaching objects and are inhibited by receding objects. Interestingly, the DCMD also interacts with widefield motion which diminishes the loom response (Rind & Simmons 1992).

LGMDs are the presynaptic neuron to DCMDs taking their inputs from different regions of the lobula and synapsing almost exclusively onto the DCMD (O'Shea & Williams 1974).

2.3.11.3 Male Lobula Giants (MLG)

Flesh flies have been shown to possess small-target tracking neurons. These neurons dubbed Male Lobula Giants (MLG) are as the name suggests, specific to males, indicating their likely role in conspecific or mate pursuit (Gilbert & Strausfeld 1991, Trischler et al., 2007). This proposal is evidenced by their projection to the pre-motor areas of flight control and receptive fields subtending the acute zone of the visual field. Most interestingly, MLGs are direction selective and tuned for small targets. There is also some evidence that MLGs reuse the same EMDs as their wide-field counterparts with MLG inputs arising from identical layers to transmedullary (Tm) neurons implicated in the wide-field motion system (Gilbert & Strausfeld 1991).

2.3.11.4 Target-Selective Descending Neurons of Odonata

Between the brain and the thoracic ganglion, dragonflies and damselflies possess a set of eight descending neurons two of which are target selective (Olberg 1981, Frye & Olberg 1995). These neurons exhibit a central receptive field, are insensitive to widefield motion and are selective for small objects subtending 4-16° (Frye & Olberg 1995).

2.3.11.5 Small Target Motion Detectors of the Dragonfly

STMDs in dragonflies represent a diverse class of neurons that respond selectively to small (< 10 degrees) moving targets (O'Carroll 1993).

CSTMD1

CSTMD1 is perhaps the most studied of all STMD neurons and is one of the main neurons studied in this thesis. It is a centrifugal neuron that takes its inputs from the contralateral side (relative to the traditional axonal recording site Figure 14) and synapses into the ipsilateral distal lobula (Geurten et al., 2007). The cell body is found in the contralateral mid-brain and interestingly, it possesses two arborisations within the mid-brain one on each side. This lends the possibility that CSTMD1 communicates with its symmetric equivalent (Geurten et al 2007).

CSTMD1 is a spiking neuron (due to the recording site and size any graded potentials have well and truly dissipated) exhibiting strong bi-phasic responses up to 100mv (Geurten et al., 2007). CSTMD1 also exhibits a potent post-excitatory rebound after prolonged stimulation (as much as 10 mV).

CSTMD1 is size selective for small targets measuring less than 5 degrees across (optima ~3 degrees). This is proposed to be caused by a strong lateral inhibition (Wiederman et al., 2008). This inhibition acts on a range more distant that traditional surround inhibition thought to occur in the lamina and is responsible for the exquisite size tuning (Figure 13). CSTMD1 also exhibits a longer-range version of inhibition, which will be discussed later as part of the selective attention section.

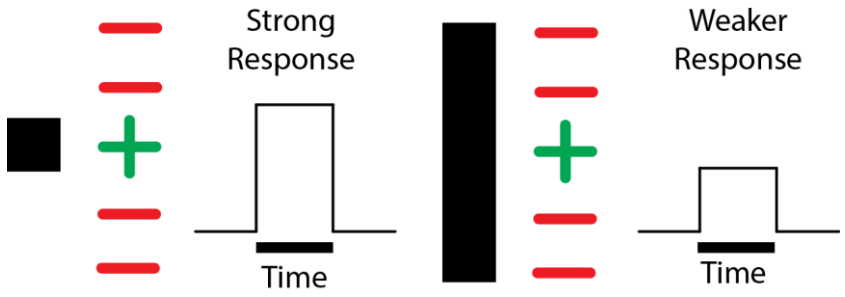

Figure 13: Example of lateral inhibition creating size selectivity. The small target (left) is not subject to any lateral inhibition (+/-) producing a strong response whereas the elongated target (right) stimulates the surrounding regions thus producing a weaker overall response.

CSTMD1 is velocity tuned, generating rising responses up to its preferred velocity of 56°/s (Geurten et al., 2007, Dunbier et al., 2012) before reducing. It is interesting how this may interact behaviourally as perching dragonflies have been shown to use a target velocity cue just prior to take-off (Olberg et al., 2000).

CSTMD1 also exhibits a unique receptive field (Figure 15), showing strong excitatory responses (with limited direction-selectivity) to small targets presented in the contralateral hemisphere. The receptive field is large (subtending > 80 degrees laterally) and when stimulated using projection distortion (Evans 2016, unpublished work) exhibits a roughly equal response out to the limits of the stimulus. On the ipsilateral side, CSTMD1 also possesses an inhibitory hemifield whose direction selectivity is suspiciously similar to a mirror-symmetric CSTMD1 (Bolzon et al., 2009), lending yet more evidence to the possibility of inter-CSTMD1 communication.

The mechanism to explain CSTMD1's preference for small targets (rather than narrow edge features) has been shown to be the result of correlations of ON and OFF edges (Wiederman et al., 2013). This also reconciles with previous experiments demonstrating that the velocity tuning of CSTMD1 is highly dependent on the effective width of the target (assuming a laterally moving target). In effect, a larger target, moving more quickly will cause the same temporal profile for underlying EMDs as a small target moving more slowly. This work also showed that the correlation is one-way, only responding to OFF->ON sequences as expected

of dark targets. This apparent correlation effect is completely independent of background motion. Targets embedded in a moving background still showed robust responses despite a lack of relative motion (Wiederman & O'Carroll 2011).

CSTMD1 exhibits potent contrast sensitivity (Geurten et al., 2007), matching the behaviour of correlator EMDs thought to underlie its behaviour (Hassenstein & Reichardt 1965, Wiederman et al., 2008). Interestingly, the contrast sensitivity observed is extremely good (50-67 – O'Carroll & Wiederman 2014) and is not dissimilar to that seen in earlier processing stages such as photoreceptors indicating that even though they only sample small regions of the receptive field at a time, the complex spatial and temporal integration improves effective contrast sensitivity.

Finally, CSTMD1 also possesses numerous higher-order properties including facilitation (Wiederman & Fabian et al., 2017) and selective attention (Wiederman & O'Carroll 2013). Due to the complexity of these processes, they will be discussed separately under *Physiology*.

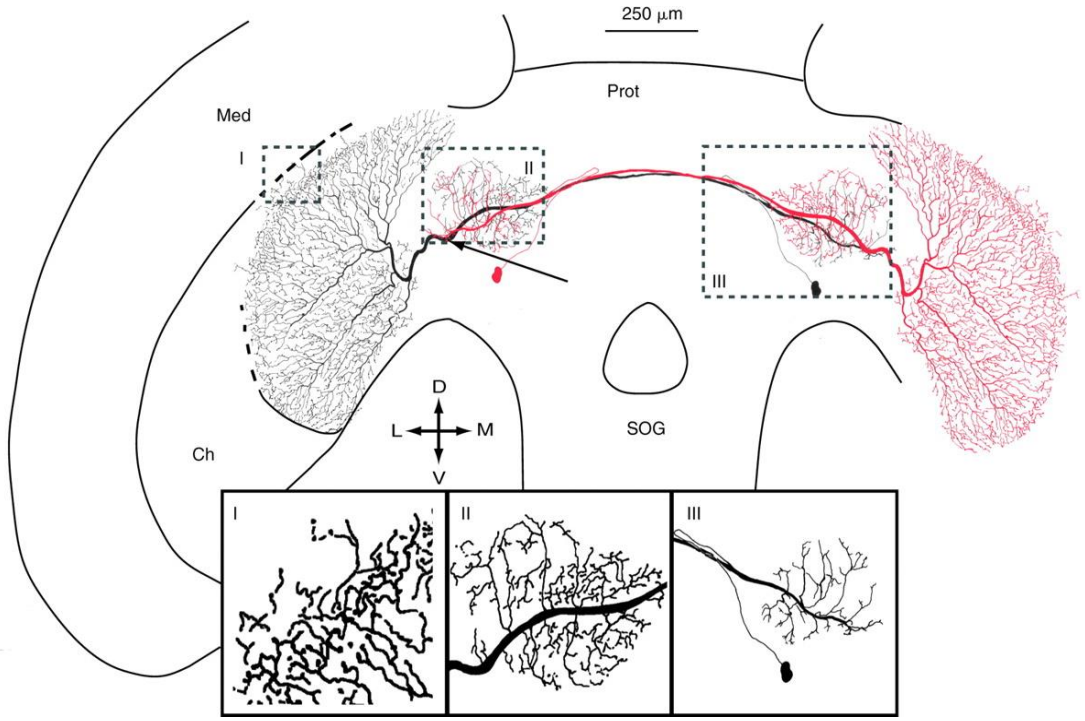

Figure 14: Reproduced from Geurten et al 2007. Shows an example of the layout of CSTMD1 and its mirror-symmetric counterpart. Cell body is contralateral to the recording site (left axon). CSTMD1 receives inputs from the midbrain and synapses onto the distal lobula.

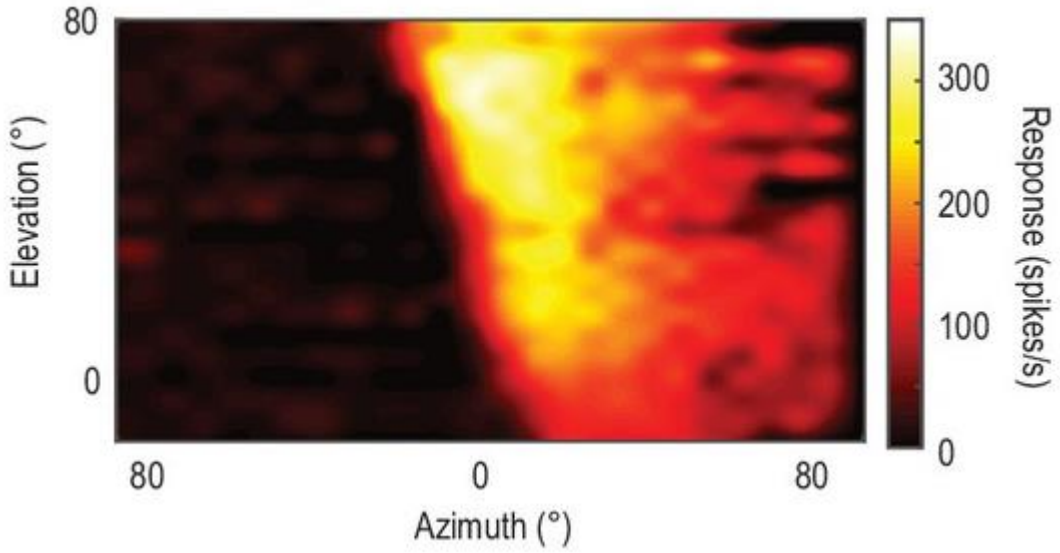

Figure 15: Reproduced from Wiederman & Fabian et al., 2017, Shows an example of a CSTMD1 receptive field. CSTMD1 is strongly excited by small targets moving in the right hemifield and inhibited by small targets moving in the left hemifield.

2.3.11.6 Small Target Motion Detectors of Flies

STMDs in Eristalis are very similar to their dragonfly counterparts preferring targets subtending small angles (cells were found with the following optimums: 0.8, 1.6 or 3.0°, Nordström & O'Carroll 2006, Barnett et al., 2007). These neurons occur across sexes (though with some differences) which implicates them in more than just mate pursuit behaviours. STMDs in Eristalis are divided by their receptive field sizes (large and small – described below).

Large Field STMDs

Large Field STMDs (simply referred to as STMDs) have receptive fields that subtend more than 10° of the visual field and represent a broad class of subtypes. In general, these STMDs have no spontaneous behaviour but can respond robustly (up to 200 spikes/s) to small targets moving anywhere within their receptive field (Nordström et al., 2006).

While very little formal classification of Eristalis STMDs exist, they vary significantly in the size and shape of their receptive fields, direction-selectiveness, preferred direction and mode (spiking/graded, Nordström et al., 2006). Several possess large receptive fields, subtending extensive regions of the ipsilateral hemifield with no apparent inter-hemispheric interactions (Nordström & O'Carroll 2006).

Perhaps the most interesting distinctions in these large-field STMDs occurs with their responses in cluttered scenarios. While some detect targets robustly when subjected to

background motion, others exhibit full inhibition. Others instead respond to feature-like regions of the background stimulus (Nordström et al., 2006).

Small Field STMDs

Within the family of STMDs, there are the so-called Small-Field STMDs (sfSTMDs, Nordström & O'Carroll, 2006). These neurons are defined primarily by their small receptive fields (<10° of visual space, Barnett et al., 2007). Anatomically, these neurons take their inputs from the dorsal Lobula (at least those investigated) and output into the ipsilateral and contralateral midbrain depending on subtype (Barnett et al., 2007).

Though normally dormant (no spontaneous), they respond robustly (up to 400 spikes/s) to small targets (1.6 degrees) moving through their receptive field. These neurons can be both direction-selective and non-direction selective and maximally respond to small targets, showing lower response to larger objects (bars) and no response to gratings.

2.3.11.7 Drosophila Variants (Lobula Columnar – LC)

Similarly tuned small object detecting neurons have been found in *Drosophila* (Keles & Frye 2017). These neurons are end-stopped (i.e. non-responsive to bars or gratings) and have omnidirectional direction preference. Interestingly, upon blocking surrounding inhibitory connections, these neurons not only regained their preference for elongated objects (i.e. reverting to pure motion detectors) but also lost some of their response to small targets (Keles & Frye 2017).

2.3.12 Vision Alternatives

Though vision is perhaps the most vital component for navigation, other senses such as olfaction can play a huge role. Of particular note is the existence of halteres (Figure 16) in flies (Nalbach & Hengstenberg 1994; Dickinson 1999). These vestigial wings-turned-gyros can detect the rate of angular turn at a lower latency than visual inputs. Both halteres and vision both play a role in gaze stabilization (Hengstenberg 1991) working in tandem to allow flies to encode a broad range of velocities (Sherman & Dickinson 2003; Schwyn et al., 2011).

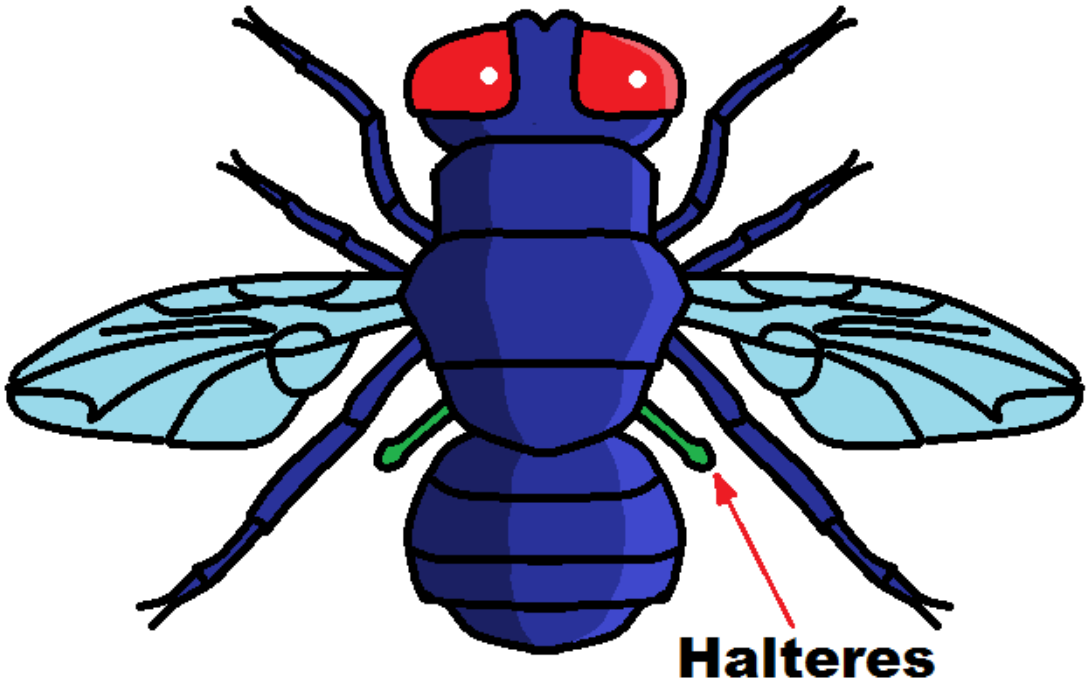

Figure 16: Fly Halteres are used in tandem with visual input for course stabilization via attitude control and gaze stabilization

2.3.13 Central Complex

Once outside the optic lobes, information becomes more and more abstract, the result of the integration of large amounts of information. None-the-less, certain regions of the central brain contain neuropils whose function maps well to behavioural and visual input. One of these better understood regions is the central complex, a collection of modular neuropils across the midline of the insect brain which receives the bulk of its input from visual areas (Pfeiffer & Homberg 2014).

2.4 Physiology

2.4.1 General Topics

2.4.1.1 Adaptation

One common process, which occurs in many neurons in many different contexts and for many purposes, is that of neuronal adaptation. Adaptation occurs when a neuron is exposed to an unchanging stimulus for an extended period of time. While the neuron may have an initially stronger response, over time the magnitude of the response reduces. Adaptation has many advantages including energy saving. By encoding only changes in a stimulus, the energy required to generate spikes can be minimized allowing neurons to remain silent (and thus energy efficient) in absence of changes. In addition to a simple energy saving measure,

adaptation can also allow a system to become more sensitive to small changes. Adaptation can also provide benefits in normalisation, for example in visual processing such as reducing the variation induced from changes in scene statistics (Brinkworth & O'Carroll 2009).

2.4.1.2 Contrast Sensitivity

Tied to the concept of spatial and temporal tuning is the concept of contrast sensitivity. As seen in Figure 10, at certain spatial frequencies, it is difficult to distinguish the sinusoidal pattern. This is due to differing contrast sensitivities between different spatial frequencies in human vision. This also applies to insects and one important measure used to describe (and more importantly compare) different insect vision systems is the contrast sensitivity in the most favourable conditions (i.e. matched temporal and spatial tuning).

Contrast sensitivity is measured based on the inverse of contrast itself. Detection of a lower contrast implies a higher sensitivity. As common measures of contrast range from 0-1, sensitivity can likewise vary from 1 to ∞.

Contrast sensitivity varies significantly among insects ranging as low as 25-40 in the fly (Dvorak et al., 1980) and sphingid moths (O'Carroll et al., 1996, O'Carroll et al., 1997, Straw et al., 2006) to 40-100 in *Eristalis* and *Volucella pellucens* (O'Carroll et al., 1996, O'Carroll et al., 1997, Straw et al., 2006). Contrast sensitivity is also relevant to feature detection. For example, STMDs respond to targets whose contrasts matches the sensitivity of dragonfly photoreceptors (O'Carroll & Wiederman 2014).

2.4.1.3 Environmental Tuning

The fundamental mechanisms of motion detection have already been discussed. However, one factor which governs the specific fine-tuning of any neuronal system is the environment and behavioural modes in which the animal operates. The motion detection needs of hovering are different from those of patrolling. Additionally, the visual content in any scene (and the remaining information after optics has altered it) plays an important role in shaping a creature's motion pathways. For example, in guinea pigs, OFF retinal ganglion cells have smaller receptive fields and more common distribution, reflecting the presence of dark and light features found in natural scenes (Ratliff et al., 2010).

2.4.1.4 Phi, Reverse Phi & Higher Order Motion

One of the basic principles that underlie correlation-based motion detection is the concept that a feature being detected occurs at two points in space sequentially. This can be as simple as two light flashes or two dark flashes and is often described as Fourier motion or phi

motion. However, this is not the only form of motion that exists. When the polarities of the light flashes are reversed (i.e. light-dark or dark-light), this also creates a percept of motion (Egelhaaf & Borst 1992), even in humans (Bours et al., 2009). This form of motion is called reverse-phi motion (Clark et al., 2011, Tuthill et al., 2011). Reverse-phi motion causes a reverse optomotor response in Drosophila and exhibits similar tuning properties (Tuthill et al., 2011). Interestingly, this response inverts (i.e. returning to forward motion) when a sufficiently high flicker/temporal frequency is reached (Tuthill et al., 2011).

The presence of reverse-phi motion sensitivity is evident for a correlation-based model of motion detection. Basic EMDs also exhibit this effect. The presence of reverse-phi motion thus appears to be intrinsic to motion detection found in invertebrates. Experiments involving the silencing of L1 & L2 neurons in the lamina showed that each neuron was involved in complementary reverse-phi detection with L1 detecting dark-bright and L2 detecting bright-dark responses (Clark et al., 2011).

Alternate forms of non-fourier motion such as theta motion also exist. In theta motion a textured object moves, while the texture of the object moves in the direction opposite to the objects itself. This kind of complex motion elicits both a preferred direction and non-preferred direction response in H1 neurons (Quenzer & Zanker 1991).

Non-fourier motion can become even more complex. Instead of simple inversions, motion can be generated via strict parity-like rules. These forms of motion exhibit no frame-to-frame correlations, instead relying on higher-order pattern (Hu & Victor 2010), which produce motion percepts in humans and flies (Clark et al., 2014).

2.4.1.5 Temporal vs Spatial Tradeoffs

The detection of widefield motion is inherently one of trade-offs between detection of high spatial frequency information (detail) and high temporal frequency information (speed). By integrating over longer periods of time, temporal information can be directly traded for spatial information but at the cost of introducing delays. This is perhaps most easily seen in hawkmoths, whose species, though similar operate in very different lighting environments. It has been shown for example that *Manduca sexta* (a crepuscular hawkmoth) has higher spatial acuity than its cousin *Deilephila elpenor* (a nocturnal hawkmoth). However, both are tuned to similar slow velocities, facilitating a common feeding behaviour. This is done at the expense of temporal delays in the nocturnal *Deilephila elpenor* (Theobald et al., 2010).

2.4.1.6 Spatial Pooling

Neurons that respond over a wide area of an insect's visual field necessarily involve the pooling of responses from large numbers of inputs. For example, evidence in landing flies indicates that behaviour depends on the integration of many motion units (Borst & Bahde 1988). In WFMS neurons, the size of any given visual input and the response characteristics of a variety of wide-field cells has been shown to be highly non-linear (Borst et al., 1995). Exciting only small subsections of a receptive field can generate near full-screen responses. However, this research was limited to sinusoidal grating stimuli. When considering more realistic natural image input, how might these non-linearities operate?

As previously detailed, ego-motion in insects is performed via integrating the responses of multiple small motion detector elements. Natural scenes conform to similar fundamental statistical distributions. In general the total frequency content of a scene will contain more power at lower spatial frequencies (i.e. large details) than at high spatial frequencies (i.e. fine details). This distribution is given the shorthand 1/f (where f is frequency) indicating that each decade of frequencies contains largely the same amount of power. Despite this general rule, subsections of an image exhibit significant variation. Theoretical modelling (Dror et al., 2000) of correlator-based motion has demonstrated that the similar fundamental statistics of natural scenes causes correlation-detection to effectively encode velocity. This differs from results drawn from sinusoidal gratings where tuning is defined by the temporal frequency, rather than the velocity. However, this modelling assumes a broad distribution of frequencies as seen when aggregating input over a large field of view.

One might expect that reducing the receptive field of a wide-field neuron, subjects the response of a neuron to more pattern noise (variation due to scene variation). Indeed, even neurons with large receptive fields can be subject to local variability with increased receptive field size only has a modest effect on pattern variability (O'Carroll et al., 2011). This is despite a general normalisation across images where different image phases are averaged (Straw et al., 2008).

2.4.2 Target Tracking

2.4.2.1 Spatial Facilitation

Facilitation as a general principle merely represents one system preferentially influencing the response of another system and, hence facilitating its response. In this broad definition, nearly all neurons are the subject of some form of facilitation and specifically correlation-based motion employing a non-linear component is essentially a facilitatory interaction. In the

context of target-tracking, spatial facilitation refers to a higher-order process, independent of the fundamental motion detection paradigm which causes targets which have been tracked for a period of time to experience preference over those which have not. This manifests as a preference for long paths (Figure 17).

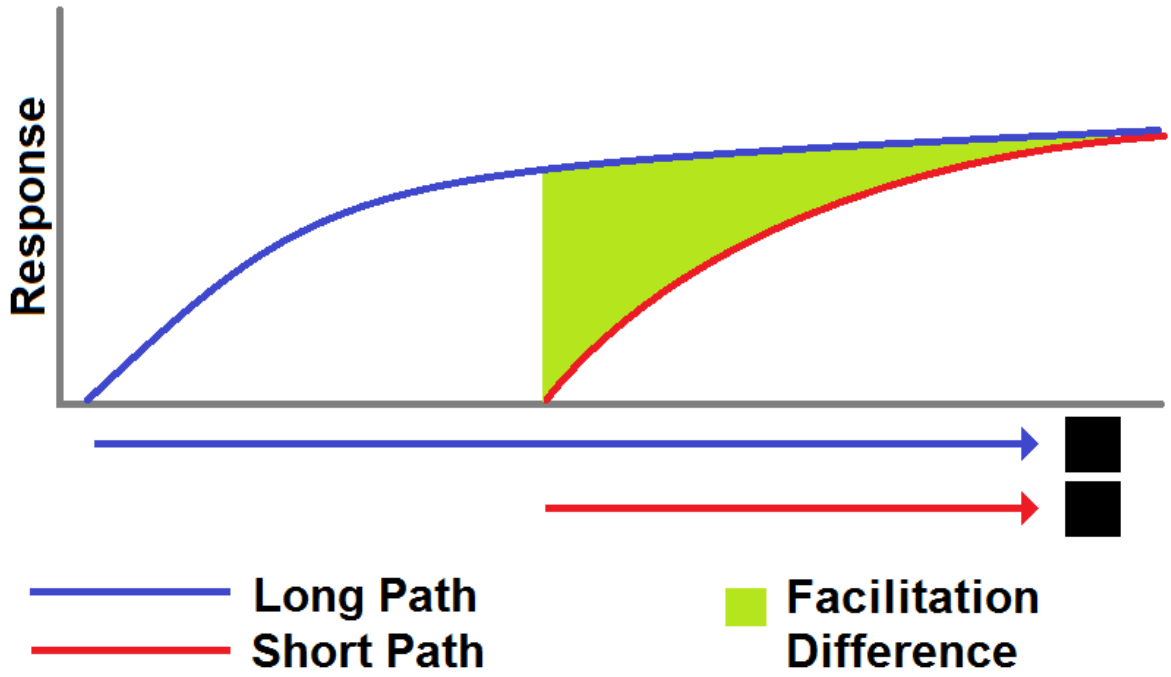

Figure 17: Facilitation causes targets to build up their response over time preferencing long paths over short paths

The first evidence of spatial facilitation in insects was found in STMD neurons in dragonflies (Nordström et al., 2011). It was found that targets that follow long trajectory paths have higher responses than those, which travel on shorter paths. In essence, presenting a target 'primer' on a path prior to the test path increases (to a limit) the response of the test path. Further research of dragonfly facilitation has identified numerous properties of the phenomenon. For example, while facilitation effectively slows the onset of maximal response to a small target, the offset response remains unchanged. Additionally, the enhancement caused by facilitation persists over time, allowing for a target to be temporarily occluded. These properties rule out a simple change to time-constants for motion detection.

The effect of facilitation is not global as paths that are discontinuous show no benefit (Dunbier et al., 2012) and can experience inhibition (Wiederman & Fabian et al., 2017). Facilitation is also direction dependent, preferring plausibly moving targets that travel in straight trajectories or exhibit only moderate turning. Contrast only plays a small role in facilitation as dim targets can generate significant facilitation (Wiederman & Fabian et al.,

2017). Finally, facilitation appears to be highly dependent on the velocity of the target, maximizing for targets that move at the optimal speed of STMD neurons (Fabian et al., in preparation). A plausible model for facilitation has been modelled through the neuron BSTMD1 (Binocular Small Target Motion Detector 1) via retinotopically arrayed NMDA receptors (Bekkouche et al., 2017).

2.4.2.2 Selective Attention

Selective Attention refers to a process where a single system can ignore a distracting stimulus in favour of an excitatory stimulus. Examples of selective attention exist in primates (Treue & Maunsell 1996) where cuing a monkey to attend to a specified object can alter the response properties of the underlying neuronal responses. In physiological terms, it describes a process where two competing stimuli are presented but the response matches closely to data as if only a single stimulus was presented at any given time (Figure 18). While this does not preclude complex switches, it limits the system to a single focus of attention at a single point in time. This is of particular relevance to predatory insects, which often benefit from hunting from swarms of prey (Combes et al., 2012).

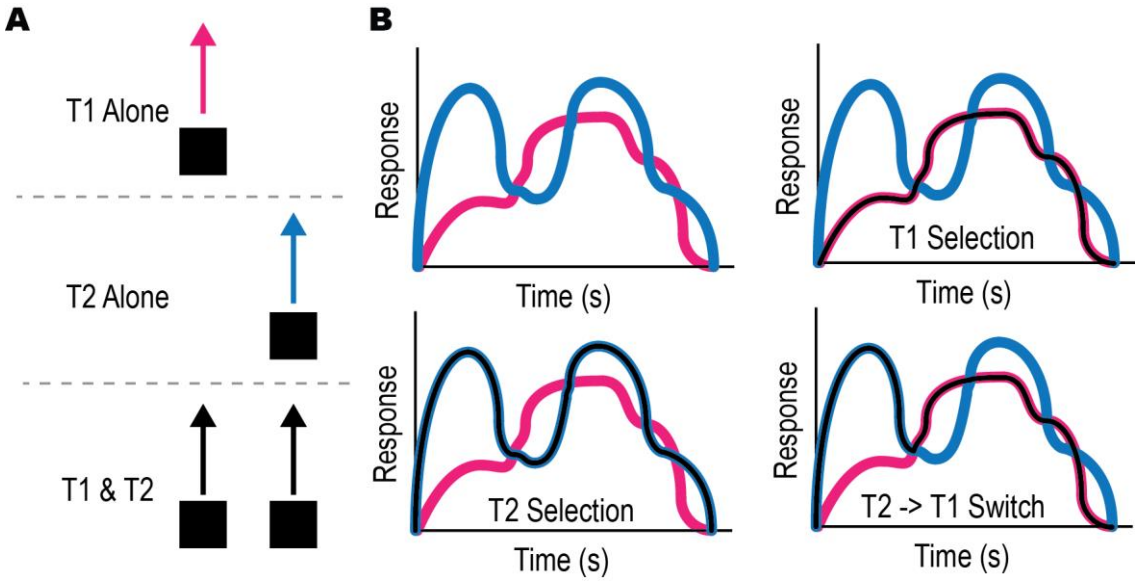

Figure 18: Selective attention example with two targets, T1 & T2. A) Example of the three stimulus conditions. T1 presented alone (top), T2 presented alone (middle), T1 & T2 presented together (bottom). B) Example of T1 and T2 responses when presented alone (top left) and three examples of what could happen when two targets are presented together. Either T1 is responded to ignoring T2 (top right), T2 is responded to ignoring T1 (bottom left) or some combination, in this case the neurons responds to T2 and then switches attention to T1.

CSTMD1 was found to be subject to long-range inhibition and contralateral inhibition (Bolzon et al., 2009). In effect, when two targets were presented, the primary target exhibited

reduced response if the secondary target was presented in the excitatory hemifield and complete suppression if it was presented in the inhibitory hemifield. Subsequent experimentation by Wiederman & O'Carroll 2013, further elucidated this inhibitory interaction demonstrating that CSTMD1 selectively attends to individual targets whilst suppressing others. It is even capable of switches (Wiederman & O'Carroll 2013).

2.4.2.3 Target Tracking in Natural Images and Clutter

One challenge for insect target tracking is the cluttered environments in which these activities often occur. Background clutter can interfere with target tracking through several different mechanisms such as changes in local contrast obscuring the target and objects in the background acting as competing features for target detection systems. The performance of STMD neurons in cluttered scenarios has been tested previously in both flies (*Eristalis*) and dragonflies. In *Eristalis*, it was found that a subset of STMD neurons could detect targets in cluttered backgrounds (Nordström et al 2006). These experiments used artificially generated clutter and showed that detection was still strong when the target and background were both in motion in a variety of different speeds (both parallel and anti-parallel motion). In dragonflies, similar experiments were performed using panoramas of natural scenes (Wiederman & O'Carroll 2014) though with targets embedded in the background and no relative motion.

To properly understand these two results requires an understanding of what constitutes a natural image and what features exist within a natural image. A first observation is that spatial frequency spectral analysis reveals a great deal of similarity between otherwise disparate images. In general, most images have a spectral density function approximating $1/f^2$ where f is the spatial frequency (Field 1987, Tolhurst 1992). In essence, this means that the total amount of 'power' (as measured in sinusoidal frequency power) in any 'octave' (i.e. a fixed unit in log-space) of frequency space is roughly constant. In layman's terms, there is more low-spatial frequency information than high-spatial frequency information.

However, while this interpretation of natural images is quite suitable for wide-field motion analysis (i.e. Dror et al., 2000), it falls down when considering target tracking. For example, in the three images presented in Figure 19, while the top two images might conform well to a $1/f^2$ power law, the bottom image contains obviously different frequency information vertically compared to horizontally. Even so, the two upper images may have similar broadband statistics, but have entirely different features within said images.

Figure 19: Examples of natural imagery

A more extreme example of this can be seen in Figure 20. The two images have identical broadband statistics having been generated from the same original image. In both images, a 2-dimensional fast Fourier transform was performed followed by the equivalent inverse transform. However, in the second image, prior to the inverse transform, the phase of the frequency information was randomized. As is apparent from the transformed image, the 'features' of the image have been entirely erased and the image is unrecognisable. This indicates that in a feature detection space, two images with similar statistics may be drastically different when considering their higher-order features.

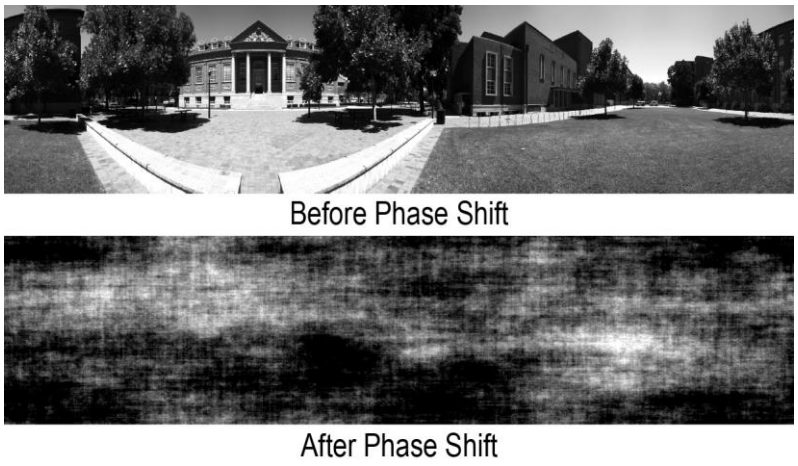

Figure 20: Example of the importance of phase information in image intelligibility. Top: An image with intact phase. Bottom: The same image with the phase of the frequency content randomized.

2.5 Modelling

2.5.1 Photoreceptors

The first role of any photoreceptor model is to reconcile the fact that the variation in photon intensity in any given scene is quite extreme, ranging of several orders of magnitude. Photoreceptors accommodate this via adaptation, allowing them to detect changes in light intensity at a large range of different mean light levels. This ability to operate at a vast range of photon density is enabled via several time-courses of adaptation as exemplified in various models (van Hateren & Snippe 2001, Brinkworth et al., 2007). Fortunately, though a completely accurate photoreceptor model is challenging to generate, its effects on downstream units are modest. Alternative computationally simpler approximations of photoreceptors include using a Lipetz function followed by a single low-pass filter (Wiederman et al., 2008) in scenarios with limited light intensity changes.

2.5.2 Elementary Motion Detectors

One of the core assumptions underlying the behaviour of WFMS neurons is that their large receptive fields exist due to the integration of many smaller motion sensitive units. These Elementary Motion Detectors (EMDs) are thought to calculate motion at the most atomic level (Borst & Euler 2011). As WFMS neurons have direction selectivity, it is necessary that any such model must integrate information that is spatially distinct, (i.e. there must be inputs from adjacent areas of space). While this theoretical model is possible, it is more likely that any neuronal system that encodes motion likely does so using more than a single pair of ommatidia (Gruntman et al., 2018). Unsurprisingly, the biological substrate for these proposed EMDs has proved elusive (Clark et al., 2011). Here I will discuss the two most prominent EMD models.

2.5.2.1 Hassenstein-Reichardt Correlators

Perhaps the most well-known form of EMD is the 1956 Hassenstein-Reichardt Correlator (HRC) (Hassenstein & Reichardt 1956). The HRC works by taking the input from two adjacent locations, delaying the input from one arm and then performing a non-linear operation (usually a multiplication) on the two inputs (Figure 21A). This generates a motion-dependent signal. Fourier analysis of the optomotor response of *Drosophila* indicates that the response can be predicted by HRC inspired motion detectors connecting nearest neighbour and next-nearest neighbour ommatidia (Buchner 1976; Schuling 1989). Certainly, basic

modelling of an HRC yields many of the basic properties of WMFS neurons such as their tuning to specific spatial and temporal frequencies (see Physiology for more information).

2.5.2.2 Barlow-Levick Correlators

An alternative model for EMDs was proposed based on recordings from rabbit retina (Barlow & Levick 1965). This EMD works in a complementary sense to the HRC, featuring two input arms with one subject to a delay and a nonlinear interaction. However, unlike the HRC the Barlow-Levick Correlator (BLC) implements an inhibitory connection (Figure 21B). This inhibition functions by acting as a gating mechanism to the undelayed arm. Unlike the HRC, the BLC works by suppressing Null-Direction (ND) motion rather than amplifying Preferred Direction (PD) motion.

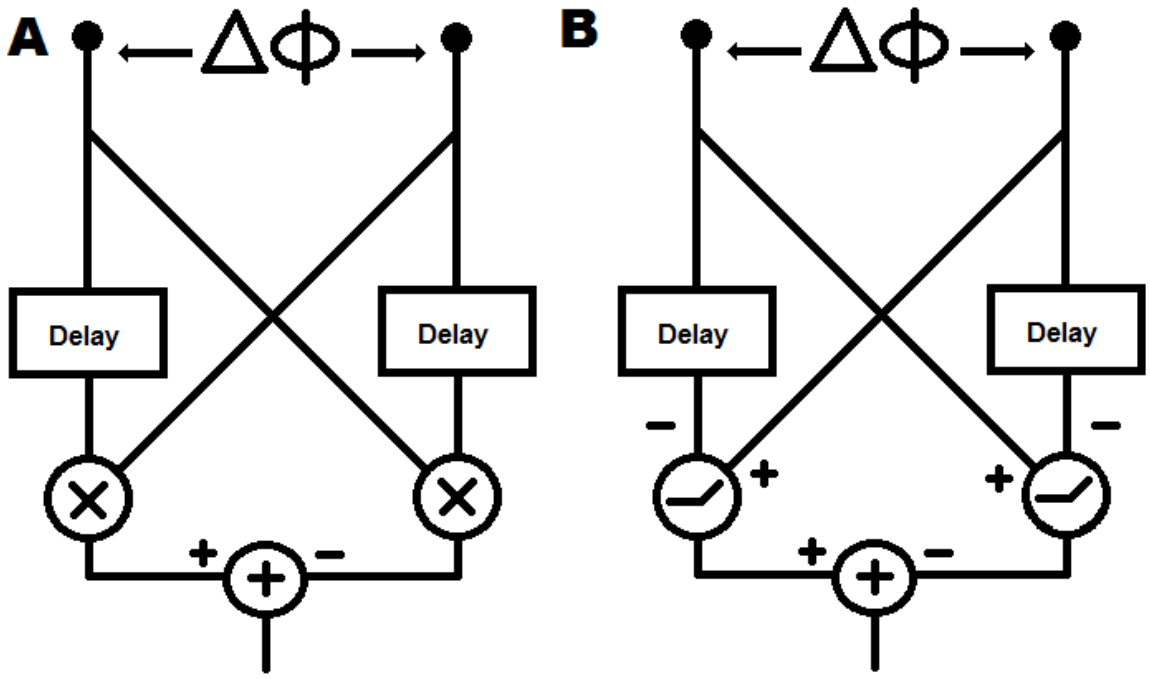

Figure 21: A) Example of a Hassenstein-Reichardt Correlator. The correlator consists of two half correlators which are subtracted at the end (to create motion asymmetry). Each half-correlator consists of inputs from two adjacent ommatidia, one of which is subject to a delay before multiplication. B) Example of a Barlow-Levick Correlator. The correlator consists of two half correlators which are subtracted at the end (to create motion asymmetry). Each half-correlator consists of inputs from two adjacent ommatidia, one of which is subject to a delay. The delayed arm acts as a gating mechanism for the response for the undelayed arm.

2.5.3 Spatial Pooling

A tacit assumption of the EMD model is that WFMS neurons must take their inputs from a vast array of such EMDs in order to function. How WFMS neurons integrate this information is not fully understood, though several electrophysiological and behavioural studies indicate

some limitations. Firstly, it has been shown behaviourally that the extent of a pattern when stimulating flies is only relevant for very small pattern sizes (Reichardt et al., 1983). This is despite stimuli of different velocities eliciting very different responses. In order to explain this, a form of gain control was proposed, initially via an intermediary pool cell (Reichardt et al., 1983) and later via dendritic computations (Borst et al., 1995).

2.5.4 Proposed Components of EMDs in Insects

The WFMS neuron equivalent in flies is the Lobula Plate Tangential Cell (LPTC). However, these are not the most elementary directional motion sensitive cells discovered. T4 & T5 (see Anatomy) are small-field motion sensitive neurons that connect the Lobula Plate to the Medulla and Lobula respectively. These neurons exhibit directional motion sensitivity (Maisak et al., 2013) and synapse onto LPTCs. It is thought that T4 & T5 represent the computational outputs of EMDs. The inputs to these T4 & T5 EMDs is a bit more complex, consisting of a multitude of transmedullary and intermedullary neurons. Each of these interneurons exhibit differing temporal properties some of which can be (reasonably) simplified into elementary delay filters (Arenz et al 2017). Currently, the inputs which might form HRC or BLC-like EMDs are proposed to be the following groups.

T4: Mi1, Tm3, Mi4 & Mi9

T5: Tm1, Tm2, Tm4 & Tm9

2.5.5 STMD Modelling

STMD neurons are thought to receive their inputs from many of the same early visual processing regions as their WFMS counterparts. As such, when modelling STMDs some of the subsystems (namely the retina and lamina) can be replicated. Where STMDs differ in their response to WFMS neurons is not their correlation-based responses, but rather what they correlate. It has been suggested (Wiederman et al., 2008, Wiederman et al., 2013) that STMD neurons correlate ON and OFF signals and that replacing the inputs to EMD-like units with ON and OFF inputs. While this can explain the preference for small targets in the direction parallel to motion, the orthogonal direction requires a separate explanation. As STMDs respond weakly to elongated features a parsimonious explanation is that they receive inhibitory input from their surround. With both inhibitory surround and ON-OFF correlation, the size preference of STMDs is easily reproduced.

How the higher-order properties of STMD neurons might be accurately modelled remains an area of active research, however it has been shown that modelling these higher order

properties (such as facilitation) can result in improved model performance (Bagheri et al., 2015).

2.6 Aims & Scope

This introduction has covered an extensive range of topics surrounding the main work of this PhD. Within this general framework, my PhD has focussed on a broad range of subtopics, all joined by the theme, dragonfly perception in natural environments. Here I present my main research questions and some related work performed during my PhD.

2.7 Research Questions

My main research questions are as follows:

1. How do dragonflies encode wide-field motion and how do these systems operate in naturalistic environments?
2. How do natural images interact with the dragonfly target-tracking systems?
3. How might insects recover information from ambiguous input?

The following chapters detail the work I performed during my project to address theses aims.

2.8 Topic Change

My PhD was originally established to investigate the context of dragonfly behaviour in natural scenes via behavioural capture. This project progressed significantly during my first six months, however, a combination of competing research laboratories, a very weak initial dragonfly season and my general interests steered me away from this area. None-the-less, this project produced two published research articles which are included in appendices. These articles describe methods for multi-camera integration for insect behaviour capture using low-cost cameras.

2.9 References

1) Arenz A, Drews MS, Richter FG, Ammer G & Borst A, 2017, The Temporal Tuning of the Drosophila Motion Detectors is Determined by the Dynamics of Their Input Elements, Current Biology (vol. 27, no. 7, pp. 929-944), https://doi.org/10.1016/j.cub.2017.01.051

2) Arikawa K, Mizuno S, Scholten DG, Kinoshita M, Seki T, Kitamoto J & Stavenga DG, 1999, An ultraviolet absorbing pigment causes a narrow-band violet receptor and a single-peaked green receptor in the eye of the butterfly Papilio, Vision Research (vol. 39, no. 1, pp. 1-8), https://doi.org/10.1016/S0042-6989(98)00070-4

3) Arnett DW, 1971, Receptive Field Organization of Units in the First Optic Ganglion of Diptera (vol. 173, no. 4000, pp. 929-931)

4) Autrum H & Kolb G, 1968, Spektrale Empfindlichkeit einzelner Sehzellen der Aeschniden, Zeitschrift für vergleichende Physiologie (vol. 60, no. 4, pp. 450-477). https://doi.org/10.1007/BF00297940

5) Bahl A, Ammer G, Schilling T & Borst A, 2013, Object tracking in motion-blind flies, Nature Neuroscience (vol. 16, pp. 730-738), https://doi.org/10.1038/nn.3386

6) Baird JM & May ML, 1997, Foraging behavior of Pachydiplax longipennis (Odonata Libellulidae), Journal of Insect Behaviour (vol. 10, no. 5, pp. 655-678), https://doi.org/10.1007/BF02765385

7) Barlow HB & Levick WR, 1965, The mechanism of directionally selective units in rabbit's retina, Journal of Physiology (vol. 178, pp. 477-504) https://doi.org/10.1113/jphysiol.1965.sp007638

8) Barnett PD, Nordstrom K & O'Carroll DC, 2007, Retinotopic organization of Small-Field-Target-Detecting Neurons in the Insect Visual System, Current Biology (vol. 17, pp. 569-578), https://doi.org/10.1016/j.cub.2007.02.039

9) Barnett PD, Nordstrom K & O'Carroll DC, 2010, Motion Adaptation and the Velocity Coding of Natural Scenes, Current Biology (vol. 20, pp. 994-999), https://doi.org/10.1016/j.cub.2010.03.072

10) Behnia R, Clark DA, Carter AG, Clandinin TR & Desplan C, 2014, Processing properties of ON and OFF pathways for Drosophila motion detection, Nature (512, pp. 427-430), https://doi.org/10.1038/nature13427

11) Bekkouche B., Shoemaker P.A., Fabian J., Rigosi E., Wiederman S.D., O'Carroll D.C. (2017) Multicompartment Simulations of NMDA Receptor Based Facilitation in an Insect Target Tracking Neuron. In: Lintas A., Rovetta S., Verschure P., Villa A. (eds) Artificial Neural Networks and Machine Learning – ICANN 2017. ICANN 2017. Lecture Notes in Computer Science, vol 10613. Springer, Cham, https://doi.org/10.1007/978-3-319-68600-4_46

12) Boeddeker N, Dittmar L, Sturzl W & Egelhaaf M, 2010, The fine structure of honeybee head and body yaw movements in a homing task, Proceedings Biological Sciences (vol. 277, no. 1689, pp. 1899-1906) https://doi.org/10.1098/rspb.2009.2326

13) Bolzon DM, Nordstrom K & O'Carroll DC, 2009, Local and Large-Range Inhibition in Feature Detection, Journal of Neuroscience (vol. 29, no. 45, pp. 14143-14150), https://doi.org/10.1523/JNEUROSCI.2857-09.2009

14) Borst A & Bahde S, 1986, What kind of movement detector is triggering the landing response of the housefly?, Biological Cybernetics (vol. 55, no. 1, pp. 59-69).

15) Borst A & Bahde S, 1988, Spatio-Temporal Integration of Motion, Naturwissenschaften (vol. 75, pp. 265-267)

16) Borst A, Egelhaaf M & Haag J, 1995, Mechanisms of Dendritic Integration Underlying gain Control in Fly Motion-Sensitive Interneurons, Journal of Computational Neuroscience (vol. 2, pp. 5-18), https://doi.org/10.1007/BF00962705

17) Borst A, Haag J & Reiff DF, 2010, Fly Motion Vision, Annual Review of Neuroscience (vol. 33, pp. 49-70) https://doi-org.proxy.library.adelaide.edu.au/10.1146/annurev-neuro-060909-153155

18) Borst A & Euler T, 2011, Seeing Things in Motion: Models, Circuits, and Mechanisms, Neuron (vol. 71, pp. 974-994) https://doi.org/10.1016/j.neuron.2011.08.031

19) Borst A & Helmstaedter, 2015, Common circuit design in fly and mammalian motion vision, Nature Neuroscience (vol. 18, pp. 1067-1076), https://doi.org/10.1038/nn.4050

20) Bours RJE, Kroes MCW & Lankheet MJ, 2009, Sensitivity for reverse-phi motion, Vision Research (vol. 49, no. 1, pp. 1-9), https://doi.org/10.1016/j.visres.2008.09.014

21) Brinkworth RS, Mah EL, Gray JP & O'Carroll DC, 2007, Photoreceptor processing improves salience facilitating small target detection in cluttered scenes, Journal of Vision (vol. 8, no. 11, pp. 1-17), https://doi.org/10.1167/8.11.8

22) Brinkworth RSA & O'Carroll DC, 2009, Robust Models for Optic Flow Coding in Natural Scenes Inspired by Insect Biology, PloS Computational Biology (vol. 5, no. 11, e1000555), https://doi.org/10.1371/journal.pcbi.1000555

23) Buchner E, 1976, Elementary movement detectors in an insect visual system, Biological Cybernetics (vol. 24, no. 2, pp. 85-101).

24) Buchner E & Buchner S, 1984, Neuroanatomical Mapping of Visually Induced Nervous Activity in Insects by 3H-Deoxyglucose. In: Ali M.A. (eds) Photoreception and Vision in Invertebrates. NATO ASI Series (Series A: Life Sciences), vol 74. Springer, Boston, MA, https://doi.org/10.1007/978-1-4613-2743-1_17

25) Chappell RL & DeVoe RD, 1975, Action spectra and chromatic mechanisms of cells in the median ocelli of dragonflies, Journal of General Physiology (vol. 65, no. 4, pp. 399-419), https://doi.org/10.1085/jgp.65.4.399

26) Chiappe ME, Seelig JD, Reiser MB & Jayaraman V, 2010, Walking Modulates Speed Sensitivity in Drosophila Motion Vision, Current Biology (vol. 20, no. 16, pp. 1470-1475), https://doi.org/10.1016/j.cub.2010.06.072

27) Clark DA, Bursztyn L, Horowitz MA, Schnitzer MJ & Clandinin TR, 2011, Defining the Computational Structure of the Motion Detector in Drosophila, Neuron (vol. 70, no. 6, pp. 1165-1177) https://doi.org/10.1016/j.neuron.2011.05.023

28) Clark DA, Fitzgerald JE, Ales JM, Gohl DM, Silies MA, Norcia AM & Clandinin TR, 2014, Flies and humans share a motion estimation strategy that exploits natural scene statistics, Nature Neuroscience (vol. 17, pp. 296-303), https://doi.org/10.1038/nn.3600

29) Clifford CWG & Langley K, 1996, A Model of Temporal Adaptation in Fly Motion Vision, Vision Research (vol. 36, no. 16, pp. 2595-2608).

30) Collett TS & Land MF, 1978, How Hoverflies Compute Interception Courses, Journal of Comparative Physiology, (vol. 125, pp. 191-204).

31) Combes SyA, Rundle DE, Iwasaki JM & Crall JD, 2012, Linking biomechanics and ecology through predator-prey interactions: flight performance of dragonflies and their prey, Journal of Experimental Biology (vol. 215, pp. 903-913), https://doi.org/10.1242/jeb.059394

32) Combes SA, Salcedo MK, Pandit MM & Iwasaki JM, 2013, Capture Success and Efficiency of Dragonflies Pursuing Different Types of Prey, Integrative and Comparative Biology (vol 53, no. 5, pp. 787-798), https://doi.org/10.1093/icb/ict072

33) Coombe PE & Heisenberg M, 1986, The Structural Brain Mutant Vacuolar medulla of Drosophila melanogaster with Specific Behavioral Defects and Cell Degeneration in the Adult, Journal of Neurogenetics (vol. 3, pp. 135-158) https://doi.org/10.3109/01677068609106845

34) Corbet PS & May ML, 2008, Fliers and perchers among Odonata: dichotomy or multidimensional continuum? A provisional reappraisal, International Journal of Odonatology (vol. 11, pp. 155-171), https://doi.org/10.1080/13887890.2008.9748320

35) Dickinson MH, 1999, Haltere-mediated equilibrium reflexes of the fruit fly, Drosophila melanogaster, Philosophical Transactions of the Royal Society B (vol. 354, pp. 903-916), https://doi.org/10.1098/rstb.1999.0442

36) Dror RO, O'Carroll DC & Laughlin SB, 2000, The Role of Natural Image Statistics in Biological Motion Estimation, Biologically Motivated Computer Vision (pp. 492-501), https://doi.org/10.1007/3-540-45482-9_50

37) Dubner R & Zeki SM, 1971, Response properties and receptive Fields of cells in an anatomically defined region of the superior temporal sulcus in the monkey, Brain Research (vol. 35, no. 2, pp. 528-532), https://doi.org/10.1016/0006-8993(71)90494-X

38) Duffy CJ & Wurtz RH, 1991, Sensitivity of MST neurons to optic flow stimuli. I. A continuum of response selectivity to large-field stimuli, Journal of Neurophysiology (vol. 65, no. 6, pp. 1329-1345), https://doi.org/10.1152/jn.1991.65.6.1329

39) Duffy CJ, 1998, MST Neurons Respond to Optic Flow and Translational Movement, Journal of Neurophysiology (vol. 80, no. 4, pp. 1816-1827), https://doi.org/10.1152/jn.1998.80.4.1816

40) Dunbier JR, Wiederman SD, Shoemaker PA & O'Carroll DC, 2012, Facilitation of dragonfly target-detecting neurons by slow moving features on continuous paths, Frontiers in Neural Circuits (vol. 6, 79) https://doi.org/10.3389/fncir.2012.00079

41) Dvorak D, Srinivasan MV & French AS, 1980, The Contrast Sensitivity of Fly Movement-Detecting Neurons, Vision Research (vol. 20, pp. 397-407), https://doi.org/10.1016/0042-6989(80)90030-9

42) Eckert H & Dvorak DR, 1983, The centrifugal horizontal cells in the lobula plate of the blowfly, Phaenicia sericata, Journal of Insect Physiology (vol. 29, no. 7, pp. 547-553, 555-560), https://doi.org/10.1016/0022-1910(83)90020-3

43) Egelhaaf, 1985a, On the neuronal basis of figure-ground discrimination by relative motion in the visual system of the fly. I Behavioural constraints imposed on the neuronal network and the roles of the optomotor system, Biological Cybernetics (vol. 52, no. 2, pp. 123-140), https://doi.org/10.1007/BF00364003

44) Egelhaaf, 1985b, On the neuronal basis of figure-ground discrimination by relative motion in the visual system of the fly. II. Figure-detection cells, a new class of visual interneurones, Biological Cybernetics (vol. 52, no. 2, pp. 123- 140) https://doi.org/10.1007/BF00339948

45) Egelhaaf M & Borst A, 1992, Are there separate ON and OFF channels in fly motion vision?, Visual Neuroscience (vol. 8, pp. 151-164)

46) Egelhaaf M, Borst A, Warzecha AK, Flecks S & Wildemann A, 1993, Neural circuit tuning fly visual neurons to motion of small objects. II. Input organization of

inhibitory circuit elements revealed by electrophysiological and optical recording techniques. Journal of Neurophysiology (vol. 69, no 2. pp. 340-351), https://doi.org/10.1152/jn.1993.69.2.340

47) Eichner H, Joesch M, Schnell B, Reiff DF & Borst A, 2011, Internal Structure of the Fly Elementary Motion Detector, Neuron (vol. 70, no. 6, pp. 1155-1164), https://doi.org/10.1016/j.neuron.2011.03.028

48) Fabian JM, El Jundi B, Wiederman SD, and O'Carroll DC, The Complex Optic Lobes of Dragonflies, in preparation

49) Fain Gl, Hardie R & Laughlin SB, 2010, Phototransduction and the Evolution of Photoreceptors, Current Biology (vol. 20, no. 3, pp. R114-R124), https://doi.org/10.1016/j.cub.2009.12.006

50) Field DJ, 1987, *Relations between the statistics of natural images and the properties of cortical cells.* Journal of the Optical Society of America A, (vol. 4, pp. 2379-2394)

51) Fischbach KF & Dittrich APM, 1989, The optic lobe of Drosophila melanogaster. I. A Golgi analysis of wild-type structure, Cell and Tissue Research (vol. 258, pp. 441-475).

52) Fisher YE, Leong JCS, Sporar K, Ketkar MD, Gohl DM, Clandinin TR & Silies M, 2015, A Class of Visual Neurons with Wide-Field Properties is Required for Local Motion Detection, Current Biology (vol. 25, no. 24, pp. 3178-3189) https://doi.org/10.1016/j.cub.2015.11.018

53) Fisher YE, Silies M & Clandinin, TR, 2015, Orientation selectivity sharpens motion detection in Drosophila, Neuron (vol. 88, pp. 390-402)

54) Fotowat H & Gabbiani F, 2011, *Collision Detection as a Model for Sensory-Motor Integration*, Annual Review of Neuroscience (vol. 34, pp. 1-19)

55) Franceschini N, Hardie R, Ribi W & Kirschfield K, 1981, Sexual dimorphism in a photoreceptor, Nature (vol. 291, pp. 241-244), https://doi.org/10.1038/291241a0

56) Frye MA & Olberg RM, 1995, Visual receptive field properties of feature detecting neurons in the dragonfly, Journal of Comparative Physiology A (vol. 177, no. 5, pp. 569-576), https://doi.org/10.1007/BF00207186

57) Gauck V, Egelhaaf M & Borst A, 1998, Synapse distribution on VCH, an inhibitory, motion-sensitive interneuron in the fly visual system, Journal of Comparative Neurology (vol. 381, no. 4, pp. 489-499), https://doi.org/10.1002/(SICI)1096-9861(19970519)381:4<489::AID-CNE8>3.0.CO;2-Z

58) Geurten BRH, Nordstrom K, Sprayberry JDH, Bolzon DM & O'Carroll DC, 2007, Neural Mechanisms underlying target detection in a dragonfly centrifugal neuron, The Journal of Experimental Biology (vol. 210, pp. 3277-3284), https://doi.org/10.1242/jeb.008425

59) Gilbert C & Strausfeld NJ, 1991, The functional organization of male-specific visual neurons in flies, Journal of Comparative Physiology A (vol. 169, pp. 395-411), https://doi.org/10.1007/BF00197653

60) Gjorgjieva J, Sompolinsky H & Meister M, 2014, Benefits of Pathway Splitting in Sensory Coding, Journal of Neuroscience (vol. 34, no. 36, pp. 12127-12144), https://doi.org/10.1523/JNEUROSCI.1032-14.2014

61) Götz KG, 1964, Optomotrische Untersuchung des visuellen systems einiger Augenmutanten der Fruchtfliege Drosophila, Kybernetik (vol. 2, no. 2, pp. 77-92)

62) Gruntman E, Romani S & Reiser MB, 2018, Simple integration of fast excitation and offset, delayed inhibition computes directional selectivity in Drosophila, Natura Neuroscience (vol. 21, pp. 250-257)

63) Haag J & Borst A, 1997, Encoding of Visual Motion Information and Reliability in Spiking and Graded Potential Neurons, Journal of Neuroscience (vol. 17, no. 12, pp 4809-4819)

64) Haag J & Borst A, 2001, Recurrent Network Interactions Underlying Flow-Field Selectivity of Visual Interneurons, Journal of Neuroscience (vol. 21, no. 15, pp. 5685-5692), https://doi.org/10.1523/JNEUROSCI.21-15-05685.2001

65) Haag J, Arenz A, Serbe E, Gabbiani F & Borst A, 2016, Complementary mechanisms create direction selectivity in the fly, eLife (vol. 5, e17421) https://doi.org/10.7554/eLife.17421

66) Haag J, Mishra A & Borst A, 2017, A common directional tuning mechanism of Drosophila motion-sensing neurons in the ON and in the OFF pathway, eLife (6) e29044. http://doi.org/10.7554/eLife.29044

67) Hardie RC, 1986, The photoreceptor array of the dipteran retina, Trends in Neuroscience (vol. 9, pp. 419-423), https://doi.org/10.1016/0166-2236(86)90136-0

68) Hardie RC & Raghu P, 2001, Visual transduction in Drosophila, Nature(vol. 413, pp. 186-193), https://doi.org/10.1038/35093002

69) Haikala V, Joesch M, Borst A & Mauss AS, 2013, Optogenetic Control of Fly Optomotor Responses, The Journal of Neuroscience (vol. 33, no. 34, pp. 13927-13934).

70) Harris RA, O'Carroll DC & Laughlin SB, 1999, Adaptation and the temporal delay filter of fly motion detectors, Vision Research (vol. 39, no. 16, pp. 2603-2613) https://doi.org/10.1016/S0042-6989(98)00297-1

71) Harris, RA, O'Carroll DC & Laughlin SB, 2000, Contrast Gain Reduction in Fly Motion Adaptation, Neuron (vol. 28, pp. 595-606) https://doi.org/10.1016/S0896-6273(00)00136-7

72) Hassenstein B & Reichardt W, 1956, Systemtheoretische Analyse der Zeit-, Reihenenfolgen- und Vorzeichenauswertung bei der Bewegungspertzeption des Rüsselkäfers Chlorophanus, Zeitschrift für Naturforschung B (vol. 11, no. 9-10, pp. 513–524) https://doi.org/10.1515/znb-1956-9-1004

73) Hausen K, 1982a, Motion sensitive interneurons in the optomotor system of the fly I The Horizontal Cells: Structure and Signals, Biological Cybernetics (vol. 45, no. 2, pp. 143-156), https://doi.org/10.1007/BF00335241

74) Hausen K, 1982b, Motion sensitive interneurons in the optomotor system of the fly. II The horizontal cells: Receptive field organization and response characteristics, Biological Cybernetics (vol. 46, no. 1, pp. 67-79), https://doi.org/10.1007/BF00335352

75) Heisenberg M & Buchner E, 1977, The Role of Retinula Cell Types in Visual Behavior of Drosophila melanogaster, Journal of Comparative Physiology (vol. 117, pp. 127-162)

76) Hengstenberg R, 1982, Common Visual Response Properties of Giant Vertical Cells in the Lobula Plate of the Blowfly Calliphora, Journal of Comparative Physiology A (vol. 149, pp. 179-193).

77) Hengstenberg, 1991, Gaze control in the blowfly Calliphora: a multisensory, two-stage integration process, Seminars in Neuroscience (vol. 3, no. 1, pp. 19-29), https://doi.org/10.1016/1044-5765(91)90063-T

78) Homberg U, 2004, In search of the sky compass in the insect brain, Naturwissenshaften (vol. 91, no. 5, pp. 199-208), https://doi.org/10.1007/s00114-004-0525-9

79) Horridge GA, 1978, The Separation of Visual Axes in Apposition Compound Eyes, Philosophical Transactions of the Royal Society (vol. 285, no. 1003, pp. 1-59), http://www.jstor.org/stable/2418048

80) Horstmann W, Egelhaaf M & Warzecha AK, 2001, Synaptic interactions increase optic flow specificity, European Journal of Neuroscience (vol. 12, no.6, pp. 2157–2165), https://doi.org/10.1046/j.1460-9568.2000.00094.x

81) Horvath G, Malik P, Kriska G & Wildermuth H, 2007, Ecological traps for dragonflies in a cemetery: the attraction of Sympetrum species (Odonata: Libellulidae) by horizontally polarizing black gravestones, Freshwater Biology (vol. 52, no. 9, pp. 1700-1709), https://doi.org/10.1111/j.1365-2427.2007.01798.x

82) Hu Q & Victor JD, 2010, A set of high-order spatiotemporal stimuli that elicit motion and reverse-phi percepts, Journal of Vision (vol. 10, 9.1-9.16), https://doi.org/10.1167/10.3.9

83) Ibbotson MR & Goodman LJ - 1990 - Response characteristics of four wide-field motion-sensitive descending interneurones in apis mellifera Journal of Experimental Biology (vol. 148, pp. 255-279)

84) Ibbotson MR, 1991, Wide-field motion-sensitive neurons tuned to horizontal movement in the honeybee, Apis mellifera, Journal of Comparative Physiology A (vol. 168, no. 1, pp. 91-102), https://doi.org/10.1007/BF00217107

85) Joesch M, Schnell B, Raghu SV, Reiff DF & Borst A, 2010, ON and OFF pathways in Drosophlia motion vision, Nature (vol. 468, pp. 300-304) https://doi.org/10.1038/nature09545

86) Judge S & Rind F, 1997, The locust DCMD, a movement-detecting neurone tightly tuned to collision trajectories, Journal of Experimental Biology (vol. 200, pp. 2209-2216)

87) Juusola M, Uusitalo RO & Weckström M, 1995, Transfer of Graded Potentials at the Photoreceptor Synapse, Journal of General Physiology (vol. 105, pp. 117-148) https://doi.org/10.1085/jgp.105.1.117

88) Kathman ND, Kesavan M & Ritzmann RE, 2014, Encoding wide-field motion and direction in the central complex of the cockroach Blaberus discoidalis, Journal of Experimental Biology (vol. 217, pp. 4079-4090) https://doi.org/10.1242/jeb.112391

89) Katsov AY & Clandinin TR, 2008, Motion Processing Streams in Drosophila are Behaviourally Specialized, Neuron (vol. 59, no. 2, pp. 322-335) https://doi.org/10.1016/j.neuron.2008.05.022

90) Keles MF & Frye MA, 2017, Object-Detecting Neurons in Drosophila, Current Biology (vol. 27, no. 5, pp. 680-687), https://doi.org/10.1016/j.cub.2017.01.012

91) Kim AJ, Fitzgerald JK & Maimon G, 2015, *Cellular evidence for efference copy in Drosophila visuomotor processing*, Nature Neuroscience (vol. 18, no. 9, pp. 1247-1255), https://doi.org/10.1038/nn.4083

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

93) Krapp HG, Hengstenberg B & Hengstenberg R, 1998, *Dendritic Structure and Receptive-Field Organization of Optic Flow Processing Interneurons in the Fly*, Journal of Neurophysiology (vol. 79, no. 4, pp. 1902-1917), https://doi.org/10.1152/jn.1998.79.4.1902

94) Krapp HG, Hengstenberg R & Egelhaaf M, 2001, *Binocular Contributions to Optic Flow Processing in the Fly Visual System*, Journal of Neurophysiology (vol. 85, no. 2, pp. 724-734), https://doi.org/10.1152/jn.2001.85.2.724

95) Land MF & Collett TS, 1973, Chasing behaviour of houseflies (Fannia canicularis), Journal of comparative physiology (vol. 89, no. 4, pp. 331-357), https://doi.org/10.1007/BF00695351

96) Land MF, 1989, Variations in the Structure and Design of Compound Eyes. In: Stavenga D.G., Hardie R.C. (eds) Facets of Vision. Springer, Berlin, Heidelberg

97) Land MF, 1997, Visual Acuity in Insects, Annual Review of Entomology (vol. 42, pp. 147-177), https://doi.org/10.1146/annurev.ento.42.1.147

98) Laughlin SB, 1974, Neural integration in the first optic neuropile of dragonflies. III. The transfer of angular information. Journal of Comparative Physiology A Neuroethology, Sensory, Neural, Behavioural Physiology (vol. 92, pp. 377–396), https://doi.org/10.1007/BF00694708

99) Laughlin SB, 1976, The Sensitivities of Dragonfly Photoreceptors and the Voltage Gain of Transduction, Journal of Comparative Physiology (vol. 111, no. 3, pp. 221-247), https://doi.org/10.1007/BF00606466

100) Laughlin SB & McGinness S, 1978, The Structures of Dorsal and Ventral Regions of a Dragonfly Retina, Cell and Tissues Research (vol. 188, pp. 427-447).

101) Laughlin SB & Lillywhite PG, 1982, Intrinsic noise in locust photoreceptors, The Journal of Physiology (vol. 332, no. 1, pp. 25-45), https://doi.org/10.1113/jphysiol.1982.sp014398

102) Laughlin SB & Osorio D, 1989, Mechanisms for Neural Signal Enhancement in the Blowfly Compound Eye, Journal of Experimental Biology (vol. 144, pp. 113-146)

103) Li JL, Lindemann JP & Egelhaaf M, 2017, Location motion adaptation enhances the representation of spatial structure at EMD arrays, PLoS Computational Biology (vol. 13, no. 12, e1005919), https://doi.org/10.1371/journal.pcbi.1005919

104) Longden KD & Krapp HG, 2009, *State-Dependent Performance of Optic-Flow Processing Interneurons*, Journal of Neurophysiology (vol. 102, no. 6, pp. 3606-3618), https://doi.org/10.1152/jn.00395.2009

105) Longden KD, Muzzu T, Cook DJ, Schultz SR & Krapp HG, 2014, *Nutritional State Modulates the Neural Processing of Visual Motion* (vol. 24, no. 8, pp. 890-895), https://doi.org/10.1016/j.cub.2014.03.005

106) Maddess T & Laughlin SB, 1985, Adaptation of the motion-sensitive neuron H1 is generated locally and governed by contrast frequency, Proceeding of the Royal Society Biological Sciences (vol. 225, no. 1239, pp. 251-275).

107) Maddess T, 1986, Afterimage-like effects in the motion-sensitive neuron H1, Proceedings of the Royal Society Biological Sciences (vol. 228, pp. 433-459), https://doi.org/10.1098/rspb.1986.0062

108) Maisak MS, Haag J, Ammer G, Serbe E, Meier M, Leonhardt A, Schilling T, Bahl A, Rubin GM, Nern A, Dickson BJ, Reiff DF, Hopp E & Borst A, 2013, A directional tuning map of Drosophila elementary motion detectors, Nature (vol. 500, pp. 212-216), https://doi.org/10.1038/nature12320

109) Mauss AS, Meier M, Serbe E & Borst A, 2014, Optogenetic and Pharmacologic Dissection of Feedforward Inhibition in Drosophila Motion Vision, Journal of Neuroscience (vol. 34, no. 6, pp. 2254-2263), https://doi.org/10.1523/JNEUROSCI.3938-13.2014

110) Mauss AS, Pankova K, Arenz A, Nern A, Rubin GM & Borst A, 2015, Neural Circuit to Integrate Opposing Motions in the Visual Field, Cell (vol. 162, no. 2, pp. 351-362), https://doi.org/10.1016/j.cell.2015.06.035

111) Meinertzhagen IA & O'Neil SD, 1991, Synaptic Organization of Columnar Elements in the Lamina of the Wild Type in Drosophila melanogaster, The Journal of Comparative Physiology (vol. 305, pp. 232-263).

112) Mertes M, Dittmar L, Egelhaaf M & Boeddeker N, 2014, Visual motion-sensitive neurons in the bumblebee brain convey information about landmarks during a

navigational task, Frontiers of Behavioural Neuroscience: https://doi.org/10.3389/fnbeh.2014.00335

113) Mischiati M, Lin HT, Herold P, Imler E, Olberg R & Leonardo A, 2015, Internal models direct dragonfly interception steering, Nature (vol. 517, pp. 333-338), https://doi.org10.1038/nature14045

114) Mongeau JM & Frye MA, 2017, Drosophila Spatiotemporally Integrates Visual Signals to Control Saccades, Current Biology (vol. 27, no. 19, pp. 2901-2914.e2), https://doi.org/10.1016/j.cub.2017.08.035

115) Nalbach G & Hengstenberg R, 1994, *The halteres of the blowfly Calliphora*, Journal of Comparative Physiology A (vol. 175, no. 6, pp. 695-708), https://doi.org/10.1007/BF00191842

116) Nordström K, Barnett PD & O'Carroll DC, 2006, Insect Detection of Small Targets Moving in Visual Clutter, PLoS Biology (vol. 4, no. 3, e54), https://doi.org/10.1371/journal.pbio.0040054

117) Nordström K & O'Carroll DC, 2006, Small object detection neurons in female hoverflies, Proceedings of the Royal Society B (vol. 273, pp. 1211-1216), https://doi.org/10.1098/rspb.2005.3424

118) Nordstrom K & O'Carroll DC, 2009, The motion after-effect: local and global contributions to contrast sensitivity, Proceedings of the Royal Society B (vol. 276, pp. 1545-1554) https://doi.org/10.1098/rspb.2008.1932

119) Nordström K, Bolzon DM & O'Carroll DC, 2011, Spatial facilitation by a high-performance dragonfly target-detecting neuron, Biology Letters (vol. 7, no. 4, pp. 588-592), https://doi.org/10.1098/rsbl.2010.1152

120) Nordström K, Moyer de Miguel I & O'Carroll DC, 2011, Rapid contrast gain reduction following motion adaptation, Journal of Experimental Biology, (vol. 214, pp. 4000-4009), https://doi.org/10.1242/jeb.057539

121) O'Carroll DC, 1993, Feature-detecting neurons in dragonflies, Nature (vol. 362, pp. 541-543), https://doi.org/10.1038/362541a0

122) O'Carroll DC, Bidwell NJ, Laughlin SB & Warrant EJ, 1996, Insect motion detectors matched to visual ecology, Nature (vol. 382, pp. 63-66), https://doi.org/10.1038/382063a0

123) O'Carroll DC, Laughlin SB, Bidwell NJ & Harris RA, 1997, Spatio-Temporal Properties of Motion Detectors Matched to Low Image Velocities in Hovering

Insects, Vision Research (vol. 37, no. 23, pp. 3427-3439), https:/doi.org/10.1016/S0042-6989(97)00170-3

124) O'Carroll DC, Barnett PD & Nordstrom K, 2011, Local and global responses of insect motion detectors to the spatial structure of natural scenes, Journal of Vision (vol. 11, no. 20.) https://doi.org/10.1167/11.14.20

125) O'Carroll DC & Wiederman SD, 2014, Contrast sensitivity and the detection of moving patterns and features, Philosophical Transactions of The Royal Society B (369: 20130043) https://doi.org/10.1098/rstb.2013.0043

126) O'Shea M & Williams JLD, 1974, The anatomy and output connection of a locust visual interneurone; the lobular giant movement detector (LGMD) neurone, Journal of comparative physiology (vol. 91, no. 3, pp. 257-266), https://doi.org/10.1007/BF00698057

127) Olberg RM, 1981, Object- and Self-Movement Detectors in the Ventral Nerve Cord of the Dragonfly, Journal of Comparative Physiology (vol. 141, pp. 327-334), https://doi.org/10.1007/BF00609935

128) Olberg RM, Worthington AH & Venator KR, 2000, Prey pursuit and interception in dragonflies, Journal of Comparative Physiology A (vol. 186, no. 2, pp. 155-162), https://doi.org/10.1007/s003590050015

129) Olberg RM, Worthington AH, Fox JL, Bessette CE & Loosemore MP, 2005, Prey size selection and distance estimation in foraging adult dragonflies, Journal of Comparative Physiology A (vol. 191, no. 9, pp. 791-797), https://doi.org/10.1007/s00359-005-0002-8

130) Olberg RM, Seaman RC, Coats MI & Henry AF, 2007, Eye movements and target fixation during dragonfly prey-interception flights, Journal of Comparative Physiology A (vol. 193, no. 7, pp. 685-693), https://doi.org/10.1007/s00359-007-0223-0

131) Paulk AC, Phillips-Portillo J, Dacks, AM, Fellous JM & Gronenberg W, 2008, The Processing of Color, Motion, and Stimulus Timing Are Anatomically Segregated in the Bumblebee Brain, The Journal of Neuroscience (vol. 28, no. 25, pp. 6319-6332), https://doi.org/10.1523/JNEUROSCI.1196-08.2008

132) Pfeiffer K & Homberg U, 2014, Organization and functional roles of the central complex in the insect brain, Annual Review of Entomology (vol. 59, pp. 165-184). https://doi.org/10.1146/annurev-ento-011613-162031

133) Pichaud F, Briscoe A & Desplan C, 1999, Evolution of color vision, Current Opinion in Neurobiology (vol. 9, no. 5, pp. 622-627).

134) Quenzer T & Zanker JM, 1991, *Visual detection of paradoxical motion in flies*, Journal of Comparative Physiology A (vol. 169, pp. 331-340), https://doi.org/10.1007/BF00206997

135) Ratliff CP, Borghuis BG, Kao YH, Sterling P & Balasubramanium V, 2010, Retina is structured to process an excess of darkness in natural scenes, Proceedings of the National Academy of Sciences (vol. 107, no. 40, pp. 17368017373), https://doi.org/10.1073/pnas.1005846107

136) Rauschecker JP, von Grunau MW & Poulin C, 1987, Centrifugal organization of direction preferences in the cat's lateral suprasylvian visual cortex and its relation to flow field processing, Journal of Neuroscience (vol. 7, no. 4, pp. 943-958), https://doi.org/10.1523/JNEUROSCI.07-04-00943.1987

137) Reichardt W & Wenking H, 1969, Optical detection and fixation of objects by fixed flying flies, Naturwissenschaften (vol. 56, pp. 424–425), https://doi.org/10.1007/BF00593644

138) Reichardt W, Poggio T & Hausen K, 1983, Figure-ground discrimination by relative movement in the visual system of the fly: Part II: Towards the neural circuitry, Biological Cybernetics (vol. 46, supp 1, pp. 1-30), https://doi.org/10.1007/BF00595226

139) Reiff DF, Plett J, Mank M, Griesbeck O & Borst A, 2010, Visualizing retinotopic half-wave rectified input to the motion detection circuitry of Drosophila, Nature Neuroscience (vol. 13, pp. 973-978), https://doi.org/10.1038/nn.2595

140) Reiser MB & Dickinson MH, 2010, Drosophila fly straight by fixating objects in the face of expanding optic flow, Journal of Experimental Biology (vol. 213, pp. 1771-1781), https://doi.org/10.1242/jeb.035147

141) Remsburg AJ, Olson AC & Samways MJ, 2008, Shade Alone Reduces Adult Dragonfly (Odonata: Libellulidae) Abundance, Journal of Insect Behaviour (vol. 21, no. 6, pp. 460-468), https://doi.org/10.1007/s10905-008-9138-z

142) Reynolds JH, Pasternak T & Desimone R, 2000, Attention Increases Sensitivity of V4 Neurons, Neuron (vol. 26, no. 3, pp. 703-714), https://doi.org/10.1016/S0896-6273(00)81206-4

143) Riehle A & Franceschini N, 1984, *Motion detection in flies: Parametric control over ON-OFF pathways*, Experimental Brain Research (vol. 54, no. 2, pp. 390-394), https://doi.org/10.1007/BF00236243

144) Rigosi E, Wiederman SD & O'Carroll DC, 2017, Visual acuity of the honey bee retina and the limits for feature detection, Scientific Reports (vol. 7, 45972), https://doi.org/10.1038/srep45972

145) Rind FC & Simmons PJ, 1992, Orthopteran DCMD neuron: a reevaluation of responses to moving objects. I. Selective responses to approaching objects, Journal of Neurophysiology (vol. 68, no. 5, pp. 1654-1666), https://doi.org/10.1152/jn.1992.68.5.1654

146) Rister J, Pauls D, Schnell B, Ting CY, Lee CH, Sinakevitch I, Morante J, Strausfeld NJ, Ito K & Heisenberg M, 2007, Dissection of the Peripheral Motion Channel in the Visual System of Drosophila melanogaster, Neuron (vol. 56, pp. 155-170) https://doi.org/10.1016/j.neuron.2007.09.014

147) Rister J & Desplan C, 2011, The retinal mosaics of opsin expression in invertebrates and vertebrates, Developmental Neurobiology (vol. 71, no. 12, pp. 1212-1226), https://doi.org/10.1002/dneu.20905

148) Ritzmann RE, Ridgel AL & Pollack AJ, 2008, Multi-unit recording of antennal mechano-sensitive units in the central complex of the cockroach Blaberus discoidalis, Journal of Comparative Physiology, Neuroethology, Sensory, Neural and Behavioural Physiology (vol. 194, no. 4, pp. 341-360) https://doi.org/10.1007/s00359-007-0310-2

149) Robertson RM & Johnson AG, Collision Avoidance of Flying Locusts: Steering Torques and Behaviour, Journal of Experimental Biology (vol. 183, pp. 35-60)

150) Sauseng M, Pabst MA & Kral K, 2003, The dragonfly Libellula quadrimaculata (Odonata: Libellulidae) makes optimal use of the dorsal fovea of the compound eyes during perching, European Journal of Entomology (vol. 100, no. 4, pp. 475-479), https://doi.org/10.14411/eje.2003.071

151) Schnell B, Joesch M, Forstner F, Raghu SV, Otsuna H, Ito K, Borst A & Reiff DF, 2010, Processing of horizontal optic flow in three visual interneurons of the Drosophila brain, Journal of Neurophysiology (vol. 103, no. 3, pp. 1646-1657), https://doi.org/10.1152/jn.00950.2009

152) Schnell B, Raghu SV, Nern A & Borst A, 2012, Columnar cells necessary for motion responses of wide-field visual interneurons in Drosophila, Journal of

Computational Physiology A (vol. 198, pp. 389-395), https://doi.org/10.1007/s00359-012-0716-3

153) Schoenlein RW, Peteanu LA, Mathies RA & Shank CV, 1991, The first step in vision: femtosecond isomerization of rhodopsin (vol. 254, no. 5030, pp. 412-415)

154) Schuling FH, Mastebroek HAK, Bult R & Lenting BPM, 1989, *Properties of elementary movement detectors in the fly Calliphora Erythrocephala*, Journal of Comparative Physiology A (vol. 165, no. 2, pp. 179-192), https://doi.org/10.1007/BF00619192

155) Schwegmann A, Lindemann JP & Egelhaaf M, 2014, Depth information in natural environments derived from optic flow by insect motion detection system: a model analysis, Frontiers in Computation Neuroscience (8:83) https://doi.org/10.3389/fncom.2014.00083

156) Schwyn DA, Heras FJH, Bolliger G, Parsons MM, Krapp HG & Tanaka RJ, 2011, *Interplay between Feedback and Feedforward Control in Fly Gaze Stabilization*, International Federation of Automatic Control Proceedings (vol. 44, no. 1, pp. 9674-9679), https://doi.org/10.3182/20110828-6-IT-1002.03809

157) Sherman A & Dickinson MH, 2003, A comparison of visual and haltere-mediated equilibrium reflexes in the fruit fly Drosophila melanogaster, Journal of Experimental Biology (vol. 206, pp. 295-302), https://doi.org/10.1242/jeb.00075

158) Shoemaker PA, O'Carroll DC & Straw AD, 2005, Velocity constancy and models for wide-field visual motion detection in insects, Biological Cybernetics (vol. 93, no. 4, pp. 275-287), https://doi.org/10.1007/s00422-005-0007-y

159) Silies M, Gohl Dm, Fisher YE, Freifeld L, Clark DA & Clandinin TR, 2013, Modular Use of Peripheral Input Channels Tunes Motion-Detecting Circuitry, Neuron (vol. 79, no. 1, pp. 111-127), https://doi.org/10.1016/j.neuron.2013.04.029

160) Spalthoff C, Egelhaaf M, Tinnefeld P & Kurtz R, 2010, Localized direction selective responses in the dendrites of visual interneurons of the fly, BMC Biology (8: 36), https://doi.org/10.1186/1741-7007-8-36

161) Stavenga DG, 1995, Insect retinal pigments: Spectral characteristics and physiological functions. Progress in Retinal and Eye Research (vol. 15, no. 1, pp. 231-259).

162) Strausfeld NJ, 1971, The organization of the insect visual system (Light microscopy) I. Projections and arrangements of neurons in the lamina ganglionaris of Diptera,

Zeitschrift für Zellforschung und Mikroskopische Anatomie (vol. 121, no. 3, pp. 377-441)

163) Strausfeld NJ & Campos-Ortega JA, 1973, L3 the 3rd 2nd Order Neuron of the 1st Visual Ganglion in the “Neural Superposition” Eye of Musca domestica, Zeitschrift für Zellforschung und Mikroskopische Anatomie (vol. 139, no. 3, pp. 397-403). https://doi.org/10.1007/BF00306593

164) Strausfeld NJ & Campos-Ortega JA, 1977, Vision in Insects: Pathways Possibly Underlying Neural Adaptation and Lateral Inhibition, Science (vol. 195, no. 4281, pp. 894-897),

165) Strausfeld NJ & Lee JK, 1991, Neuronal basis for parallel visual processing in the fly, Visual Neuroscience (vol. 7, no. 1-2, pp. 13-33) https://doi.org/10.1017/S0952523800010919

166) Strausfeld, 2005, The evolution of crustacean and insect optic lobes and the origins of chiasmata, Arthropod Structure & Development (vol. 34, no. 3, pp. 235-256), https://doi.org/10.1016/j.asd.2005.04.001

167) Straw AD, Warrant EJ & O'Carroll DC, 2006, A 'bright zone' in male hoverfly (Eristalis tenax) eyes and associated faster motion detection and increased contrast sensitivity, Journal of Experimental Biology (vol. 209, pp. 4339-4354), https://doi.org/10.1242/jeb.02517

168) Straw AD, Rainsford T, O’Carroll DC, 2008, Contrast sensitivity of insect motion detectors to natural images, Journal of Vision (vol. 8, no. 32, pp. 1-9), https://doi.org/10.1167/8.3.32

169) Strother JA, Nern A & Reiser MB, 2014, Direct Observation of ON and OFF Pathways in the Drosophila Visual System, Current Biology (vol. 24, no. 9, pp. 976-983), https://doi.org/10.1016/j.cub.2014.03.017

170) Suver MP, Mamiya A & Dickinson MH, Octopamine Neurons Mediate Flight-Induced Modulation of Visual Processing in Drosophila, Current Biology (vol. 22, no. 24, pp. 2294-2302)

171) Switzer PV & Eason PK, 2000, Proximate Constraints on Intruder Detection in the Dragonfly Perithemi tenera (Odonata: Libellulidae): Effects of Angle of Approach and Background, Annals of the Entomological Society of America (vol. 93, no. 2, pp. 333-339), https://doi.org/10.1603/0013-8746(2000)093[0333:PCOIDI]2.0.CO;2

172) Takemura SY, Lu Z & Meinertzhagen IA, 2008, Synaptic circuits of the Drosophila optic lobe: The input terminals to the medulla, The Journal of Comparative Neurology (vol. 509, no. 5, pp. 493-513), https://doi.org/10.1002/cne.21757

173) Takemura SY, Bharioke A, Lu Z, Nern A, Vitaladevuni S, Rivlin PK, Katz WT, Olbris DJ, Plaza SM, Winston P, Zhao T, Horne JA, Fetter RD, Takemura S, Blazek K, Chang LA, Ogundeyi O, Saunders MA, Shapiro V, Sigmund C, Rubin GM, Scheffer LK, Meinertzhagen IA & Chklovskii DB, 2013, A visual motion detection circuit suggested by Drosophila connectomics, Nature (vol. 500, pp. 175-181), https://doi.org/10.1038/nature12450

174) Takemura SY, Nern A, Chklovskii DB, Scheffer LK, Rubin GM & Meinertzhagen IA, 2017, The comprehensive connectome of a neural substrate for 'ON' motion detection in Drosophila, eLife(6:e24394), https://doi.org/10.7554/eLife.24394

175) Tanaka K & Saito H, 1989, Analysis of motion of the visual field by direction, expansion/contraction, and rotation cells clustered in the dorsal part of the medial superior temporal area of the macaque monkey, Journal of Neurophysiology (vol. 62, no. 2, pp. 626-641), https://doi.org/10.1152/jn.1989.62.3.626

176) Tatler B, O'Carroll DC & Laughlin SB, 2000, Temperature and temporal resolving power of fly photoreceptors, Journal of Comparative Physiology A (vol. 186, no. 4, pp. 399-407), https://doi.org/10.1007/s003590050439

177) Theobald JC, Duistermars BJ, Ringach DL & Frye MA, 2008, Flies see second-order motion, Current Biology (vol. 18, no. 11, pp. 464-465), https://doi.org/10.1016/j.cub.2008.03.050

178) Theobald JC, Warrant EJ & O'Carroll DC, 2010, Wide-field motion tuning in nocturnal hawkmoths, Proceedings of the Royal Society Biological Sciences (vol. 277, no. 1683, pp. 853-860)

179) Tolhurst DF, Tadmor Y & Chao T, 1992, *Amplitude spectra of natural images*, Ophthalmology and Physiological Optics (vol. 12, pp. 229-232)

180) Treue S & Maunsell JHR, 1996, Attentional modulation of visual motion processing in cortical areas MT and MST, Nature (vol. 382, pp. 539-541), https://doi.org/10.1038/382539a0

181) Trischler C, Boeddeker N & Egelhaaf M, 2007, Characterisation of a blowfly male-specific neuron using behaviourally generated visual stimuli, Journal of Comparative Physiology A (vol. 193, no. 5, pp. 559-572), https://doi.org/10.1007/s00359-007-0212-3

182) Tuthill JC, Nern A, Holtz SL, Rubin GM & Reiser MB, 2013, Contributions of the 12 Neuron Classes in the Fly Lamina to Motion Vision, Neuron (vol. 79, no. 1, pp 128-140), https://doi.org/10.1016/j.neuron.2013.05.024

183) Tuthill JC, Nern A, Rubin GM & Reiser MB, 2014, Wide-field Feedback Neurons Dynamically tune Early Visual Processing, Neuron (vol. 82, no. 4, pp. 887-895), https://doi.org/10.1016/j.neuron.2014.04.023

184) van Hateren JH & Schilstra C, 1999, Blowfly flight and optic flow. II. Head movements during flight, Journal of Experiemental Biology (vol. 202, pp. 1491-1500)

185) van Hateren JH & Snippe HP, 2001, Information theoretical evaluation of parametric models of gain control in blowfly photoreceptor cells, Vision Research (vol. 41, pp. 1851-1865)

186) von Frisch K, 1967 The Dance Language and Orientation of Bees. Cambridge, MA: Harvard Univ. Press.

187) Wang YC & Frost BJ, 1992, Time to collision is signalled by neurons in the nucleus rotundus of pigeons, Nature (vol. 356, pp. 236-238), https://doi.org/10.1038/356236a0

188) Warzecha AK, Egelhaaf M & Borst A, 1993, Neural circuit tuning fly visual interneurons to motion of small objects. I. Dissection of the circuit by pharmacological and photoinactivation techniques, Journal of Neurophysiology (vol. 69, no. 2, pp. 329-339) https://doi.org/10.1152/jn.1993.69.2.329

189) Wicklein & Varju - 1999 - Visual system of the european humming bird hawkmoth Macroglossum stellatarum (sphingidae, lipidoptera): Motion-sensitive interneurons of the lobula plate, Journal of Comparative Neurology (vol. 408, no. 2, pp. 272-282). https://doi.org/10.1002/(SICI)1096-9861(19990531)408:2<272::AID-CNE8>3.0.CO;2-9

190) Wicklein M & Strausfeld NJ, 2000, Organization and significance of neurons that detect change of visual depth in the hawk moth Manduca sexta, Journal of Comparative Neurology (vol. 424, no. 2, pp. 356-376).

191) Wiederman SD, Shoemaker PA & O'Carroll DC, 2008, A model for the detection of moving targets in visual clutter inspired by insect physiology. PLoS One (3:e2784), https://doi.org/10.1371/journal.pone.0002784

192) Wiederman SD & O'Carroll DC, 2011, Discrimination of Features in Natural Scenes by a Dragonfly Neuron, Journal of Neuroscience (vol. 31, no. 19, pp. 7141-7144)

193) Wiederman SD & O'Carroll DC, 2013, Selective Attention in an Insect Visual Neuron, Current Biology (vol. 23, pp. 156-161), http://dx.doi.org/10.1016/j.cub.2012.11.048

194) Wiederman SD, Shoemaker PA & O'Carroll DC, 2013, Correlation between OFF and ON Channels Underlies Dark Target Selectivity in an Insect Visual Systems, The Journal of Neuroscience (vol. 33, no. 32, pp. 13225-13232) https://10.1523/JNEUROSCI.1277-13.2013

195) Wiederman SD, Fabian JM, Dunbier JR & O'Carroll DC, 2017, A predictive focus of gain modulation encodes target trajectories in insect vision, eLife, (vol. 6, e26478), https://doi.org/10.7554/eLife.26478

196) Wildermuth H, 1998, Dragonflies Recognize the Water of Rendezvous and Oviposition Sites by Horizontally Polarized Light: A Behavioural Field Test, Naturwissenshaften (vol 85, pp. 297-302)

197) Yang EC & Osorio D, 1991, Spectral sensitivities of photoreceptors and lamina monopolar cells in the dragonfly Hemicordulia Tau, Journal of Comparative Physiology A (vol. 169, no. 6, pp. 663-669).

198) Zettler F & Järvilehto M, 1971, Lateral Inhibition in an Insect Eye, Zeitschrift für vergleichende Physiologie (vol. 76, no. 3, pp. 233-244)

199) Zheng L, de Polavieja GG, Wolfram V, Asyali MH, Hardie RC & Juusola M, 2006, Feedback Network Controls Photoreceptor Output at the Layer of First Visual Synapses in Drosophila, (Journal of General Physiology (vol. 127, no. 5, pp. 495) https://doi.org/10.1085/jgp.200509470

3 Differential Adaptation to Visual Motion Allows Robust Encoding of Optic Flow in the Dragonfly

3.1 Preamble

Much of science represents happy accidents. It is not uncommon for discoveries to be made on the journey to another destination, a scientific detour if you will. In this sense, my research of the wide-field motion pathways in dragonflies was one such detour. Originally a spin-off of another student's work, I discovered I possessed a particular knack for obtaining wide-field neurons. As there was no published literature on this account, I thought it a waste to ignore data staring me right in the face and thus continued to collect data.

What started as a simple story of characterising a set of neurons found in numerous other insect species turned into a complex story all of its own. Following on from work done by Joseph Fabian examining the anatomical structure of the dragonfly lobula, dragonfly anatomy proved to be quite different from other species. Dragonflies possess both a lobula plate and a sub-lobula with each neuropil being associated with wide-field motion detection in different species. Anatomical investigation represents only a small component of my research, but the few examples I obtained were sufficient to demonstrate a very complex underlying system. Still the question remained, which lobe did dragonfly wide-field neurons arborize in and what might that mean? For evolutionary enthusiasts, this remains an interesting question given the dragonfly's significant age as a species.

The story became even more complex with the collection of what appeared to be highly inconsistent data. Most other species have quite consistent physiology across their widefield systems. Dragonflies seem to be an exception to the trend.

Spatiotemporal tuning curves (a basic characterising stimuli) did not seem to line up with one another across cells. This was in part due to the presence of three different cell subtypes all exhibiting broadly similar responses but differing markedly in their adaptation properties. This made traditional approaches to analysis ineffective and required a new way of visualizing the data to fully capture the individual uniqueness of each subclass of WFMS neuron. These differences also manifest themselves anatomically (though this is an ongoing story). Unlike other insects, dragonfly wide-field neurons can arborize in both lobula plate and sub-lobula, or simply the sub-lobula alone (no evidence exists yet for lobula plate-only

wide-field neurons in dragonflies). These structural differences appear to correlate (there are only three successful stains) with the functional differences elicited from physiology.

Apart from the reconstruction of a single wide-field neuron presented in the proto-paper below (thanks Joseph), the remaining work here is solely of my own collection and analysis. I would also like to thank Sam Polacek whose project inspired me to continue down this pathway.

Statement of Authorship

Title of Paper	Differential Adaptation to Visual Motion Allows Robust Encoding of Optic Flow in the Dragonfly
Publication Status	Submitted to BioRxiv
Publication Details	Evans BJE, O'Carroll DC, Fabian JM & Wiederman SD, unpublished

Principle Author

Name of Principal Author	Bernard Evans		
Contribution to the Paper	Experiment conceptualisation, animal collection, data collection, data analysis, data interpretation, figure generation, manuscript authorship.		
Overall percentage (%)	60%		
Certification	This paper reports on original research I conducted during the period of my Higher Degree by Research candidature and is not subject to any obligations or contractual agreements with a third party that would constrain its inclusion in this thesis. I am the primary author of this paper.		
Signature		Date	14/12/2018

Co-Author Contributions

By signing the Statement of Authorship, each author certifies that:

i. the candidate's stated contribution to the publication is accurate (as detailed above);
ii. permission is granted for the candidate in include the publication in the thesis; and
iii. the sum of all co-author contributions is equal to 100% less the candidate's stated contribution.

Name of Co-Author	David C O'Carroll		
Contribution to the Paper	Experiment conceptualisation, data interpretation, manuscript authorship		
Signature		Date	4/12/2018

Name of Co-Author	Joseph M Fabian		
Contribution to the Paper	Animal collection, data collection, immunohistochemistry processing, neuron imaging, figure generation (anatomy), manuscript evaluation		
Signature		Date	28/11/18

Name of Co-Author	Steven D Wiederman		
Contribution to the Paper	Experiment conceptualisation, data interpretation, manuscript evaluation		
Signature		Date	14/12/18

3.2 Introduction

Flying insects live in complex and varied 3-dimensional environments and display diverse flight behaviour, from near stationary hovering, to territorial patrolling and rapid pursuits of prey or conspecifics. This diversity places conflicting demands on the neuronal networks underlying self-motion detection. Neurons that respond robustly to patterns of wide-field motion have been extensively studied in several insect groups, including Dipteran flies (Hausen 1982, Hausen & Egelhaaf 1989), moths (Wicklein & Varju 1999, Theobald et al., 2010, Stöckl et al., 2016) and bees (DeVoe et al., 1982, Ibbotson 1991, Mertes et al., 2014). Typified by Lobula Plate Tangential Cells (LPTCs) of Dipteran flies, these neurons take input from local elementary motion detection (EMDs) elements located in the medulla (Borst et al 2010) and employ local correlation of spatially separated inputs with asymmetric delay mechanisms, consistent with influential computational motion models (Hassenstein & Reichardt 1956, Barlow & Levick 1965, Gruntman et al., 2018). Such neurons are tuned to specific spatial and temporal frequency ranges by their underlying spatial sampling and temporal delay filters. Because this places fundamental limitations on the velocity range of motion that neurons can individually encode, insects have evolved strategies for motion analysis that match their distinctive behaviour. For example, diurnal and nocturnal hawkmoths are precise hoverers when flower feeding and use wide field motion sensitive neurons specialized for such slow velocities (O'Carroll et al., 1996, 1997, Wicklein & Varju 1999, Theobald et al., 2010, Stöckl et al., 2016, 2017). By contrast, fast flying butterflies and bees show tuning to higher image speeds (Ibbotson 1991, O'Carroll et al., 1996).

In Dipteran flies, the conflicting demands of diverse flight modes that involve switches between slow hovering to high speed pursuit flight are in part met by multimodal integration of fast input to descending visual pathways from the ocelli (Parsons et al., 2006) and specialized hindwing mechanosensory organs (halteres) that detect rapid accelerations, allowing compound eye neurons (the LPTCs) to focus on slower motion (Hengstenberg 1991).

Although dragonflies have recently emerged as an important model for studying visual target tracking, very little is known about their neural tuning to wide-field motion. Dragonflies exhibit a similar behavioural repertoire to Dipterans, but have a lower wingbeat frequency, and lack specialized halteres for detecting gyroscopic forces. As an essentially visual creature, how do dragonflies encode the large velocity ranges demanded by their behavior? One potential strategy is to process the same retinal input using parallel pathways employing

spatiotemporal filters tuned to different speed ranges, as seen in mammals (Movshon & Newsome 1996, Nassi & Callaway 2008). In many insects, however, replicating such parallel pathways may be constrained by their size and weight. Indeed, in species studied to date, motion tuning at a behavioral level appears to reflect a single common EMD mechanism (Buchner 1976). Nevertheless, we hypothesize that parallel processing may be viable for dragonflies, which possess among the largest eyes and brain of extant insects. Alternatively, useful coding of different speed ranges may result from additional downstream processing. Motion adaptation, for example, can improve velocity contrast via relief from saturation (Maddess & Laughlin 1985; Barnett et al., 2010) and on a timescale similar to the stimulus response (Nordström et al., 2011). It also improves velocity encoding of natural images (Shoemaker et al., 2005, Straw et al., 2008, Barnett et al., 2010) and enhances differentiation between foreground and background features (Li et al., 2017).

We tested these two alternative strategies by recording from widefield motion-sensitive neurons in the lobula of dragonflies. We found evidence for several unique subclasses of widefield motion-sensitive neurons, some of which differ substantially from their counterparts in other species. We found evidence that these neurons likely share common input pathways (i.e. using the same EMD inputs) but differ radically in their adaptation to image motion. This differential motion-adaptation tunes otherwise similar neurons to significantly different velocity ranges, providing very robust encoding of motion over several decades of image speed.

3.3 Materials and Methods

3.3.1 Electrophysiology.

93 wild-caught, dragonflies (*Hemicordulid*) were immobilized with a 1:1 beeswax and rosin mixture, with the head tilted forward to access the posterior surface. A hole was cut above the brain to gain access to the lobula and lateral midbrain, but the preparation was otherwise left with the perineural sheath and overlying haemolymph sacs intact. We penetrated the sheath and recorded intracellularly using strong aluminosilicate micropipettes, pulled on a Sutter Instruments P-97 puller and backfilled either with KCl (2M, electrode tip resistance typically 50-150 MΩ) or 4% Lucifer Yellow solution in 0.1M LiCl.

3.3.2 Visual Stimuli

We presented stimuli on high definition LCD monitors (120 -165 Hz). The animal was placed 20 cm away and centred on the visual midline. Contrast stimuli were presented at screen

centre to minimize off-axis artefacts. The display was corrected for distortions using OpenGL to ensure each 1° onscreen was 1° from the animal's perspective. The visual field was 104° by 58.5° (azimuth and elevation respectively). All temporal frequencies tested were limited to one quarter of the monitors frame rate. Stimulus scripts were written using MATLAB's Psychtoolbox and integrated into the data acquisition system.

To classify neurons as widefield motion sensitive, a sequence of characterising stimuli were presented to the dragonfly. These included a gyrated, randomly generated texel pattern (1°), grey to black and grey to white full screen flicker (White – 338 cd/m^2, Black 0.5 cd/m^2), moving edges (up, down, left and right, 25°/s), moving bars (2° width, up, down, left and right, 25°/s) and a square-wave grating pattern moving up, down left and right (0.025 cycles/°, 6.25Hz). Neurons were categorised as widefield motion sensitive based on robust responses to the gyrated texel pattern and square-wave gratings. Following, sinusoidal gratings were presented to dragonflies which had a linear increase in contrast for 1 s (0 to 0.25, Weber) followed by a 1 s exponential rise (0.25 to 1).

3.3.3 Neuroanatomy

The morphology of a widefield motion-sensitive neuron was visualized by intracellular labelling with Lucifer Yellow (Figure 22a, b, c). Iontophoresis was achieved by passing 1nA negative current through electrodes tip-filled with Lucifer Yellow for 12 minutes. Brains were then carefully dissected, fixed overnight in 4% paraformaldehyde at 4°C, dehydrated in ethanol series (70%, 90%, 100%, 100%), cleared in methyl salicylate and mounted using Permount on a slide using three spacer rings and covered with a cover slip for imaging. The sample was scanned using a confocal microscope using a 10x objective and the 3D slices reconstructed using Neutube.

3.3.4 Experimental Design and Statistical Analysis

All analysis was completed in MATLAB. Spike-counting was done using a custom-written spike-counting script. Curve fits used MATLABs in-built curve-fitting tools. To find peaks of tuning curves, repeated measures were averaged followed by the application of a 5-point moving average filter to smooth data before finding the maximum. All statistical tests were either paired-sample non-parametric tests (Wilcoxon sign test, for paired data) or two-sample non-parametric tests (Mann-Whitney U test, for unpaired data) with appropriate multiple-comparison corrections (Bonferroni). All means are calculated from biological replicates (i.e. repeated measurements from identified neurons in different animals). Each biological

replicate represents the mean of between 1 and 5 technical replicates. P values are reported as raw numbers in text if significant differences were found (unmarked otherwise) or as < 0.0001 if sufficiently small. Box and whisker plots represent the 75th, 50th and 25th quartiles (lines) with raw data shown.

3.4 Results

3.4.1 Neuroanatomical characterisation

While we were not able to anatomically identify every neuron that we recorded from in the midbrain/lobula complex in the dragonfly, we nevertheless stained a subset of these neurons, typified by the example in Figure 22. This confirms a similar general organisation to that in Diptera and several other insect orders (Hausen 1982, Egelhaaf 1989, Fabian 2017), where optic flow is integrated within specialised subregions of the 3rd optic ganglion (the lobula complex) by a set of *tangential* neurons, the well-studied wide-field motion sensitive 'Lobula Plate Tangential Cells' (Hausen 1982). These neurons have input dendrites that integrate tangentially across arrays of retinotopically-organised inputs from underlying local motion detectors (presumptive EMDs) at earlier stages of visual processing. Figure 22A-C shows the reconstructed morphology of a dragonfly neuron that exhibited wide-field motion sensitivity. The overall morphology of these neurons strongly resembles their Dipteran counterparts, with tree-like input arborisations within the lobula complex, and outputs in the lateral midbrain. However, as with this individual example, several neurons described in this study have their inputs originating solely from a deep neuropil on the anterior side of the lobula, similar to the 'sublobula' identified in bees (Devoe et al., 1982, Strausfeld 2005, Strausfeld et al., 2006), rather than from a posterior Lobula Plate. Until the homologies between these different lobula subregions with their counterparts in other insect groups are more clearly identified, we label these neurons more generally as 'Lobula Tangential Cells' (LTCs).

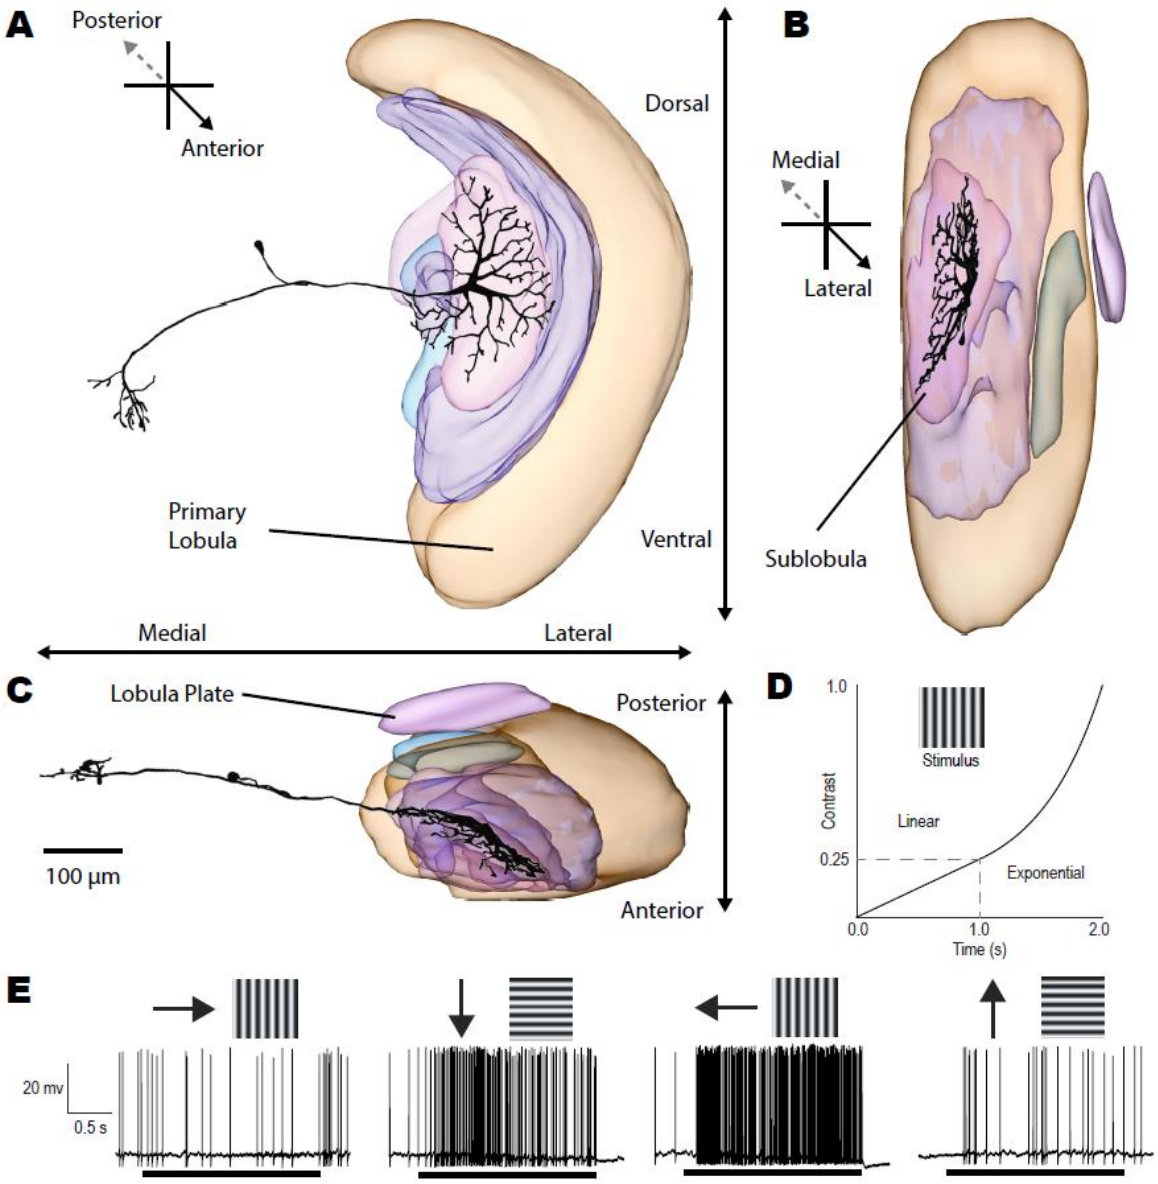

Figure 22: ***A,*** *Anterior,* ***B,*** *Longitudinal and* ***C,*** *Horizontal projections for a 3D model reconstruction of a dragonfly lobula tangential cell (LTC).* ***D,*** *description of the waveform of the contrast ramp used to characterise response tuning and contrast sensitivity in LTC neurons. At cessation of a pre-stimulus period, the contrast rises linearly for a 1 second, from 0 to a contrast of 0.25. Then for the remaining 1 second, the contrast continues to rise exponentially from 0.25 to a final contrast of 1.0.* ***E,*** *Example responses to sinusoidal gratings in four directions for a direction opponent LTC. Black bars indicate the period of the 2 second contrast ramp. Direction opponent cells are inhibited by motion in the 'anti-preferred' direction (left) while responding strongly when stimulated by the opposite direction of motion.*

3.4.2 Direction selectivity and opponency

We tested the motion sensitivity of LTCs using sinusoidal gratings drifted in eight directions (presented randomly in 45° increments). Each grating was displayed as a 'contrast ramp' with a nonlinear increase in contrast from zero over a 2 s period (Figure 22D). This stimulus

avoids onset-transients inherent with step changes in contrast. The ramp also weights more time around the threshold contrast values whilst still providing a stimulus that contains the entire contrast range (O'Carroll et al., 1996, 1997). Figure 22E shows an individual neuron's spiking response to four directions of motion. This LTC response exhibits clear direction opponency, with excitation to preferred motion and inhibition to the anti-preferred direction of motion. The response time course in these neurons typically shows high initial sensitivity to low contrast, indicated by a rapid rise in firing rate. The ramp responses show variable degrees of saturation as the contrast continues to increase and, in many cases, a subsequent decrease in response due to temporal adaptation.

We measured LTC directionality by presenting a gyrated random (binary) texture pattern (Figure 23A). This continuous stimulus is composed of a broad range of spatial and temporal frequencies and permits precise calculation of direction selectivity. The full-screen texture was gyrated (not rotated) in a clockwise and counter-clockwise direction for two complete cycles. In direction opponent neurons, this elicits a sinusoidal pattern of excitation and inhibition (Figure 23B). We thus fitted a sinusoid (Figure 23C) to the instantaneous spike rate (inverse inter-spike interval). Responses to clockwise and counter-clockwise rotations were then averaged to eliminate any phase-lag due to response latency. The maximum and minimum values of the fitted sinusoid were taken as the neuron's preferred (R_p) and anti-preferred (R_a) responses. We then defined a direction opponency metric relative to the spontaneous activity level (R_s) as:

$$\text{Direction Opponency (DO)} = 1 - (R_a - R_s)/(R_p - R_s)$$

We used this basic stimulus to characterise the preferred direction selectivity in 93 neurons (recorded from the dragonfly lobula complex) that gave strong responses to wide-field motion (Figure 23D). These opponent LTCs exhibit both vertical and horizontal preferred directions, clustering around all four cardinal directions (i.e. left, right, up and down). A similar alignment of neuronal sensitivity to the different directional components of ego-motion are also observed in the frontal visual fields of LPTCs in other insect species (Strausfeld & Lee 1991, Krapp & Hengstenberg 1996).

This initial neuron selection was likely biased by repeated recordings from this stereotyped location in the optic lobes where we had previously located direction selective cells. Hence the resulting distribution of DO values probably underestimates the number of neurons that are not direction selective. Nevertheless, while the recorded neurons displayed a large range of DO (from near zero to 2.8), a histogram of DO reveals a clear peak in the distribution

above 1.0, indicating that most of the neurons were strongly direction opponent (Figure 23E). Strong direction selectivity and opponency is a characteristic of many of the LPTC neurons seen in other taxa, so the role of the weakly directional neurons from this group in optical flow analysis is unclear at this stage. We therefore limit our subsequent analysis in this paper to a subset of recorded neurons with strong direction selectivity (DO>0.75), which are more likely to be analogues or homologues of the LPTC neurons in other species (Figure 23E).

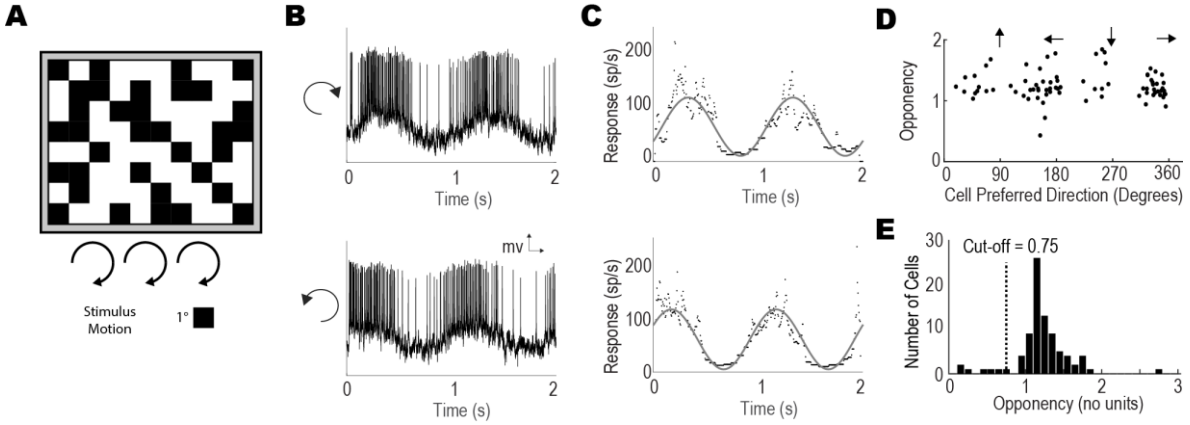

*Figure 23: **A** Moving texel patterns (1° texels) at a constant orientation were presented to the dragonfly whilst recording from widefield motion sensitive neurons. The pattern was gyrated twice in a clockwise circle (not rotated) and then twice anti-clockwise at 50°/s. **B.** An example spike trace during texel pattern stimulus shows a clear periodic response corresponding to the texel pattern direction of motion. **C.** The inverse interspike interval (ISI) reveals spiking activity to the stimulus, which is fitted with a sinusoidal curve for both a clockwise (top) and counter-clockwise (bottom) texel gyration. This curve fit provides both the peak response phase (i.e. preferred direction) and response characteristic at the preferred and anti-preferred direction establishing the neuron's directionality **D.** Each point indicates the neuron's preferred direction and DO. Dragonfly LTCs exhibit preference for four directions (progressive, regressive, upward, downward). **E.** A histogram showing the direction selectivity (in total neurons) of the widefield-motion sensitive neurons.*

3.4.3 Spatial and Temporal Tuning

In many animal models, motion sensitive neurons exhibit distinctive tuning to both the spatial and temporal frequency of drifted gratings (Hausen 1982, Devoe et al., 1982, Arenz et al., 2017). Here we established LTC spatiotemporal tuning by presenting two series of sinusoidal gratings. To establish spatial tuning, we presented a randomised series of 30 gratings of varying spatial frequencies logarithmically spaced between 0.01 and 1 cycles/°, using a fixed temporal frequency (5 Hz). To establish temporal tuning, we presented a series with varying temporal frequencies logarithmically spaced between 0.1 and 30 Hz, using a fixed spatial frequency (0.1 cpd). In each case the grating stimulus was a ramp of contrast as previously described (Figure 22D).

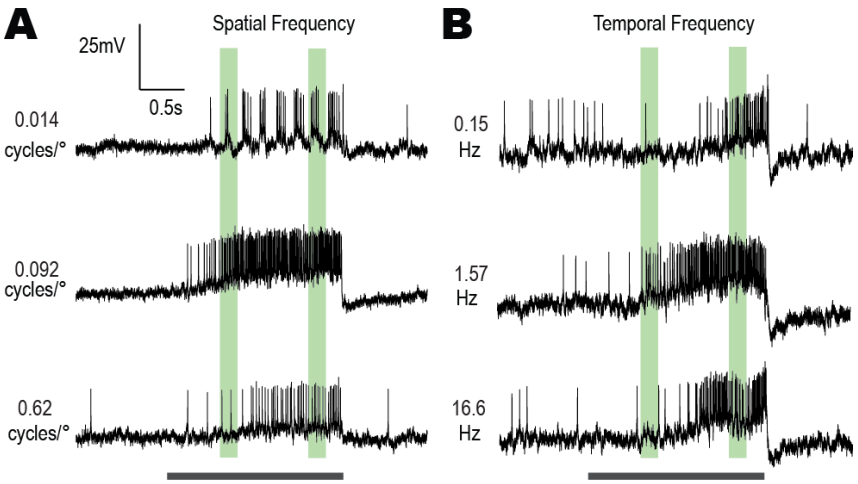

Figure 24: Responses from a neuron that exemplifies a subset of LTCs that show strong consistency in their spatial and temporal tuning over time (and thus increasing contrast). The period of the contrast ramp stimulus is indicated by the black bar. Green bars represent early and late window periods used for later analysis. ***A,*** *raw responses to 3 different spatial frequencies, 0.014, 0.092 and 0.62 cycles/° using gratings with a temporal frequency of 5 Hz.* ***B,*** *corresponding raw responses to 3 different temporal frequencies (0.15, 1.57 and 16.6 Hz) at a constant spatial frequency of 0.1 cycles/°.*

Figure 24 shows data for a single neuron that exemplifies a subset of LTCs that show consistency in their spatial and temporal tuning across time, and thus to increasing contrast as the ramp stimulus progresses. This particular neuron gave mixed mode responses, with spikes that ride on graded depolarization when the stimulus was excitatory (Figure 24A). Separate quantitative analysis of such mixed-mode responses revealed general consistency between the graded and spiking responses across a wide range of stimulus conditions. Since other neurons showed only biphasic (axonal) action potentials (e.g. Figure 22) we thus limit our subsequent analysis to the spiking component of the activity. At very low spatial frequencies, responses are often phase-locked to the 1st harmonic of the stimulus waveform (i.e. the original 5 Hz frequency), particularly at high contrast (e.g. Figure 24A top trace). This phase-locking is not evident, however at higher spatial frequencies. Stimulus conditions that elicit the strongest responses during presentation of the grating also lead to a strong rebound response on motion cessation, i.e. a motion after effect (Anstisa et al., 1998; Nordstrom et al., 2009) very evident in the graded response, but also in a reduction in spike firing rate in the post-stimulus period.

3.4.4 Temporal Adaptation and Neuron Classification

Do the direction opponent LTCs show evidence of more than one class? 3-dimensional plots of different LTC responses over time show remarkably different characteristics, particularly in the temporal frequency domain (Figure 25A-C, note colour is the difference between the stimulus induced response and spontaneous activity). Such plots reveal responses over a

broader range of frequencies as time progresses, reflecting recruitment of activity for less-optimal stimuli as grating contrast increases. This initial broadening was observed in all neurons recorded, although we found that they varied widely both in their initial contrast thresholds and in their subsequent responses due to differences in temporal adaptation. To account for these differences and yet still compare initial (weakly adapted) with later motion-adapted responses we also derived tuning curves (Figure 25D-F) based on two 200 ms duration analysis windows (as indicated in Figure 24). The *early* window (black line) starts when responses exceed the spontaneous activity by 2 standard deviations to an optimal stimulus. A *late* window (red line) begins 250 ms prior to the end of the ramp stimulus.

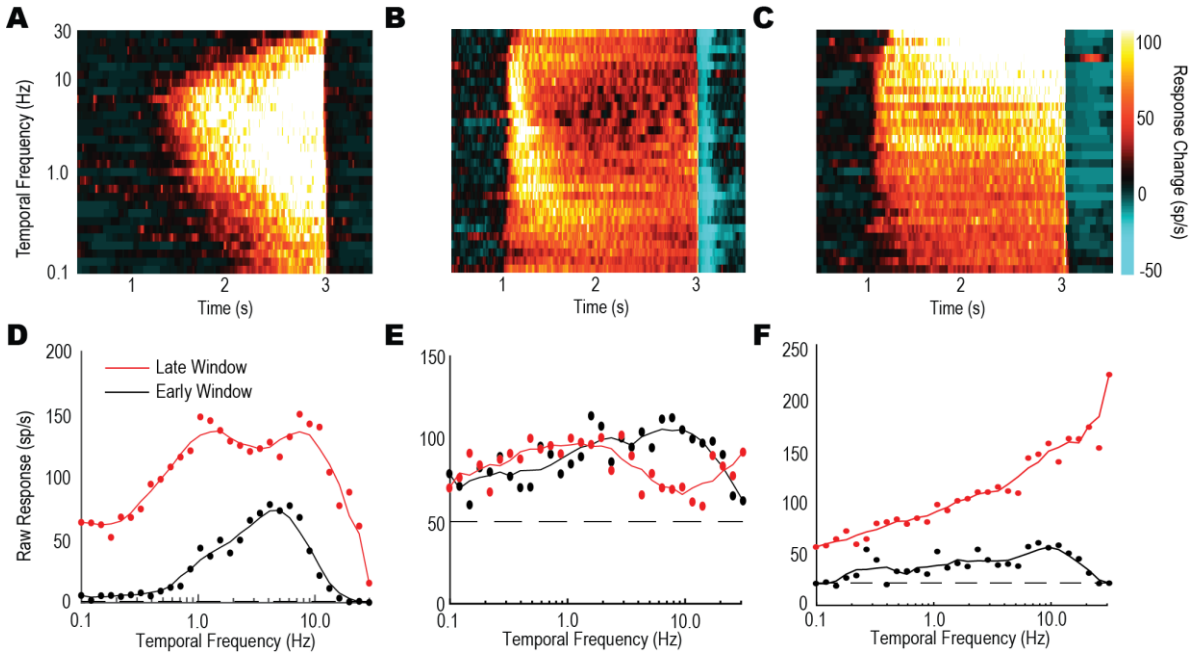

*Figure 25: **A-C,** 3-dimensional plots of the spike rate above spontaneous activity over time in response to contrast ramps at 30 different temporal frequencies (at a spatial frequency of 0.1 cycles/°) for examples of 3 different putative neuron sub-classes; (**A** slow adapting tangential cells (SATC), **B** selective fast adapting tangential cells (SFATC), **C** fast adapting tangential cells (FATC). Each plot shows the spiking activity during the 2 second contrast-ramp gratings (colour coded for the inverse interspike interval i.e. spike/s). **D-F** Temporal tuning curves derived by averaging responses from plots as in **A-C** within short (200ms) windows either early in the ramp (low contrast, black line) or late (high contrast, red line).*

In the first example LTC (Figure 25A, D), responses continue to increase with higher contrast, saturating across a range of temporal frequencies for the near-optimal spatial frequency used here (0.1 cycles/°), but with no obvious shift in the location of the optimum. Consistency of spatial and temporal optima, and robustness of the basic shape with respect to adaptation state of neurons (as in Figure 25A, D) is also observed in *Dipteran* LPTCs (Hausen 1982, Harris et al., 1999). However, not all LTCs exhibited such LPTC-like tuning

properties, however. Figure 25B and Figure 25C show data for two additional neurons, highlighting large variation in the evolution of responses over time for different temporal frequencies (though all still exhibiting direction selective responses to wide-field motion). We first defined a subclass of neurons like that in Figure 25A and Figure 25D, showing little change in their tuning properties over time. We termed these *slow adapting tangential cells* (SATCs). Qualitatively, the other neurons (Figure 25B, C) exhibited much stronger motion-adaptation, though in different ways. A second subclass exhibit very strong motion adaptation at their initial preferred temporal frequency, so that even high contrasts at the end of the ramp at such frequencies elicit only weak responses (Figure 25B, E). This gives rise to a distinctive 'notch' in the centre of the late window temporal tuning (Figure 25E) which is also evident from the dark region in the 3 dimensional plot (Figure 25B). Despite this adaptation at the initial optimum frequency, these neurons retain strong responses at both higher and lower temporal frequencies. We refer to these as *selective fast adapting tangential cells* (SFATCs) to account for the selective nature of the motion adaptation by intermediate frequencies. A third group of neurons (Figure 25C, F) strongly adapted to low temporal frequency stimuli, shifting their most robust responses to higher temporal frequencies over time, reaching the limits of frequencies possible with our display. We termed this subclass *fast adapting tangential cells* (FATCs).

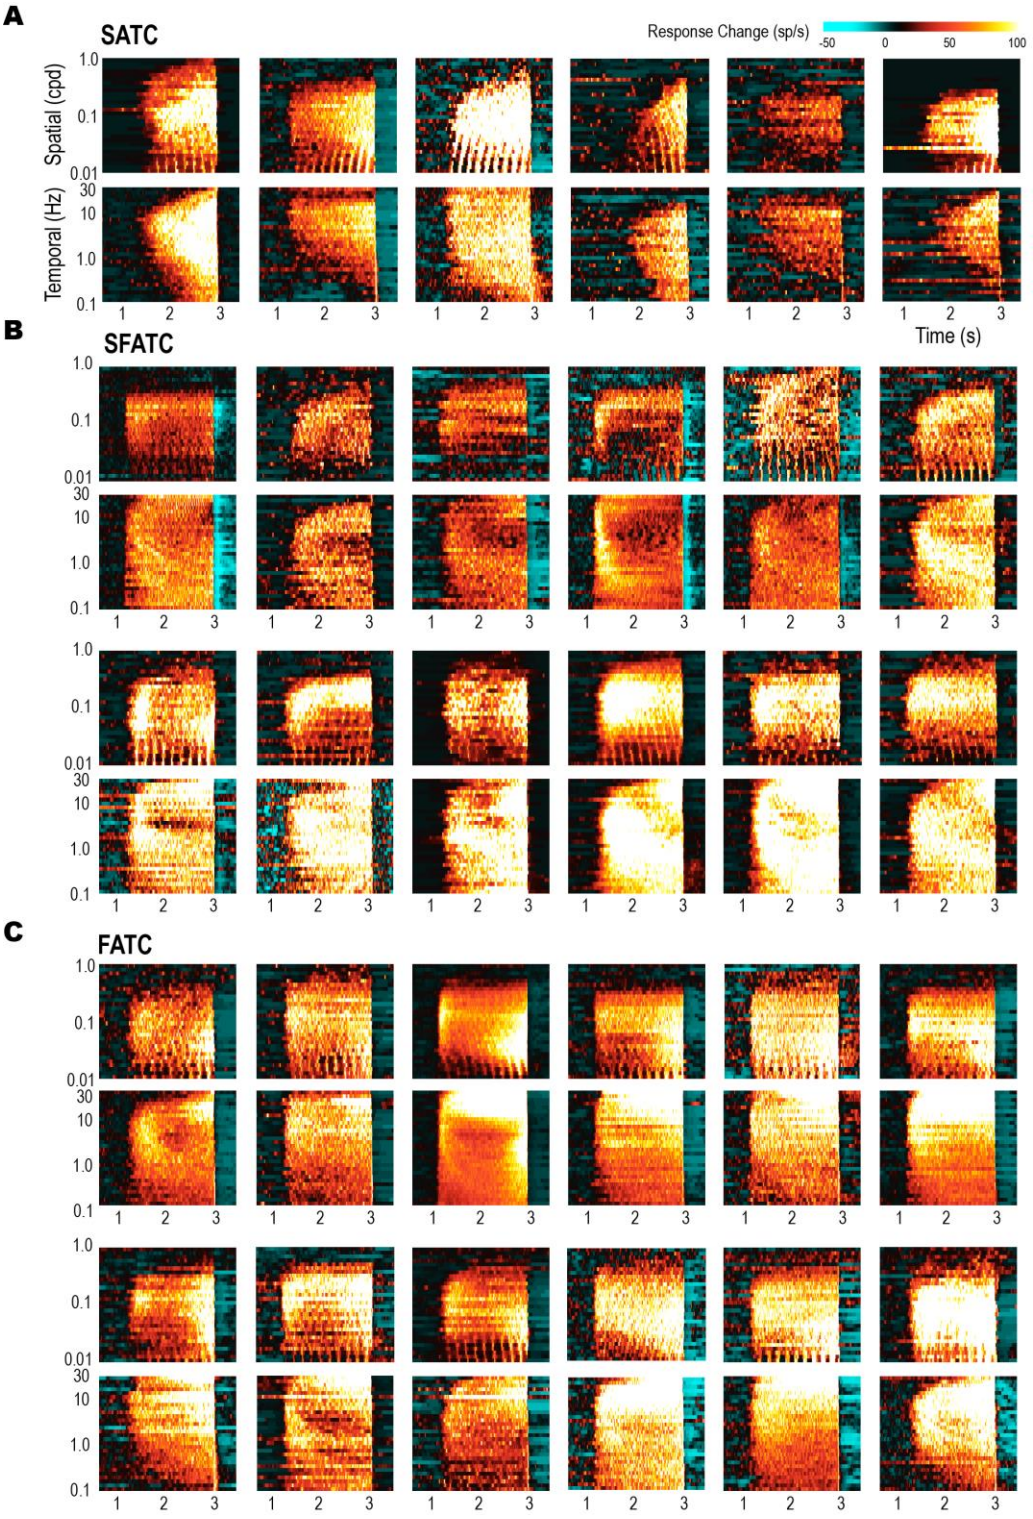

Figure 26: 3-dimensional plots of individual LTCs forming putative subclasses. The change in spike rate (spontaneous is subtracted) over time is plotted in response to contrast ramps at 30 different spatial and temporal frequencies. Plots are arranged in pairs for each individual neuron, with responses to varying spatial frequencies (temporal frequency of 5Hz) arranged above that for temporal frequencies (spatial frequency of 0.1 cycles/°). Each plot shows the spike rate during the 2 s stimulus duration (contrast-ramp gratings) which includes a 1 s pre-stimulus and 1 s post stimulus period. ***A****, Slow adapting tangential cells (SATCs, 6 examples shown) exhibit LPTC-like tuning, that broadens over time (increasing contrast) with little*

*shift in the optimal tuning **B**, Selective fast adapting tangential cells (SFATCs, 12 examples shown) with their readily identifiable adapted 'notch' forming at the earlier optimal tuning. **C,** Fast adapting tangential cells (FATCs, 12 examples shown) revealing strong shifts in spatial and temporal frequency tuning over time.*

3.4.5 Classification of Subclasses of LTC

Whilst the differences between individual neurons as illustrated in Figure 25 is observed in the examples selected, formalizing this subdivision across our population of recorded neurons based on quantitative measures is more difficult. When we plotted 3-dimensional plots of the spike rate over time in response to contrast ramps at 30 different temporal and spatial frequencies for a larger group of neurons, the patterns described above for the different subclasses was obvious in many, but not all cells (Figure 26). Each panel in this figure uses the same colour lookup table to indicate spike rate change from spontaneous, so overall responsiveness can be compared across neurons. In a subset of neurons in which we were able to record mixed mode responses, we also examined the graded response component (data not shown). We found these to be qualitatively similar to the spiking data, so our subsequent analysis is limited to spiking responses. For each individual cell we have provided panels for the spatial and temporal tuning over time (paired vertically).

The initial division into subclasses is supported by consistent and characteristic differences in the pattern of motion adaptation, as described for the examples provided in Figure 25. Most SATCs (Figure 26A) show little motion adaptation, responding more strongly at any given frequency as pattern contrast increases and thus to a broader range of frequencies as time progresses. SFATCs and FATCs (Figure 26B, C) both show more complex time courses, with darker coloured areas often following the recruitment of responses at the initial optima in different bands of both spatial and temporal frequency.

The SFATC and FATC subclasses have more complex interactions between motion adaptation (leading to weaker responses) and increasing contrast (leading to stronger responses) during the contrast ramp. This diverse group may ultimately prove to represent a spectrum of response characteristics. Nevertheless, to permit comparison of other response parameters between them, we sorted all recorded neurons into the 3 groups as follows. First, after visual inspection of the late versus early windowed data, neurons with clearly 2-peaked adapted temporal tuning (i.e. the 'adaptation notch' identified in Figure 25E) were considered to be SFATCs. For the remaining 42 neurons, which all exhibited a single peak in temporal frequency tuning in both the early and late windows, we quantified the shift in optimum as a fold change from the original value described in the log-domain (thus 0 represents no change

and 1 a ten-fold change). The distribution of these temporal frequency 'peak change' values is shown in Figure 27A. From Figure 27A it is difficult to see a clear cut-off point between the two cell subtypes and it is likely that the two populations overlap using this metric. Despite this, by choosing an arbitrary threshold (0.32 log units, ~2.07-fold change), the two subtypes were well separated when compared qualitatively (SATC green bars, FATC red bars). Such shifts in temporal frequency optima over time (and therefore increasing contrast) is a novel observation in insect, wide-field motion sensitive neurons, with responses observed in Dipteran LPTCs more like those classified here as SATCs (Figure 27A, green bars, < 0.32 log peak change).

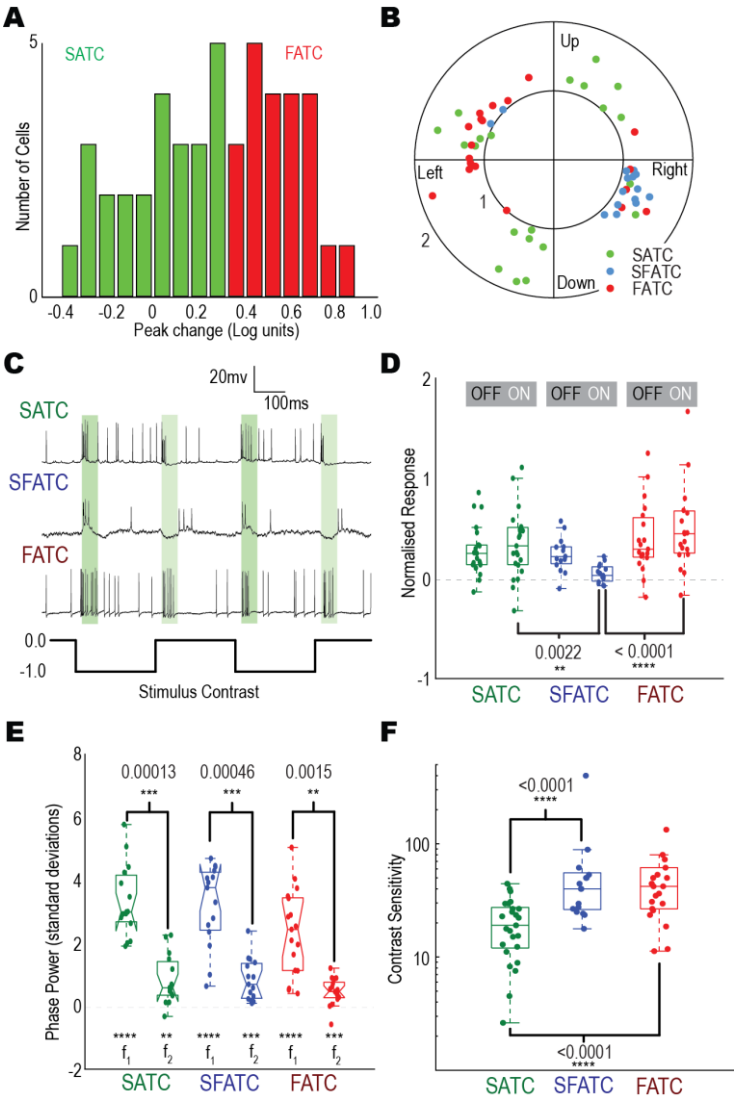

Figure 27: ***A,*** *Histogram depicting the logarithmic shift in temporal frequency optimum in the group of LTC neurons identified as putative SATCs and FATCs.* ***B,*** *Polar plot of direction tuning (angle; across up down, left and right) values plotted against direction opponency (magnitude, where DO>0.75, see Figure 23). FATC and SFATC cells appear to code horizontal motion directions only, while SATCs include individual neurons with peak responses towards any of the 4 cardinal directions.* ***C,*** *Raw traces from each subclass of LTC*

when stimulated by full-screen flicker (grey to black, 0.0 to -0.99 Weber contrast). Green bars indicate window periods (20 to 70 ms) for analysis. These individual SATC and FATC cells exhibit full-wave rectification (increased responses to both ON and OFF increments) while SFATCs were inhibited to ON flicker. ***D,*** *Boxplot distributions of windowed responses to full screen flicker (0.73 and -0.99 Weber contrast), following either luminance increments (ON) or decrements (OFF). SFATC cells show no response to ON flicker, while FATC and SATCs generally show responses to both ON and OFF components.* ***E,*** *boxplot distributions of the fundamental frequency (f_1) and second harmonic (f_2) responses estimated in the late part of ramp responses at low spatial frequencies (mean of results from 0.137-0.260 cycles/°), at a temporal frequency of 5 Hz. All 3 classes of neurons show significant modulation at both f_1 and f_2* ***F,*** *Boxplot distributions of contrast sensitivity (to a drifting grating) based on the contrast required to evoke a neuronal response (2x the standard deviation of the spontaneous activity). SFATC & FATC cells exhibit significantly higher peak contrast sensitivity than SATCs (Mann-Whitney U test, n=65).*

3.4.6 Differences in response sensitivity and tuning across subclasses of LTC

Does our classification of these cells into several subclasses (based primarily on differences in their temporal adaptation properties) correlate with other physiological response attributes? Figure 27B shows the direction tuning data as defined earlier (Figure 23) replotted for the different subclasses of LTCs on polar axes. Neurons selective for horizontal motion include examples from all 3 subclasses, but interestingly all the LTCs with vertical preferred directions (i.e. sensitive for upwards or downwards motion) fell into the SATC subclass. Given their frontal receptive fields and response characteristics resemble those of Dipteran LPTCs, these horizontal and vertical sensitive SATC neurons may indeed be the dragonfly equivalent of the VS (vertical system) and HS (horizontal system) neurons that play an important role in analysis of pitch, roll and yaw rotations (Hausen & Egelhaaf 1989; Krapp & Hengstenberg 1996).

In our initial characterisation, we noted that all LTCs give transient responses to full-screen (square wave) flicker at low temporal frequencies (Figure 27C). Much recent work on Dipteran LPTCs supports a model that integrates inputs to local motion detectors from separate ON and OFF pathways that originate in early visual processing (Borst & Helmstaedter 2015). By contrast, our own prior work suggests that motion detection in the feature-selective pathways of this same dragonfly species can be strongly selective for the OFF pathway (Wiederman & O'Carroll 2013). Do we see any clear segregation of these flicker components in the responses of different LTC subclasses? Figure 27C shows example data traces of individual LTCs in response to a 2 Hz full-screen flicker pattern. We analysed response windows 20 to 70 ms after the onset of each ON and OFF phase (shaded green

regions) for the various LTC subclasses. Figure 27D shows these ON and OFF response components for our population of LTCs, separated into the subclasses. For both SATC (green) and FATC (red) responses are full-wave rectified, with similar spiking responses to both ON and OFF transitions. However, SFATCs responded predominantly to the OFF transition in the full-screen flicker stimulus with hyperpolarisation during ON stimuli (luminance intensity increases). This resulted in a statistically significant difference between the OFF and ON responses of SFATCs (n=15) and a statistically significant difference between SFATC ON responses and both SATC and FATC ON responses (Kruskal-Wallis with multiple comparisons, SATC n=23, SFATC n=15, FATC n=19). This result is somewhat at odds with our observation that this subclass, like the SATCs and FATCs, exhibit a clear phase-locked modulation of the response to very low spatial frequency grating pattern drifted in the preferred direction (as in Figure 24A). That is, in response to low spatial frequency motion, responses modulate at the grating's temporal frequency (1^{st} harmonic, f_1). However, the transient changes of full-screen flicker results in frequency-doubled, 'breakthrough' responses to both ON and OFF intensity changes (2^{nd} harmonic, f_2).

To interrogate these phase-locked responses, we performed a fast Fourier transform on the mean response of five of the lowest spatial frequencies tested (0.0137-0.0259 cycles/°). The frequency analysis of these low spatial-frequency motion responses (Figure 27E, SATC n=15, SFATC n=15, FATC n=18) shows that they are dominated by power at the fundamental temporal frequency of the grating (i.e. the first harmonic) with weaker but still statistically significant power (Kruskal-Wallis with multiple comparisons) at double the fundamental temporal frequency (i.e. the second harmonic), suggesting a fully rectified input (i.e. both ON and OFF responses) in all 3 subclasses of LTC.

A clear distinction between the SATC and FATC/SFATC classes is also evident from their contrast sensitivity. This is observed from the 3D plots in Figure 26, where many FATC and SFATC cells show an abrupt transition from the pre-stimulus (spontaneous) response level to strong excitation within a few hundred milliseconds of ramp onset. This indicates a very high sensitivity to very low contrast across a large range of frequencies. To quantify this sensitivity further (Figure 27F), we estimated the standard deviation of the spontaneous activity of the neuron when viewing a blank screen in the pre-stimulus period and then determined the time during the ramp at which the response exceeds 2x this value and back-calculated the contrast based on the ramp waveform. We took the inverse of this threshold value as the contrast sensitivity. This method equivalent to the detectability criterion used in

previous studies of contrast sensitivity in insect motion sensitive neurons (Dvorak et al 1980). Based on this criterion, all 3 neuron classes showed high contrast sensitivity, with peak values in the range from 20 (for SATCs) to 40 (for FATCs), similar to peak sensitivity values reported for other insect species (O'Carroll & Wiederman 2014). Nevertheless, peak contrast sensitivity was significantly higher in FATCs and SFATCs (SATC n=25, SFATC n=17, FATC n=22, Kruskal-Wallis with multiple comparisons test) that display clear motion adaptation during the ramp, suggesting that adaptation may help compensate for higher initial contrast gain to avoid saturation.

3.4.7 Spatial tuning

Figure 29 shows spatial tuning (temporal frequency = 5Hz) data averaged across all cells allocated to the SATC, SFATC and FATC subclasses. All three classes exhibit qualitatively similar spatial tuning in the early analysis window, before adaptation at higher contrasts begins to affect the curve shapes. The optimum is centred near 0.1 cycles/°, with high sensitivity extending responses to the lowest frequencies tested (0.01 cycles/°) but rolling off above 0.5 cycles/°. This roll-off is consistent with the theoretical predictions based on the interommatidial angle, which is a little below 1° for this species (Horridge 1978, Buchner 1976). In the late windows, however, higher contrast and differential adaptation as ramps progress alters the shape of spatial tuning curves among different cells in ways that reinforce our earlier segregation of the different LTC subclasses based on other parameters. Firstly, consistent with their contrast sensitivity (Figure 28E) SFATCs show stronger responses than SATCs in the early window and less response increase in the later window (i.e. as responses saturate). Apart from this saturation, the late and early window responses have similar shape in SFATCs, with little evidence of a 'notch' near the initial optimum as seen in the temporal frequency domain. Although the apparent centre of the distributions is similar for the averaged tuning data (Figure 28A) individual cells do show a weak increase in the optimum spatial frequency for the SFATC cells (Figure 28C). Interestingly, although our distinction between SATCs and FATCs was made solely on the basis of the shift in their temporal optima, the group we identified as FATCs clearly also show a difference in the shape of their spatial tuning as contrast increases, with a significant reduction in the optimum towards a 3-fold lower spatial frequency (Figure 28C, n=18, Kruskal-Wallis with multiple comparisons test).

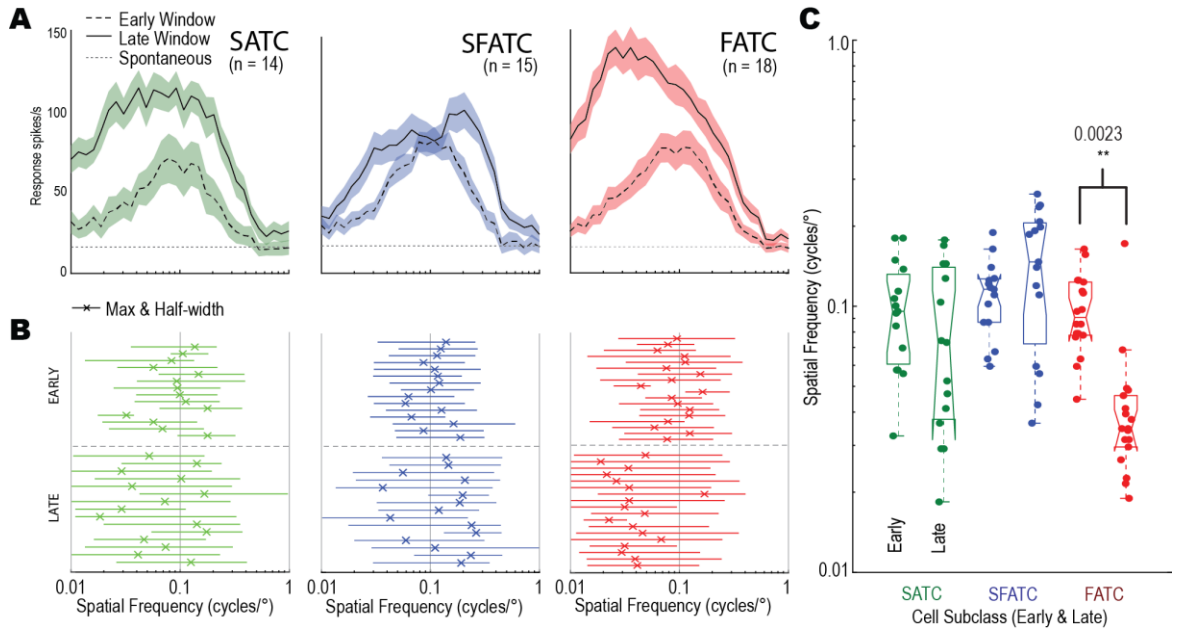

*Figure 28: Spatial tuning of different LTC subclasses. **A**, Average spatial tuning for different cell classes in two response windows (early, late) as defined in Figure 25. Each plot shows the mean response as a line with the shaded area denoting standard error. The mean spontaneous across cells is indicated by the grey dotted line. FATC cells exhibit a clear shift towards lower spatial frequencies in their late window compared with SATCs & SFATCs. **B**, Individual neuron summary data showing the position of the peak in early/late windows (denoted by an X) with the full width at half maximum (width of responses greater than 50% of maximum response) indicated by horizontal lines. **C**, Changes in spatial frequency peaks shown as box plots. FATCs show significant decrease in peak spatial tuning.*

3.4.8 Temporal Tuning

Figure 29 shows the temporal tuning data (spatial frequency = 0.1 cycles/°) averaged across all cells allocated to the SATC, SFATC and FATC subclasses, using a similar analysis as for the spatial frequency domain (Figure 28). This largely supports our qualitative observations based on individual cells. The SATC temporal frequency tuning response is broad, with an optimum in both the early and late window of approximately 6 Hz. Apart from some broadening of the late window response (Figure 29A) the shape of the tuning function is very similar in both analysis windows.

SFATCs exhibit a pronounced change in temporal frequency tuning over the two windows measured. Reflecting their high contrast sensitivity, the temporal frequency tuning is already very broad in the early window, spanning more than 2 decades of temporal frequency (Figure 29A). This broadness is even more evident in the late window, with a prominent notch at intermediate temporal frequencies. This combines with additional recruitment of stronger responses to less optimal stimuli by higher contrasts to the full range of temporal frequency tested (0.1 to 30 Hz), such that the overall tuning is remarkably flat. Indeed, the response

never falls to below 50% maximum in this range in many cells, confounding our attempt to identify a distinct temporal optimum in the late window for these cells. This made it impossible to quantify either a change in the temporal optimum in the late window or a useful measure of response half-width (Figure 29B). Interestingly, if we identify the location of the notch in individual cells (indicated by 'O' symbols in Figure 29B) this coincides closely with the peak in the early window, suggesting that it arises from adaptation recruited more strongly at what was the initially optimal stimuli.

In SATC and FATC cells by contrast, a distinct maximum is seen in the response analysed in both early and late windows. In SATCs, the late window optimum is not significantly different than in the early window. In FATCs however we see a significant (at least a 4-fold) increase in the temporal frequency optima (Figure 29C, n=22, Kruskal-Wallis with multiple comparisons test). This new temporal frequency optimum was also significantly different to the temporal frequency optimum of the other two cell subtypes (Figure 29C, SATC n=25, SFATC n=17, FATC n=22, Kruskal-Wallis with multiple comparisons test). Note that due to the frame rate of our stimulus display, we limited stimuli to below 30 Hz in order to ensure that the phase shift for gratings on successive video frames never exceeded 90 degrees. In individual FATCs, the response is still rising at this upper limit, so our data set most likely underestimates the true magnitude of this shift in tuning. It seems likely that had we used a display with infinite frame rate, the roll off in response above 30 Hz would be primarily determined by the limits of early visual processing.

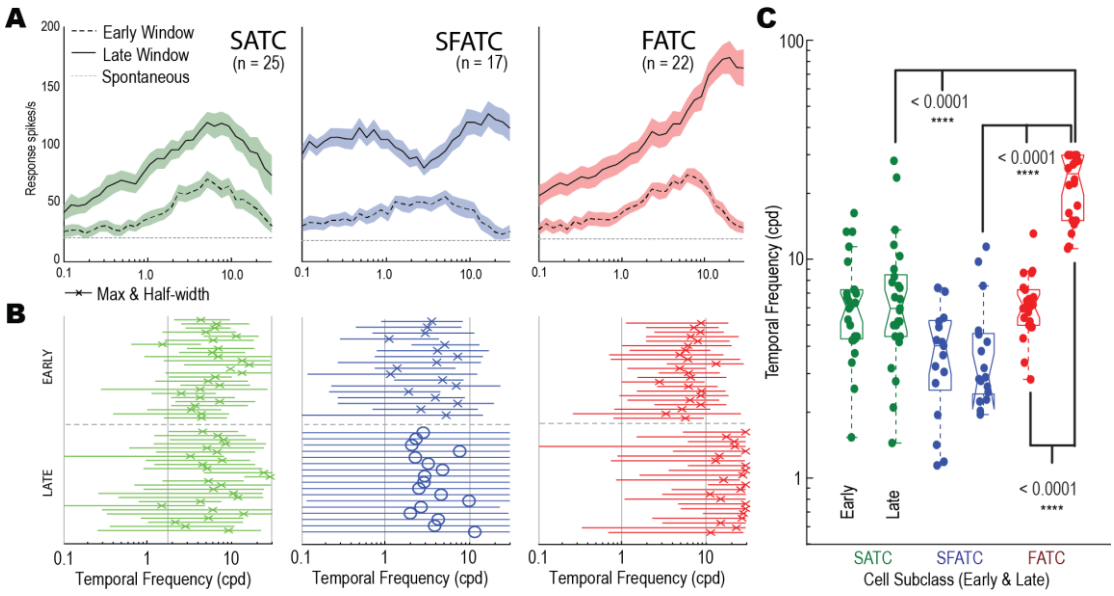

*Figure 29: Temporal tuning of different LTC subclasses. **A**, Average temporal tuning for different cell classes in two response windows (early, late) as defined in Figure 25. Each plot shows the mean response as a line with the shaded area denoting standard error. FATCs*

exhibit a clear shift towards higher temporal frequencies in their late window compared with SATCs & SFATCs. ***B****, Individual neuron summary data showing the position of the peak in early/late windows (denoted by an X) with the full width at half maximum (width of responses greater than 50% of maximum response) indicated by horizontal lines. In this case SATC cells exhibit little change in their tuning properties, FATC cells exhibit a large shift towards higher temporal frequencies and SFATC cells exhibit a characteristic adaptation at their former preferred temporal frequency. In (B), the position of the 'notch' in the adapted state is denoted by O rather than an X for the optimum for the SFATCs.* ***C,*** *Boxplot distributions for changes in temporal frequency optima. This metric cannot be defined in the late window for the SFATCs due to the notch. FATC neurons show a large increase in temporal frequency optima compared to their early window and a higher temporal frequency optimum in the late window than both SATC and SFATC cells (using the notch location as a stand in for peak).*

3.4.9 Velocity coding by LTCs for natural scenes

Our analysis of response tuning using narrow-band sinusoidal gratings suggests that all 3 LTC classes use fundamentally similar spatial and temporal filters in their underlying motion detectors, evidenced by their similar spatial and temporal tuning in the early windows. In other insects, such optima for sinusoidal patterns provide robust predictions for the velocity range over which the same neurons respond to broad-band images, including natural scenes (Dror et al., 2000, Barnett et al., 2010). However, as the contrast ramps progress, the large differences in adaptation to the stronger motion stimuli among different LTC classes may have a substantial influence on responses to moving natural patterns.

To test this we estimated velocity tuning using prolonged exposure to motion for a suite of either six or an extended set of sixteen natural image panoramas, depending on the experiment recording duration (Figure 30A). The stimulus comprised a sequence of brief periods of test motion across a wide range of velocities, interleaved with a constant adapting stimulus, but always moved continuously in the preferred direction for the neuron (Figure 30B). The adapting periods are longer than the brief test speeds (500ms versus 200ms) to ensure that the adaptation state is kept reasonably constant at the start of each test period. The test velocities cycle through an ascending and descending order to evaluate any hysteresis that may reflect differential adaptation to the test pulses themselves.

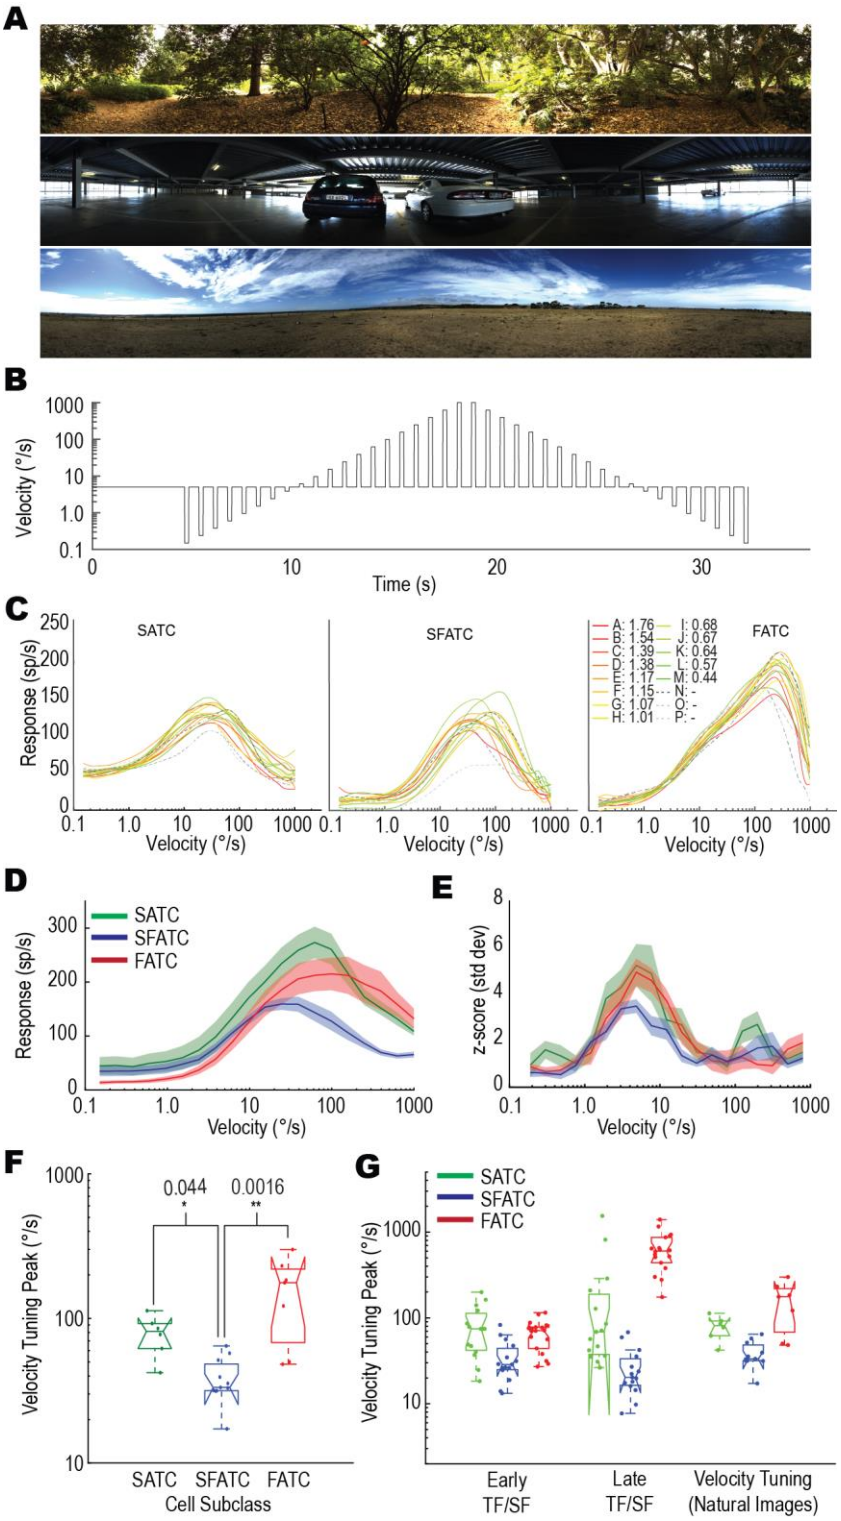

Figure 30: Velocity tuning of LTC responses to natural scenes. ***A,*** *3 examples from the set of 16 panoramic images used as stimuli.* ***B,*** *The velocity for each panoramic image is modulated over time, translated horizontally on the screen oriented in the preferred direction of the neuron. Following an initial 4 second period, brief test periods (200ms) of varying velocities were interleaved with longer (500ms) periods of the adapting speed (5°/s).* ***C,*** *Individual examples from three different cell subclasses, with each line a velocity tuning profile for a different natural image averaged across different starting phases and forward/backward sequences. Images A-P are color coded in order of their image contrast (C_{EMD}). These LTCs*

exhibit consistency between images despite large changes in image contrast. In particular, note the consistency across images over decades of velocity for the example FATC. ***D,*** *Responses of the three LTC subclasses to the test portion of the modulated velocity shown in (B) 50ms after the velocity change (100ms window, averaged across images/phases of forward and backward sequence). Each cell type exhibits a different velocity tuning profile with SATCs peaking at an intermediate velocity (62.8°/s), SFATCs peaking at lower velocities (24.5°/s) while FATCs peak at higher velocities (98.6°/s).* ***E,*** *A plot showing the z-scores (see text) at each velocity for each cell subtype (a measurement of information content. SFATC cells show the lowest information content****F****, Box plot showing the distributions of velocity optima for different LTC subtypes. P values indicate significantly lower optima in the SFATCs than the other two types (Mann-Whitney U).* ***G****, A box plot showing three estimates of velocity tuning for each of the three cell subtypes. Both early and late estimates use the tuning peaks calculated in Figure 28 and Figure 29 via the formula: Velocity = Temporal Frequency / Spatial Frequency. The final instead uses the late windows. The third shows the velocity tuning derived from natural images (Figure 30F). Velocity tuning from natural images falls between that calculated from the early and late windowed spatial and temporal tuning curves*

For each neuron, we calculated the mean response across all image phases and produced a velocity-tuning curve for each background image. An exemplar of each cell subtype with 16 different images is shown in Figure 30C. As can be seen from these individual neurons, there is close agreement in response across images indicating that like LPTCs in Diptera (Straw et al., 2008), LTCs in dragonflies exhibit remarkable velocity constancy.

For each neuron, we then averaged the responses across images and grouped subtypes. We then plotted the mean and standard error for each subtype (Figure 30D, SATC n=6, SFATC n=11, FATC n=7). The motion adaptation we earlier described (Figure 29) manifests itself as differences in the optimal velocity tuning for each of the LTC subclasses. While all three subclasses of LTC begin responding to the moving background at ~1°/s, their behaviour at fast velocities is markedly different. SFATC neurons exhibit the lowest velocity peak, occurring at 24.5 °/s. SATC neurons are next, at 62.8 °/s. FATC neurons show the broadest tuning range peaking at 98.6 °/s. The ordering is in line with our previous results of temporal frequency analysis (Figure 29) indicating that the motion adaptation seen in the grating experiments translates into different velocity-tuning optima for natural images.

We also analysed the velocity constancy (lack of variation between images) seen in Figure 30C to see if there was a difference between the three subclasses. To assess the velocity constancy we used a metric, called a z-score. The z-score captures the response variation between different velocities (i.e. the useful information) compared to the response variation at a single velocity due to the change in background image. Z-score is defined here as:

- $k(\mu_n - \mu_{n-1})/(s.d_n + s.d_{n-1})$

Where μ_i is the mean response (across images) at the i^{th} velocity and $s.d_i$ is the standard deviation at the i^{th} velocity. The change in mean response between successive velocities (i.e. the information) is divided by the variation between images for a single velocity pair. This value is then normalized by a factor (k) to account for the number of velocity samples per decade of velocity (higher resolution samples show smaller mean differences between subsequent velocity measures). This final measure (z-score) gives an indication of which velocities ranges LTCs convey the most information (i.e. their dynamic range).

Figure 30E shows that FATC and SATC neurons exhibit a much higher z-score at their peaks than SFATC cells indicating that they convey more information about velocity (potentially improving the precision of the velocity estimate for example). Though subtle, the z-scores also seem to reiterate the slight difference in peaks between FATC & SATC neurons (with FATC being faster). Figure 30E shows three individual examples of 16-image trials for a single neuron. Here we see the largest difference observed thus far. While the SFATC and SATC are qualitatively similar, the FATC neuron exhibits a significantly higher optimum velocity, around 150°/s, and a remarkably asymmetric velocity tuning function that rises almost monotonically over more than 2 decades of image velocity.

Figure 30F shows a boxplot of the peaks of the individual neurons separated by type. In agreement with Figure 30D both SATC and FATC cells have higher peak velocity than SFATC cells (SATC n=6, SFATC n=11, FATC n=7, Kruskal-Wallis with multiple comparisons). We did not find a statistically significant difference between peak of the the SATC and FATC neurons, though this is may be a consequence of insufficient power. We note the challenging nature of conducting intracellular recordings, particularly with respect to holding a neuron long enough to perform the velocity experiment across multiple images.

How do these velocity-tuning peaks relate to the temporal and spatial tuning of the neurons? To test this we calculated the predicted velocity tuning by dividing the temporal tuning peak by the spatial tuning peak (TF/SF) for both the early and late windows from Figure 28 and Figure 29 (Figure 30G). These were highly informative calculations. As can be seen from the boxplots, the velocity tuning peaks are in line with directions expected from temporal and spatial tuning peaks. The real velocity tuning data for FATC and SFATC neurons lies between the estimates derived from the early and late peaks. The SATC neurons appear fairly constant across all three methods of calculation.

	TF/SF Early (°/s)	TF/SF Late (°/s)	Natural Images (°/s)
SATC	74.5	68.8	81.3
SFATC	28.7	20.3	33.3
FATC	71.2	602.5	176.6

Table 1: *Tables showing the estimated velocity tuning peaks (median) of the three LTC subtypes. The estimates are as follows: 1, an estimate based on the early-window spatial and temporal tuning peaks, 2, an estimate based on the late-window spatial and temporal tuning peaks, 3, the peak calculated from our experiments in response to natural images of varying velocities (Figure 30)*

3.5 Discussion

3.5.1 Effect of adaptation on velocity tuning

Overall, our results show some similarities between the dragonfly LTCs and their Dipteran counterparts, but also several intriguing differences. In Dipteran LPTCs such as the H1 neuron of the blowfly for example, previous work suggests that exposure to motion at high temporal frequency (20 Hz) leads to a rapid decay in response compared with the initial level, while at 1 Hz, responses are sustained for many seconds (Jung et al., 2011). We observe precisely the opposite effect in the FATC cells of the dragonfly, with the strongest recruitment of responses by the end of the stimulus at the highest temporal frequency tested. Does the apparent shift of maximal sensitivity to very high temporal frequencies and to very low spatial frequencies that we observed towards the end of the contrast ramps for these neurons reflect an adaptation mechanism that shifts maximal sensitivity from very low speed to very high speed? Such a mechanism has previously been proposed as an explanation for motion adaptation in different species, including humans (Clifford et al 1997). Contrary to the predictions of this model, however, subsequent estimates of temporal frequency tuning before and after adaptation to a *constant* motion stimulus (typically a fast-moving pattern) showed that the time constant of the underlying delay filter responsible for temporal tuning is not altered, at least in Dipteran LPTCs. Rather, adaptation primarily reduces the response via a decrease in the contrast gain of the neuron (Harris et al 1999, 2000).

In FATCs the very weak responses at low temporal frequencies late in the ramp, but very strong responses at high frequencies certainly suggests an adaptation mechanism that is more strongly recruited by slowly changing patterns in these cells – i.e. some form of potent non-linear high pass filtering recruited primarily by higher contrasts. However, because the adapter is not constant in each trial, we cannot directly infer from our data that the delay time

constant of the motion detector itself has changed. Nevertheless, since the contrast of natural scenes is also very high and dragonflies experience prolonged exposure to motion across a large range of velocities during natural behaviour, it is interesting to consider how this complex adaptation interacts with the underlying spatiotemporal tuning to shape velocity tuning during natural image motion among the different neuron subclasses.

While the higher velocity optimum in the FATCs would at first seem to be correlated with a difference in adaptation that leads to a higher optimum in the temporal frequency tuning and corresponding lower optimum in spatial frequency in our late window ramp data, more careful evaluation of the theoretical relationship between responses to narrow band sinusoids and broad band natural scenes would suggest otherwise. If we first consider the spatial and temporal optima from the early window in ramp data around 0.1 cycles/° and 6 Hz respectively for FATCs, these would predict an optimum speed of ~60°/s for such narrow band sinusoidal patterns (Figure 30G). Natural scenes, however, have their highest contrast at low spatial frequencies, with contrast then declining at higher frequencies according to the famous $1/f^n$ characteristic (Field 1987, Tolhurst 1992). This 1/f property leads to velocity optima for such scenes a little over twice those predicted for sinusoids for the same motion detectors (Dror et al., 2000). In other words, the observed velocity optimum for the FATCs in response to natural images, at around 120°/s is actually a good match for the predictions based on the spatial and temporal tuning in the unadapted (early) windows. Indeed, both the temporal and spatial tuning optima for sinusoids and the optimum speed for natural scenes for these dragonfly neurons are similar to those observed in LPTCs of male hoverflies (Barnett et al 2010, Straw et al 2008).

Hence it is actually the very slow velocity optima of the SATC and SFATCs, around 1/3 that of the FATCs, that is surprising. These optima are inconsistent with the predictions based on their responses to low contrast sinusoids, despite the latter being similar across all 3 groups (Figure 29C, Figure 30C, Figure 30G). This suggests that motion adaptation or other non-linear processing during prolonged exposure to motion of natural scenes leads to suppression of responses to higher velocities, which otherwise ought to be a more potent stimulus for these two cell classes. For SFATCs at least, this conclusion is consistent with the appearance of the ‘notch’ that we observe in the temporal frequency tuning for initially optimal patterns. Full resolution of this issue will require extensive future work employing a rigorous test-adapt-test approach to examine how the different components of motion adaptation previously identified in dipteran LPTCs (Harris et al., 2000) are differentially recruited by

different adapting stimuli. Our initial results in this area (not shown here) indicate that the adaptation shown here is a temporal phenomenon and occurs with a constant contrast adaptor.

3.5.2 Alternative Explanations of Motion Adaptation

Previous studies have demonstrated that differences in early versus late responses to motion at different temporal frequencies can be induced by differences in the behavioural state of the animal (e.g. during tethered flight versus restrained states) (Chiappe et al., 2010; Maimon et al., 2010). Many features of this differential adaptation to prolonged stimuli can also be induced in restrained animal preparations by exogenous application of agonists for the neuromodulator octopamine, and lead to apparent shifts in the temporal tuning of Dipteran LPTCs (Longden & Krapp 2009; Jung et al., 2011; Arenz et al., 2017; Suver et al., 2012). Could differences in octopaminergic modulation of the different classes of dragonfly LTCs potentially explain some of the differences we observe in their response time courses? Although we did not attempt to address the role of octopamine directly, it seems unlikely to explain the unusual adaptation we observe in FATCs for several reasons. Firstly, it was not uncommon to record data for two or more subclasses of LTC in the same animal exhibiting very different motion adaptation. This argues strongly against global differences in octopaminergic activity as an explanation. Secondly, either increased locomotor activity or octopamine agonists appear primarily to downregulate LPTC response *reduction* at high temporal frequencies during prolonged motion exposure (Jung et al., 2011), though nutritional state has been shown to modulate these effects (Longden et al., 2014). Hence the higher gain seen at higher temporal frequency following application of octopamine agonists in dipteran LTPCs results primarily from less adaptation to motion during the response time course. By contrast in FATCs we already see sustained, vigorous responses at the highest temporal frequencies tested, despite our animals being fully restrained. Instead we see selective adaptation for lower temporal frequencies and high spatial frequencies: both are more consistent with some form of powerful redundancy reduction in the signals at the elementary motion detector inputs. Finally, we have individual recordings where the temporal tuning was repeatedly measured that lasted over four hours with no apparent change in response shape.

Strong temporal and spatial (centre surround) antagonism have both been observed in dragonfly lamina monopolar cells (Laughlin 1974). Although linear spatial filtering ought to equally affect the early response window of our ramp stimuli, it would hardly be surprising if potent antagonism – either spatial or temporal – were recruited non-linearly as contrast

increases, e.g. via additional voltage gated or inactivating conductances in feed-forward synapses. Indeed, such a mechanism may be required for a system with high contrast sensitivity (as observed here) to regulate the gain of local motion detectors and limit saturation in the real world, where average contrasts are high. Hence the differences we see between LTC subclasses may potentially arise from differences in which classes of lamina cells (or their post-synaptic targets) lie on the inputs to underlying motion detectors, an observation further supported by the differences we see in their transient responses to flicker stimuli. Note that both strong antagonism and non-linear temporal high pass filtering are key components of models proposed to explain spatiotemporal tuning of local motion detecting elements for the small target motion detector (STMD) pathway involved in target tracking in the lobula in these same dragonflies (Wiederman et al., 2008), so it is possible that the SFATC and SATC subtypes of STC take their primary inputs from the same local motion detectors as these feature-detecting neurons.

3.5.3 Velocity constancy for natural scenes

In all three LTC classes, we observed a high degree of consistency across responses to natural images, with the majority of the curves peaking at a similar optimal velocity and the gain in response to different velocities being similar despite very large differences in the global contrast among this set of test images (Barnett et al., 2010; Brinkworth and O'Carroll 2009). Such 'velocity constancy' for highly variable natural scenes in dipteran LPTCs has previously been suggested to derive from a number of dynamic non-linear processing stages in biological vision, commencing with fast temporal adaptation in the photoreceptors and second order neurons, but also with a strong contribution from dynamic gain control within the local motion detectors (Shoemaker et al., 2005; Brinkworth and O'Carroll 2009). Whatever the underlying mechanisms responsible for this impressive velocity constancy, some of the neurons recorded (particularly FATCs) come closer to being ideal velocity estimators than anything previously described at a single neuron level, in any animal, with a progressive, monotonic rise in response over more than a 100-fold range of velocities (Figure 30C). The consistency of these responses over that range is especially remarkable considering the enormous range of global contrasts observed in the set of images used in this experiment, as described in previous papers employing the same image set (Brinkworth and O'Carroll 2009)

3.5.4 Behavioural Implication

Hemicordulia dragonflies exhibit numerous distinct behaviours including hawking, patrolling and aerobatic conspecific engagements. Each of these tasks places different constraints on any system encoding optic-flow information and this would provide selective pressure on dragonflies to either adopt an extremely flexible motion-detection system or one which specialised for different tasks. Here we have shown that Hemicordulia have several different systems for encoding different kinds of ego-motion. The question is how the different LTC subclasses lend themselves towards the dragonfly's behavioural tasks. In hawking behaviour, the detection and cancelling of subtle perturbations due to airflow requires a motion system capable of detecting slow velocities. SFATC cells exhibit a surprisingly robust response to grating patterns even at the very sluggish 0.1 °/s waveform. Such neurons would appear to be perfectly situated in a system designed to detect the slow-sustained shifts in optic flow that might occur during hawking behaviours.

Meanwhile, FATC cells exhibit extremely strong adaptation to slow-motion and are better suited to fast moving tasks such as patrolling or conspecific encounters. These neurons also show a preference for low spatial frequencies which is a more important background feature in fast-moving engagements. These neurons exhibit exquisite velocity constancy in natural image experiments despite great variations in the background, similar to those found in flies (Straw et al 2008).

Finally, SATC cells exhibit similar properties to LPTC found in flies, which have been linked to turning behaviours such as the optomotor response (Haikala et al., 2013). Unlike SFATCs, SATCs exhibit more classical tuning showing weak responses at both high and low temporal frequencies. In concert with FATCs, they may simply extend the velocity ranges over which Hemicordulia can operate (i.e. a slow and fast motion detector). As tuned neurons they are well capable of describing the magnitude of turning (ambiguities being eliminated using the faster FATC neurons) and thus could be used for relatively slower manoeuvres.

One remaining unanswered question is how all three of these cells would operate in closed loop. The presence of efference copies (Kim et al., 2015) in *Drosophila* causes motion caused by voluntary actions to be suppressed. This mechanism is appropriate for hovering and general turning behaviours, however in dragonfly conspecific engagements or target pursuit, stabilizing optic flow may become less important than target-related navigation tasks and rapid flight in complex 3-dimensional environments may result in objects creating fast motion on the eye without a corresponding voluntary action (such as passing a close large

obstacle during a turn). FATCs may be well placed for this task, preferring low-spatial frequency information moving at high speeds.

3.6 References

1) Anstisa S, Verstratenb FAJ & Matherc G, 1998, *The Motion Aftereffect, Trends in Cognitive Sciences* (vol. 2, no. 3, pp. 111-117), https://doi.org/10.1016/S1364-6613(98)01142-5

2) Arenz A, Drews MS, Richter FG, Ammer G & Borst A, 2017, *The Temporal Tuning of the Drosophila Motion Detectors is Determined by the Dynamics of Their Input Elements*, Current Biology (vol. 27, no. 7, pp. 929-944), https://doi.org/10.1016/j.cub.2017.01.051

3) Barnett PD, Nordstrom K & O'Carroll DC, 2010, *Motion Adaptation and the Velocity Coding of Natural Scenes*, Current Biology (vol. 20, pp. 994-999), https://doi.org/10.1016/j.cub.2010.03.072

4) Barlow HB & Levick WR, 1965, *The mechanism of directionally selective units in rabbit's retina*, Journal of Physiology (vol. 178, pp. 477-504) https://doi.org/10.1113/jphysiol.1965.sp007638

5) Borst A, Haag J & Reiff DF, 2010, *Fly Motion Vision*, Annual Review of Neuroscience (vol. 33, pp 49-70) https://doi-org.proxy.library.adelaide.edu.au/10.1146/annurev-neuro-060909-153155

6) Clifford CWG & Langley K, 1996, *A Model of Temporal Adaptation in Fly Motion Vision*, Vision Research (vol. 36, no. 16, pp. 2595-2608).

7) Clifford CWG, Ibbotson MR & Langley K, 1997, *An adaptive Reichardt detector model of motion adaptation in insects and mammals*, Visual Neuroscience (vol. 14, no. 4, pp. 741-749), https://doi.org/10.1017/S0952523800012694

8) de Haan R, L YJ & Nordstrom K, 2013, *Novel Flicker-Sensitive Visual Circuit Neurons Inhibited by Stationary Patterns*, Journal of Neuroscience (vol. 33, no. 21, pp. 8980-8989)

9) DeVoe RD, Kaiser W, Ohm J & Stone LS, 1982, *Horizontal movement detectors of honeybees: Directionally-selective visual neurons in the lobula and brain*, Journal of comparative physiology (vol. 147, no. 2, pp. 155-170).

10) Evans BJE, O'Carroll DC & Wiederman SD, 2016, Salience invariance with divisive normalization in higher-order insect neurons, *6th European Workshop on*

Visual Information Processing (EUVIP), Marseille, 2016, pp. 1-6. https://doi.org/10.1109/EUVIP.2016.7764588

11) Dvorak D, Srinivasan MV & French AS, 1980, *The Contrast Sensitivity of Fly Movement-Detecting Neurons*, Vision Research (vol. 20, pp. 397-407), https://doi.org/10.1016/0042-6989(80)90030-9

12) Field DJ, 1987, *Relations between the statistics of natural images and the properties of cortical cells.* Journal of the Optical Society of America A, (vol. 4, pp. 2379-2394)

13) Gruntman E, Romani S & Reiser MB, 2018, *Simple integration of fast excitation and offset, delayed inhibition computes directional selectivity in Drosophila*, Natura Neuroscience (vol. 21, pp. 250-257)

14) Harris RA, O'Carroll DC & Laughlin SB, 1999, *Adaptation and the temporal delay filter of fly motion detectors*, Vision Research (vol. 39, no. 16, pp. 2603-2613) https://doi.org/10.1016/S0042-6989(98)00297-1

15) Hausen K, 1982, *Motion sensitive interneurons in the optomotor system of the fly I The Horizontal Cells: Receptive Field Organization and Response Characteristics*, Biological Cybernetics (vol. 46, no. 1, pp. 67-79).

16) Hausen K, Egelhaaf M, 1989, *Neural Mechanisms of Visual Course Control in Insects.* In: Stavenga DG, ed. Facets of vision. Berlin: Springer; 1989: 391-424.

17) Ibbotson MR, 1991, *Wide-field motion-sensitive neurons tuned to horizontal movement in the honeybee, Apis mellifera*, Journal of Comparative Physiology A (vol. 168, no. 1, pp. 91-102), https://doi.org/10.1007/BF00217107

18) Jung SN, Borst A & Haag J, 2011, *Flight Activity Alters Velocity Tuning of Fly Motion-Sensitive Neurons*, Journal of Neuroscience (vol. 25, pp. 9231-9237), https://doi.org/10.1523/JNEUROSCI.1138-11.2011

19) Kim AJ, Fitzgerald JK & Maimon G, 2015, *Cellular evidence for efference copy in Drosophila visuomotor processing*, Nature Neuroscience (vol. 18, no. 9, pp. 1247-1255), https://doi.org/10.1038/nn.4083

20) Krapp HG, Hengstenberg R & Egelhaaf M, 2001, *Binocular Contributions to Optic Flow Processing in the Fly Visual System*, Journal of Neurophysiology (vol. 85, no. 2, pp. 724-734), https://doi.org/10.1152/jn.2001.85.2.724

21) Land MF, 1997, *Visual Acuity in Insects*, Annual Review of Entomology (vol. 42, pp. 147-177), https://doi.org/10.1146/annurev.ento.42.1.147

22) Laughlin SB, 1974, *Neural integration in the first optic neuropile of dragonflies. III. The transfer of angular information*. Journal of Comparative Physiology A Neuroethology, Sensory, Neural, Behavioural Physiology (vol. 92, pp. 377–396), https://doi.org/10.1007/BF00694708

23) Li JL, Lindemann JP & Egelhaaf M, 2017, *Location motion adaptation enhances the representation of spatial structure at EMD arrays*, PLoS Computational Biology (vol. 13, no. 12), https://doi.org/10.1371/journal.pcbi.1005919

24) Longden KD & Krapp HG, 2009, *State-Dependent Performance of Optic-Flow Processing Interneurons*, Journal of Neurophysiology (vol. 102, no. 6, pp. 3606-3618), https://doi.org/10.1152/jn.00395.2009

25) Longden KD, Muzzu T, Cook DJ, Schultz SR & Krapp HG, 2014, Nutritional State Modulates the Neural Processing of Visual Motion (vol. 24, no. 8, pp. 890-895), https://doi.org/10.1016/j.cub.2014.03.005

26) Maddess T & Laughlin SB, 1985, *Adaptation of the Motion-Sensitive Neuron H1 is Generated Locally and Governed by Contrast Frequency*, Proceedings of the Royal Society (Vol. 225, no. 1239, pp. 251-275)

27) Maisak MS, Haag J, Ammer G, Serbe E, Meier M, Leonhardt A, Schilling T, Bahl A, Rubin GM, Nern A, Dickson BJ, Reiff DF, Hopp E & Borst A, 2013, *A directional tuning map of Drosophila elementary motion detectors*, Nature (vol. 500, pp. 212-216)

28) Mertes M, Dittmar L, Egelhaaf M & Boeddekker N, 2014, *Visual motion-sensitive neurons in the bumblebee brain convey information about landmarks during a navigational flight*, Frontiers in Behavioural Neuroscience (vol. 8, 335), https://doi.org/10.3389/fnbeh.2014.00335

29) Nordstrom K & O'Carroll DC, 2009, *The motion after-effect: local and global contributions to contrast sensitivity*, Proceedings of the Royal Society B (vol. 276, pp. 1545-1554) https://doi.org/10.1098/rspb.2008.1932

30) Parsons MM, Krapp HG & Laughlin SB, 2006, *A motion-sensitive neurone responds to signals from the two visual systems of the blowfly, the compound eyes and ocelli*, Journal of Experimental Biology (vol. 209, pp. 4464-4474), https://doi.org/10.1242/jeb.02560

31) Paulk AC, Phillips-Portillo J, Dacks, AM, Fellous JM & Gronenberg W, 2008, *The Processing of Color, Motion, and Stimulus Timing Are Anatomically Segregated in*

the Bumblebee Brain, The Journal of Neuroscience (vol. 28, no. 25, pp. 6319-6332), https://doi.org/10.1523/JNEUROSCI.1196-08.2008

32) Preibisch S, Saalfeld S, & Tomancak P, 2009, *Globally optimal stitching of tiled 3D microscopic image acquisitions*, Bioinformatics, 25(11):1463-1465

33) Shoemaker PA, O'Carroll DC & Straw AD, 2005, *Velocity constancy and models for wide-field visual motion detection in insects*, Biological Cybernetics (vol. 93, no. 4, pp. 275-287), https://doi.org/10.1007/s00422-005-0007-y

34) Stöckl AL, O'Carroll DC & Warrant EJ, 2016, *Neural Summation in the Hawkmoth Visual System Extends the Limits of Vision in Dim Light*, Current Biology (vol. 26, no. 6, pp. 821-826), https://doi.org/10.1016/j.cub.2016.01.030

35) Strausfeld NJ & Lee JK, 1991, *Neuronal basis for parallel visual processing in the fly*, Visual Neuroscience (vol. 7, no. 1-2, pp. 13-33) https://doi-org/10.1017/S0952523800010919

36) Straw AD, Rainsford T & O'Carroll DC, 2008, *Contrast sensitivity of insect motion detectors to natural images*, Journal of Vision (vol. 8, no. 32, pp. 1-9)

37) Suver MP, Mamiya A & Dickinson MH, *Octopamine Neurons Mediate Flight-Induced Modulation of Visual Processing in Drosophila*, Current Biology (vol. 22, no. 24, pp. 2294-2302)

38) Theobald JC, Duistermars BJ, Ringach DL & Fry MA, 2008, *Flies see second-order motion*, Current Biology (vol. 18, no. 11, pp. 464-465),

39) Theobald JC, Warrant EJ & O'Carroll DC, 2010, *Wide-field motion tuning in nocturnal hawkmoths*, Proceedings of the Royal Society B (vol. 277, no. 1683, pp. 853-860).

40) Tolhurst DF, Tadmor Y & Chao T, 1992, *Amplitude spectra of natural images*, Ophthalmology and Physiological Optics (vol. 12, pp. 229-232)

41) Wicklein & Varju, 1999, *Visual system of the european humming bird hawkmoth Macroglossum stellatarum (sphingidae, lipidoptera): Motion-sensitive interneurons of the lobula plate*, Journal of Comparative Neurology (vol. 408, no. 2, pp. 272-282). https://doi.org/10.1002/(SICI)1096-9861(19990531)408:2<272::AID-CNE8>3.0.CO;2-9

42) Wiederman SD, Shoemaker PA & O'Carroll DC, 2008, *A model for the detection of moving targets in visual clutter inspired by insect physiology*. PLoS One (3:e2784), https://doi.org/10.1371/journal.pone.0002784

43) Wiederman SD, Shoemaker PA & O'Carroll DC, 2013, *Correlation between OFF and ON Channels Underlies Dark Target Selectivity in an Insect Visual Systems*, The Journal of Neuroscience (vol. 33, no. 32, pp. 13225-13232)

4 Further Properties of Dragonfly Lobula Tangential Cells

4.1 Preamble

Having finally gotten to the end of the subtype mystery of dragonfly tangential cells, I was able to begin investigating further properties. The existence of three subtypes immediately raises numerous questions regarding their receptive field placement and other tuning properties. The following chapter discusses several experiments designed to elucidate some of these properties. As these experiments were conducted towards the very end of my PhD, some remain in a somewhat unpolished form, but the data collected, and the analysis thereof gives some insight into the similarity and variations between these different neurons.

4.2 Anatomical Investigations

4.2.1 Introduction

Hemicordulia dragonflies possess three subclasses of optic neurons sensitive to coherent widefield motion. These subclasses have been classified purely based on their physiological properties and differential adaptation to motion of different temporal frequencies. However, a thorough investigation of how these neurons may differ anatomically does not exist. My previous work showed an example of a single dragonfly Lobula Tangential Cell (LTC) whose main input arborisation was in the sub-lobula subregion of the lobula complex. While this differs from commonly studied insect species (such as *Diptera*), bees possess a sub-lobula and lack a lobula plate (the location of most WMFS neurons). Here I present dye injections from three different WFMS neurons found in the lobula of the dragonfly. It is clear from these three examples that LTCs in dragonflies represent a diverse class of neurons, which differ both physiologically and anatomically.

4.2.2 Methodology

The morphology of a widefield motion-sensitive neuron was visualized by intracellular labelling with Lucifer Yellow. Iontophoresis was achieved by passing a negative current of at least 1nA through electrodes tip-filled with Lucifer Yellow for at least 12 minutes (up to two hours in the most detailed case). Brains were then carefully dissected, fixed overnight in 4% paraformaldehyde at 4°C, dehydrated in ethanol series (70%, 90%, 100%, 100%), cleared in methyl salicylate and mounted slide with three spacer rings and mounted in Permount for imaging. For two such neurons, a secondary anti-lucifer antibody was incubated (after using

xylene to remove the permount) and a secondary dye (dye-light) attached and incubated for 3 further days. The newly stained specimen was then cleared and mounted as above. The sample was scanned using a 10x objective (the third used a 30x) on a confocal microscope and the 3D slices reconstructed using Neutube.

4.2.3 Results

During my research three neurons qualifying as WMFS neurons were stained, two of which were recorded and processed by my colleague Joseph Fabian (Figure 31). The third (and most detailed I might add) was recorded and stained by myself. This third neuron was injected for an extensive period (3 hours) using a stepping (square wave steps of ~1Hz) current of up to 3nA, which may give an indication of the necessary conditions to achieve highly detailed neuronal reconstructions. It should also be noted, that similar injection procedures were conducted on more than a dozen WMFS neurons with no successful staining, many of which were injected for multiple hours with no effect.

Classifying these neurons was straightforward for the cells shown in Figure 32B & C as sufficient physiological data was recorded. From the temporal tuning alone, it is clear that cell B & C are of subtype SATC and SFATC respectively. Given the anatomical similarities between cell A & B, it is likely that this neuron is also of type SATC (clearly not SFATC based on morphology). Analysis of flicker responses from this neuron also rule out FATC (it only responded to one polarity of flicker – OFF).

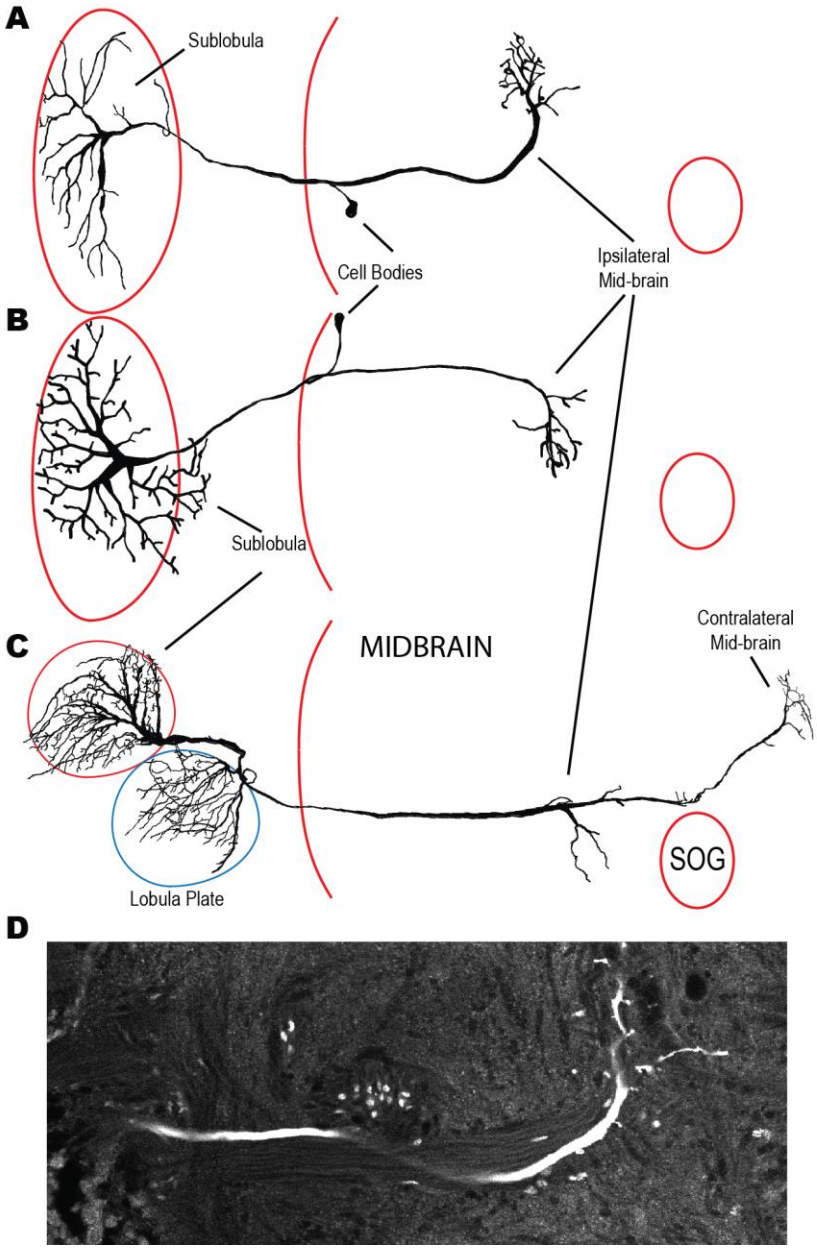

Figure 31: Example of the morphology of dragonfly LTCs (Neurons A, B stained by Joseph Fabian) with major landmarks shown in red and blue (SOG: Sub-oesophageal ganglion) ***A****, An LTC stained from an unknown subtype of LTC. It possesses a single main input arborisation (sub-lobula).* ***B****, A second LTC stained from a SATC. Like (A), this neuron possesses a single sub-lobula input arborisation. The relative orientations of the two neurons are reversed with the mid-brain arborisations forking dorsally and ventrally respectively.* ***C****, A highly detailed stain from a SFATC neuron (though the cell body remains unstained). Unlike the A & B, this neuron not only crosses the mid-brain (contralateral input/output) but also possesses two main input arborisations in the sublobula and lobula plate.* ***D,*** *Raw confocal image depicting the axon track of neuron A placed in a broader tract containing numerous parallel axons of (presumably) LTC neurons.*

If we take this then to be two SATCs and one SFATC, we can see that the SFATC is clearly the more complex neuron with two unique arborisations. Though perhaps speculative, it seems reasonable that the additional arborisation in the lobula plate and/or the contralateral

mid-brain could be involved in the adaptation mechanism seen in SFATC physiology. It should also be noted, that this neuron followed an entirely separate tract from the first two, running very superficially to the brain surface.

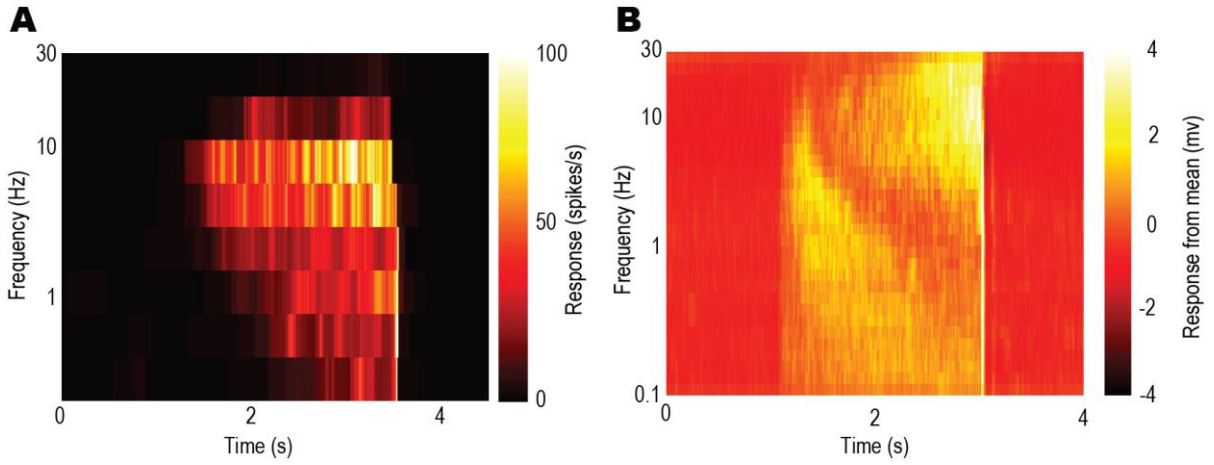

Figure 32: Physiological responses of two LTC cells shown in Figure 31 A) Temporal Tuning Curve data (response to increasing contrast gratings, see Methods from Chapter 3) taken from Cell B. Though at a lower frequency resolution, this cell conforms to the pattern exhibited in SATC neurons (consistent tuning across time) with a peak ~5-10Hz. B) Temporal Tuning Curve data taken from Cell C. Due to abrupt cessation of spiking partway through the experiment, this data was reconstructed from graded potentials by removing the mean response from each trial (to normalise across trials). The underlying signal conforms to SFATC-like behaviour.

4.2.4 Discussion

Though not a main emphasis of my research, anatomical staining is a vital technique for unequivocally distinguishing different neurons classes. The certain separation morphologically between SATC and SFATC further cements their existence as separate neuron classes. Of particular note for the SFATC neuron is the dual arborisation in both sublobula and lobula plate. This is a completely novel finding among all insect species. While flies have their WFMS neurons in the lobula plate (LPTC) and bees in the sublobula, dragonflies not only possess both neuropils (Fabian et al., in preparation) but also utilize both neuropils in their encoding of wide-field motion. The presence of neurons that arborize in both neuropils also has an additional implication as it further demonstrates the individual roles each neuropil plays. This contradicts previous theories (Strausfeld 2005), which placed the two neuropils as functional homologues. Instead, these two neuropils can function in tandem towards producing widefield-motion information and they may play complementary specialised roles towards this process.

Finally, this two-arborisation SFATC neuron has significant implications for evolutionary investigations into the origins and commonalities between species of these two neuropils. As a very old species dragonflies, predate both bees and flies (and are higher on the 'family tree') indicating it is very unlikely that the lobula plate and sublobula are directly related but are instead equivalents of one another. As such, this finding narrows the feasible models for development of these brain regions.

4.3 Receptive Fields

4.3.1 Introduction

Hemicordulia Dragonflies possess three subclasses of lobula neurons sensitive to coherent widefield motion. These Lobula Tangential Cells (LTCs) integrate local motion information over large areas of the dragonfly's vision. All three subclasses are direction opponent (that is are inhibited by anti-preferred direction motion). These neurons are distinguished by their differential adaptation to motion stimuli, specifically drifting grating motion at varied temporal frequencies. It has been shown previously (Harris et al., 2000) that prolonged stimulation of Horizontal System (HS) neurons in *Eristalis* to drifting sinusoids changed cellular responses through varying forms of motion adaptation (these will be described later).

Here I sought to investigate the receptive fields of LTCs and to see whether there was an established pattern.

4.3.2 Results

One challenge of mapping receptive fields of widefield-motion sensitive neurons is inherent in the purported function of said neurons, i.e. that they integrate motion over large regions. The key factor which limits receptive field mapping in the minimally sized stimulus that will elicit a measurable response. In other species (Barnett et al., 2007), the gain of the widefield-motion system is sufficiently large that even small targets can provide a robust (if sub-optimal) neuronal response (Krapp & Hengstenberg 1996). Hemicordulia LTCs do not appear to be as accommodating, at least, not universally so. Of the thirty-five neurons where robust widefield responses were found, only twenty-five proved amenable to receptive field mapping using a small target (Figure 33A).

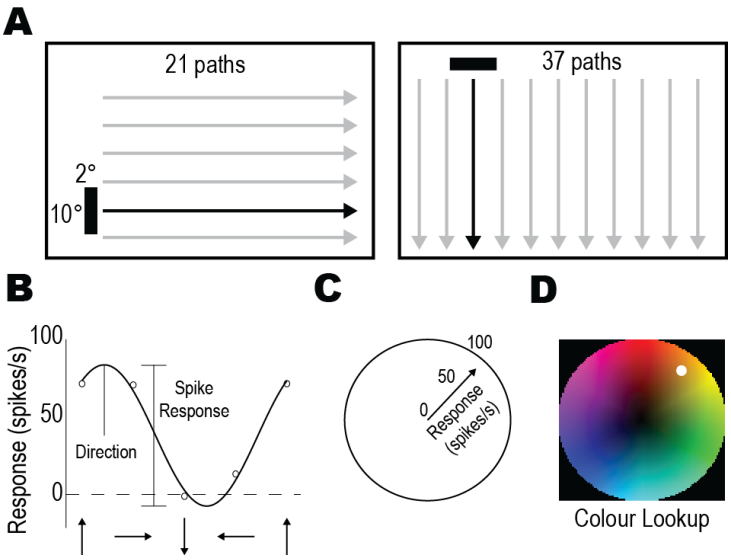

Figure 33: Example of receptive field methodology. ***A,*** *Small targets are drifted through the receptive field at different vertical (left) and horizontal (right) locations to elicit the response at that location.* ***B,*** *For each point in the receptive field shown in A, the four measurements (up, down, left, right) are used to fit a sinusoidal curve. The preferred direction and direction selectivity (preference for the preferred direction over the anti-preferred direction) are then calculated.* ***C,*** *For each point on the receptive field, the preferred direction and response are then calculated in polar coordinate terms.* ***D,*** *Finally, using the polar coordinates, the colour-lookup is based on a colour wheel where the saturation defines the response and the hue determines the direction (white dot shows example carried across from B and C).*

Another intrinsic difficulty in describing LTC receptive fields is how to present the data. It is entirely possible for an LTC to register completely different responses in different parts of a receptive field (as might occur with a neuron which is sensitive to rotational movement where the pole of rotation is in the centre of the receptive field). Thus, being able to show both the magnitude of response and the direction of responses at each spatial location is important. However, even this falls short. It is also necessary to distinguish inhibitory interactions. Some neurons might exhibit strong excitation and inhibition when exposed to preferred and anti-preferred motion. Some may exhibit strong excitation or weak or no inhibition. Being able to 'at-a-glance' capture this information along with the spatial extent is effectively impossible (there are too many dimensions to the data).

To maximize the useful information in a way that can be compared across neurons (i.e. at-a-glance) I have developed a colour/saturation model for neuronal response. To begin, I take the measurements of the response in the four directions measured and by fitting a sinusoidal curve (Figure 33B) calculate the preferred direction and magnitude of response (i.e. the difference between the maximum and minimum response). I then convert this information into polar coordinates (radius and direction, Figure 33C) and then apply a colour lookup table

(Figure 33D). In essence, the colour of a part of the receptive field (i.e. the RGB colour) indicates the direction, while the saturation of the colour represents the difference of the response from spontaneous activity.

I have presented the three subclasses of LTC together (Figure 34A (SATC), B (FATC) & C (SFATC)). It is apparent that many LTCs have large receptive fields that generally occupy the ipsilateral hemisphere. SATC & SFATC cells in particular appear to conform to this general trend. FATC cells exhibit a number of exceptions with two examples of contralateral receptive fields. All LTCs measured have single-direction sensitivity (at least measurably so) with no clear examples of rotation-sensitive neurons (characterised by different direction sensitivity in different regions of the receptive field, creating a 'spiral').

Four neurons (marked with *) produced very limited responses to the stimulus presented. To better analyse the receptive fields of these neurons (Figure 35A), the recordings taken from all four directions were summed (Figure 35B) and the average taken rather than using the directionality. As can be seen, even though the responses were very weak (measuring only a few tens of spikes in some cases) the resulting receptive fields were very much in line with those presented in (Figure 34A-C). In general, LTCs appear to have very large ispilateral receptive fields.

There are some notable single-cell curiosities. One cell (Figure 34B **) measured had a receptive field that was central and small. Previously identified as a SFATC cell, this cell with its receptive field taken into account this may represents a new subclass of LTC or even a different class of cell entirely. One possibility, is that it is a mislabelled Small-Field Small Target Motion Detector (SF-STMD) (which also possess small receptive fields mappable by small targets). When stimulated with a relatively small target (8 degrees), it achieved a response ~150 spikes/s. However, it achieved similar responses to a grating and during spatial tuning and had a strong response to a broad range of spatial frequencies (unlike SF-STMDs). Interestingly it also gave good responses to natural scenes including exhibiting velocity constancy. Without further anatomical or other physiological data it is difficult to place such a neuron in any classification scheme.

In four cells (all of type SFATC), the stimulus monitor was moved to better capture the receptive field of the neurons (Figure 36). Receptive fields taken from these four neurons show receptive fields over 100 degrees across (the size of the stimulus monitor), well lateral to the frontal part of the dragonfly eye.

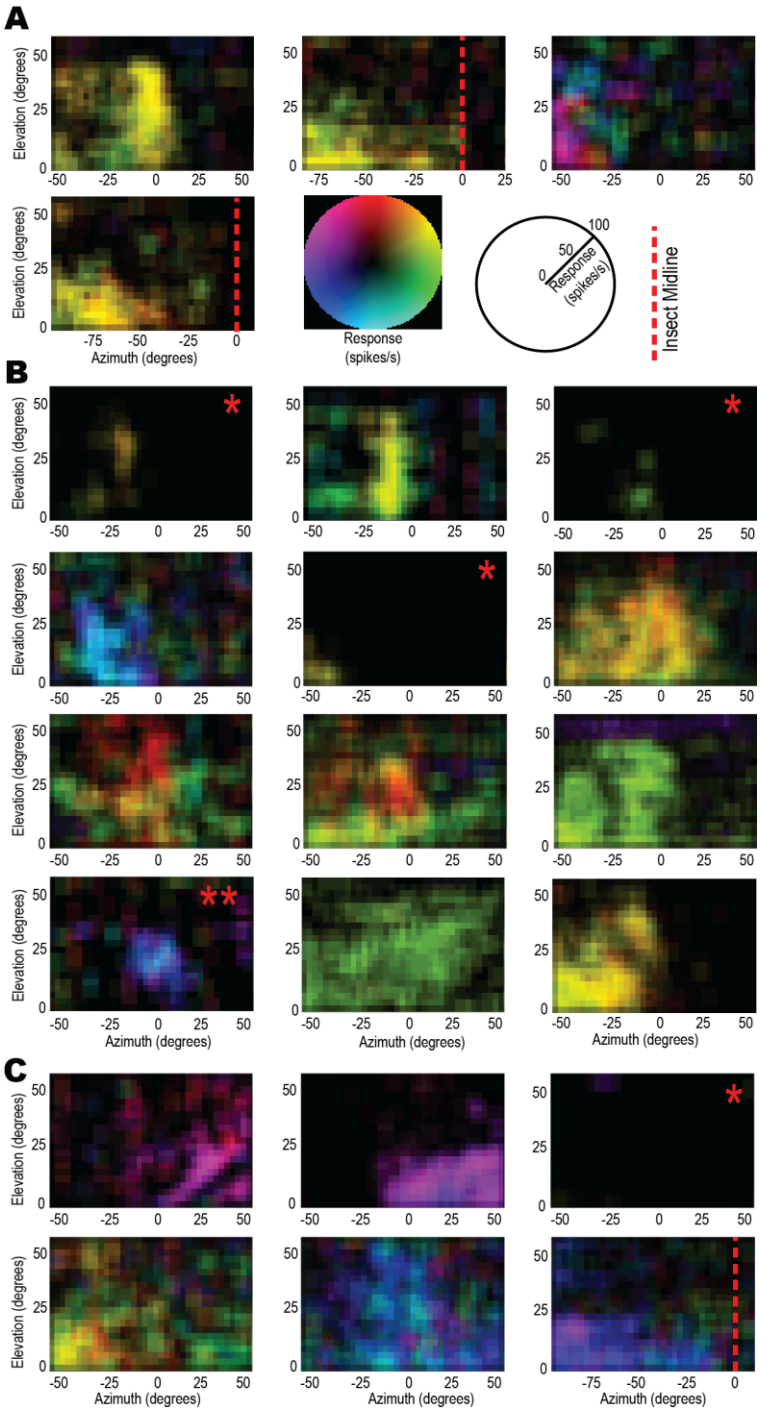

Figure 34: LTC Receptive Fields. Note the colour lookup table where hue corresponds to preferred direction and saturation refers to total mean response. Also note insect midlines where marked (for circumstances of shifted stimulus). Finally receptive fields marked with (*) or (**) will be discussed in detail later. ***A,*** *SATC receptive fields for four different cells.* ***B,*** *As (A) for FATC neurons. Here there are two examples of neurons, which have dominant right-hemisphere inputs (first two plots).* ***C,*** *As (A) for SFATC neurons. This class exhibits largely ipsilateral receptive fields.*

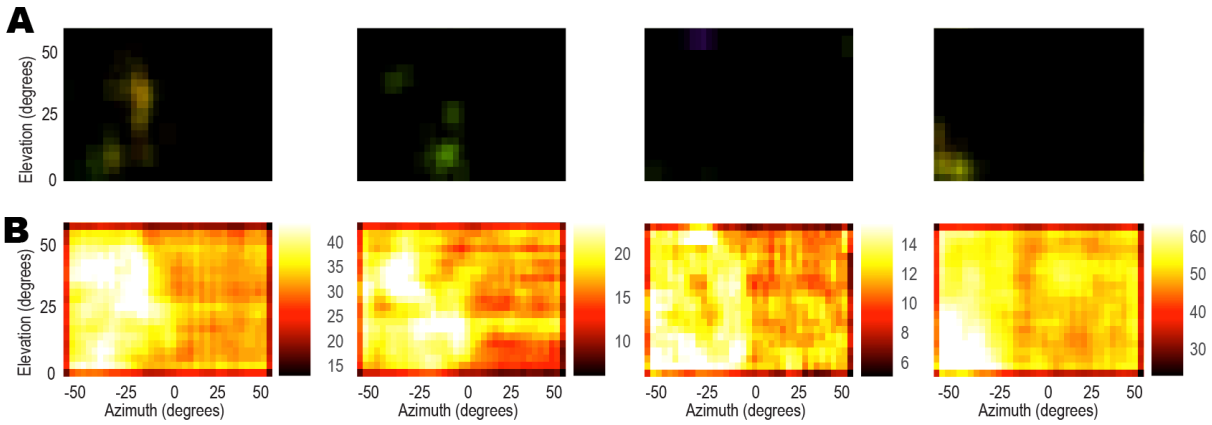

Figure 35: Alternate method for capturing receptive field shape. ***A,*** *Receptive fields reproduced from Figure 34.* ***B****, Receptive fields using the sum of all four directions (rather than directionality) to show overall sensitivity.*

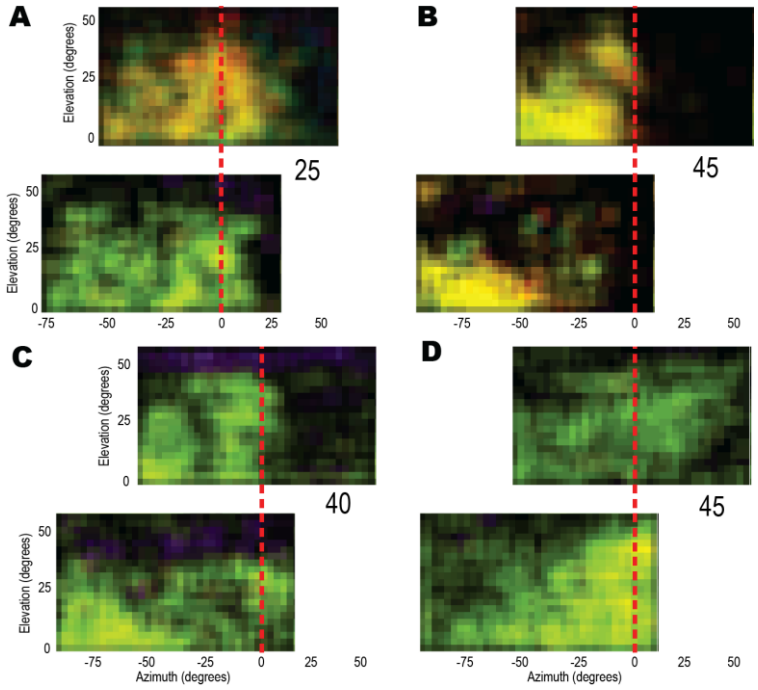

Figure 36: Example of four receptive fields, where the stimulus monitor was moved (angle moved in degrees shown in each image pair) to extend the effective range of measurement. Each pair of images (arrayed vertically) has been vertically aligned to ensure an equivalent insect midline (red dashed line). All four cells demonstrate large receptive fields, extending well into the lateral eye of the insect (>75° laterally).

Based on the receptive fields shown in (Figure 34, Figure 35 & Figure 36), the general shape of the receptive fields of dragonfly LTCs can be interpreted. LTC receptive fields can extend over large regions of the field of view. In general SATC & SFATC cells appear to be ipsilateral while some FATC cells appear to have contralateral inputs in addition.

4.3.3 Discussion

LTC receptive fields appear similar to widefield-sensitive neurons found in other species such as flies (Hausen 1982). From their relative location (due to the orientation of the dragonfly head), many of these cells might be analogues of LPTC-HSN cells found in flies (Hausen 1982) though with a larger vertical extent.

Our current stimulus setup limits the angular extent to which these neurons can be stimulated representing only a fraction of the field of view which prevents a complete comparison. In flies for example, the VS system neurons have inputs from a huge extent of the visual field (Krapp et al., 1998) which far exceeds our setup's capability which inherently limits any inter-species comparisons.

LPTC cells in flies have been shown to have ipsilateral hemispheres of excitation and only show contralateral properties in the presence of ipsilateral motion. The stimulus here would not capture that information, and this would be a sensible follow-up experiment to see if the different dragonfly LTCs can be homologously mapped to well-characterised neurons in the fly (such as H1).

It has been previously shown that LTCs can be subcategorized based on their differential motion adaptation to gratings (Chapter 3). Here there is insufficient data to draw any strong conclusions about the receptive field properties of these subtypes. However, there appears to be significant overlap in the regions of sensitivity of these different neuron subtypes. This creates questions of their interactions. For example, if SATC and SFATC neurons draw input from similar regions of space, do they then interact with each other or perhaps are integrated by higher-level ego-motion neurons?

Previous work has shown that these neurons exhibit strong velocity constancy (Chapter 3) when stimulated by moving natural scenes. All these neurons (with a single exception) exhibit large receptive fields integrating information over several tens of degrees. Despite significant differences in the contrast profiles of the images used and differing features in different parts of the images, the response of these neurons remains relatively constant for a given velocity. These large receptive fields may help achieve this velocity constancy by sampling from many parts of an image simultaneously, effectively reducing the differences between different scenes and scene sub-sections.

4.4 References

1) Barnett PD, Nordstrom K & O'Carroll DC, 2007, *Retinotopic organization of Small-Field-Target-Detecting Neurons in the Insect Visual System*, Current Biology (vol. 17, pp. 569-578), https://doi.org/10.1016/j.cub.2007.02.039

2) Harris, RA, O'Carroll DC & Laughlin SB, 2000, *Contrast Gain Reduction in Fly Motion Adaptation*, Neuron (vol. 28, pp. 595-606) https://doi.org/10.1016/S0896-6273(00)00136-7

3) Hausen K, 1982b, *Motion sensitive interneurons in the optomotor system of the fly. II The horizontal cells: Receptive field organization and response characteristics*, Biological Cybernetics (vol. 46, no. 1, pp. 67-79), https://doi.org/10.1007/BF00335352

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

5) Strausfeld, 2005, *The evolution of crustacean and insect optic lobes and the origins of chiasmata*, Arthropod Structure & Development (vol. 34, no. 3, pp. 235-256), https://doi.org/10.1016/j.asd.2005.04.001

5 Discrimination of Features in Natural Scenes Driven by Selective Attention

5.1 Preamble

Of all the work I have done in this PhD, this particular chapter binds the closest to my original intentions. It also adheres most closely to my original interest, which was the behaviour of dragonfly neurons in cluttered scenarios. While it is true that dragonflies and other predators often pick predation locations amenable to easily identifying small prey, such as using a blue sky, research in STMD neurons indicate that they can be 'tricked' by background information. When one considers conspecific engagements such as fighting or mating, it seems impossible that these do not occur with significant background motion. How the system manages to avoid being 'tricked' is a particular interest of mine.

While selective attention was not the main focus of my PhD, it did turn out to be the best answer for how natural images influence neurons proposed to underlie target-tracking. Thus, while this paper requires a decent understanding of selective attention, it doesn't seek to be a full description of the phenomenon, rather how it interacts with clutter.

This project was subject to numerous false starts. Attempting to disambiguate target and background responses proved quite challenging and the original stimulus paradigm turned out to be a hopeless failure requiring a complete redesign. Additionally, the length of recordings required for the experiments herein further complicated matters. Testing multiple images at multiple phases with sufficient controls while simultaneously varying an additional parameter (i.e. the point of the experiment) necessarily leads to long experiments. Fortunately, a sudden spike in electrophysiology proficiency saved me in my last recording season.

Statement of Authorship

Title of Paper	Discrimination of Features in Natural Scenes Driven by Selective Attention
Publication Status	Unpublished and Unsubmitted work written in manuscript style
Publication Details	Evans BJE, Fabian JM, O'Carroll DC & Wiederman SD, unpublished

Principle Author

Name of Principal Author	Bernard Evans		
Contribution to the Paper	Experiment conceptualisation, animal collection, data collection, data analysis, data interpretation, figure generation, manuscript authorship.		
Overall percentage (%)	70%		
Certification	This paper reports on original research I conducted during the period of my Higher Degree by Research candidature and is not subject to any obligations or contractual agreements with a third party that would constrain its inclusion in this thesis. I am the primary author of this paper.		
Signature		Date	14/12/2018

Co-Author Contributions

By signing the Statement of Authorship, each author certifies that:

i. the candidate's stated contribution to the publication is accurate (as detailed above);
ii. permission is granted for the candidate in include the publication in the thesis; and
iii. the sum of all co-author contributions is equal to 100% less the candidate's stated contribution.

Name of Co-Author	David C O'Carroll		
Contribution to the Paper	Experiment conceptualisation, data interpretation, manuscript evaluation		
Signature		Date	4/12/2018

Name of Co-Author	Joseph M Fabian		
Contribution to the Paper	Animal collection, data collection, immunohistochemistry processing, neuron imaging, figure generation (anatomy), manuscript evaluation		
Signature		Date	28/11/18

Name of Co-Author	Steven D Wiederman		
Contribution to the Paper	Experiment conceptualisation, data interpretation, manuscript evaluation		
Signature		Date	14/12/18

5.2 Introduction

Many aerial insects travel through visually cluttered environments whilst tracking prey and conspecifics (Collett & Land 1978; Mischiati, Lin et al., 2015). For example, dragonflies must detect a moving target that may have poor or highly variable contrast against a moving, cluttered background. Moreover, the dragonfly may need to selectively attend to a single target, which competes with similarly sized features that can occur in complex scenes (e.g. feeding in swarms). This behavior necessitates an underlying neuronal network that discriminates potential prey from distracting features of the moving background.

Neurons that are tuned to small moving targets, which are likely to underlie this behaviour, have been shown in several flying insect species (Collett 1971; O'Carroll 1993; Nordstrom & O'Carroll 2006; Geurten et al 2007). These Small Target Motion Detector (STMD) neurons respond robustly to small targets (less than 5 degrees) moving at any location within their receptive field (O'Carroll 1993). STMDs are sensitive to target contrast and are tuned to both the size and velocity of the target, irrespective of the location within the receptive field (O'Carroll & Wiederman, 2014).

CSTMD1 (Centrifugal Small Target Motion Detector 1), a well-studied STMD found in the dragonfly *Hemicordulia tau*, has proven an excellent model system for investigating the neural mechanisms underlying target detection, even when those targets are embedded in natural images (Wiederman & O'Carroll 2011). CSTMD1 possesses both excitatory and inhibitory visual hemifields on either side of the visual midline. CSTMD1 exhibits a 'predictive modulation of gain', which facilitates neuronal responses to targets moving on continuous trajectories in front of the prior path, whilst suppressing jumps into the surround (Wiederman, Fabian et al., 2017). Additonally, CSTMD1 exhibits 'selective attention', responding only to a single target, when presented with a pair, thus ignoring the distracter (Wiederman & O'Carroll 2013).

Past studies of STMD responses to the presentation of visual clutter were limited to artificially generated backgrounds with low phase congruence (Nordstrom et al 2006) or spatially constrained scenes without relative motion (Wiederman & O'Carroll 2011). How STMD neurons respond to targets in natural scenes with varying degrees of relative motion is not known, yet essential for interpreting how dragonflies pursue targets in real-world scenarios. Furthermore, the attribute of selective attention raises the question of how competitive selection processes interact with background features in natural imagery.

Here we show neuronal responses from CSTMD1, as well as a never-before described neuron, Binocular Small Target Motion Detector 2 (BSTMD2). We present various natural images at different speeds to determine which properties of naturalistic conditions most affect STMD performance in clutter. We show that both CSTMD1 and BSTMD2 have reduced performance in cluttered scenarios, particularly when the background speed is higher than the target's speed. We show that this is not simply due to the variable target contrast or inhibitory interactions from background motion, rather is the result of interplay between the target and background features competing for selection.

5.3 Methods

5.3.1 Electrophysiology

49 wild-caught, dragonflies (*Hemicordulia*) were immobilized with a beeswax and rosin (1:1) mixture with the head tilted forward to access the posterior surface. A hole was cut above the lobula where we recorded intracellularly with aluminosilicate micropipettes pulled on a Sutter Instruments P-97 puller and backfilled with KCl (2 M). The electrode tip resistance was typically 50-150 MΩ.

Data were only ever excluded due to pathological damage of the neuron or extensive habituation, resulting in experiment cessation. All means are calculated from biological replicates (i.e. repeated measurements from identified neurons in different animals). Each biological replicate represents the mean between 1 and 5 technical replicates.

5.3.2 Visual Stimuli

We presented stimuli on high definition LCD monitors (120-165 Hz). The animal was placed 20 cm away and centred on the visual midline. Contrast stimuli were presented at screen centre to minimize off-axis artefacts. The display was projection distorted (OpenGL) to present angular degrees correctly from the animal's perspective and extended 104° (horizontally) by 58.5° (vertically) of the dragonfly's visual field. Stimulus scripts were written using MATLAB's Psychtoolbox and integrated into the data acquisition system.

Natural images were taken from a series of panoramic scenes, as described in Brinkworth & O'Carroll 2009. We displayed on the monitor 8-bit versions of the high dynamic range images with linear gamma and a resolution of 8000×1600 pixels.

For classification of neurons, we presented a sequence of stimuli: a gyrated, randomly generated texel pattern (1°), grey to black and grey to white full-screen flicker (White 338 cd/m^2, Black 0.5 cd/m^2), moving edges (up, down, left and right, 25°/s), moving bars (2°

width, up, down, left and right, 25°/s) and a square-wave grating pattern (0.025 cycles/°, 6.25 Hz, up, down, left and right). Neurons were identified as CSTMD1 or BSTMD2 by their characteristic tuning responses, spike waveforms and unique receptive field.

5.3.3 Neuroanatomy

The morphology of a BSTMD2 neuron was visualized by intracellular labelling with Lucifer Yellow. Iontophoresis was achieved by passing 1 nA negative current through electrodes tip-filled with 4% Lucifer Yellow solution in 0.1 M LiCl for 12 minutes. Brains were then carefully dissected, fixed overnight in 4% paraformaldehyde at 4°C, dehydrated in ethanol series (70%, 90%, 100%, 100%), cleared in methyl salicylate and mounted on a cavity slide for imaging. The sample was scanned using a confocal microscope and the 3D slices reconstructed using Neutube.

5.3.4 Experimental Design and Statistical Analysis

All analysis was completed in MATLAB. Spike-counting was done using a custom-written spike-counting script. Curve fits used MATLAB's in-built curve-fitting tools. All statistical tests were either paired t-tests or two-sided non-parametric tests (Kruskal Wallis test with Dunn's multiple comparison correction). All p values are reported as raw numbers in text if significant differences exist (unmarked otherwise) or as < 0.0001 if sufficiently small. Box and whisker plots represent the 75^{th}, 50^{th} and 25^{th} quartiles (lines) with raw data overlaid.

5.4 Results

5.4.1 CSTMD1 & BSTMD2

We tested the ability of dragonfly visual neurons to discriminate targets in moving backgrounds by recording from STMD neurons with large receptive fields. We successfully obtained repeated intracellular recordings from two such neurons, the well-characterised CSTMD1 (Geurten et al., 2007, Wiederman and O'Carroll 2013, Wiederman, Fabian et al., 2017) and the newly described neuron named Binocular Small Target Motion Detector 2 (BSTMD2). During recordings, visual stimuli were presented on a display centered on the dragonfly's midline (Figure 37A). We recorded from CSTMD1 and BSTMD2 in the left brain hemifield, and both neurons would have their mirror-symmetric counterparts on the opposite side of the brain. Individual examples of spiking activity in response to a moving target reveals robust responses in both CSTMD1 and BSTMD2 (Figure 37B).

CSTMD1 exhibits excitation or inhibiton dependent on the target's spatial location divided along the midline of the animal (Figure 37C). Responses are robust to targets moving in

either direction but only in the right hemifield (from the dragonfly's perspective). In the left hemifield (ipsilateral to our recording sight), CSTMD1 exhibits strong inhibition to targets moving either left or right. CSTMD1 is weakly directional preferring targets moving away from the midline.

BSTMD2 responds in a strongly direction-selective manner, inhibited by rightward targets (from the dragonfly's perspective) and excited by leftwards targets (Figure 37C). The magnitude of excitation is dependent on target location with more robust activity to motion presented in the left hemifield (ipsilateral to the recording site).

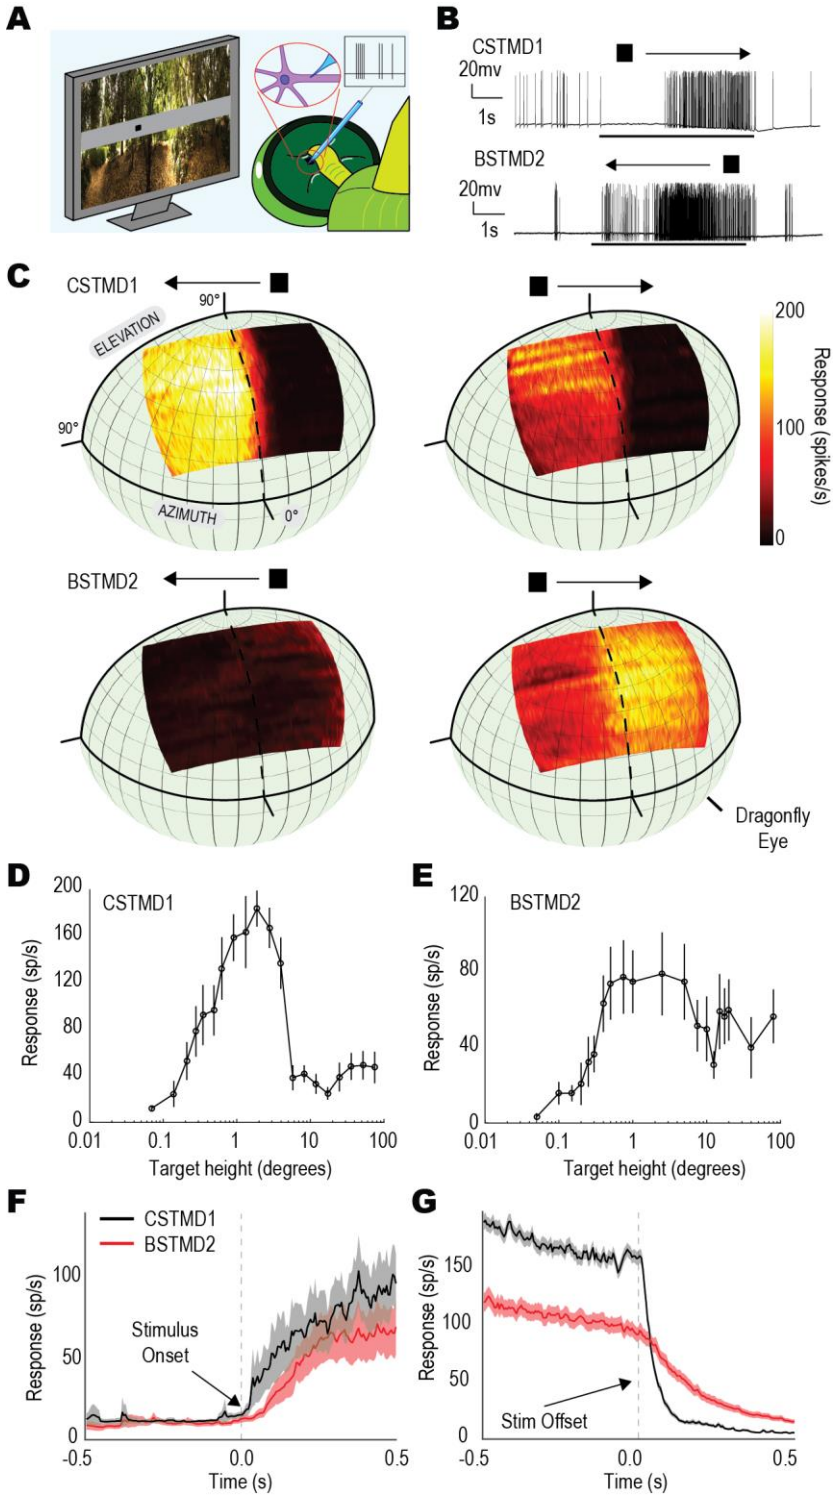

Figure 37: Response properties of BSTMD2 and CSTMD1 reveal large receptive fields and tuning to small, moving targets. ***A,*** *Illustration of intracellular recordings from the optic lobe whilst the dragonfly is presented with visual stimuli on an LCD display.* ***B,*** *Example raw traces of neuronal responses to a small square target drifting 80°/s through the receptive field of CSTMD1 (top) and BSTMD2 (bottom). Both neurons exhibit strong biphasic spiking responses (black bar indicates stimulus timing).* ***C,*** *Receptive fields of CSTMD1 and BSTMD2 as projected onto the dragonfly's eye. Mean receptive field of CSTMD1 (n = 11) to a leftward (top left) and rightward (top right) moving small target and mean receptive field of BSTMD2 (n = 5) to a leftward (bottom left) and rightward (bottom right) small target* ***D,***

Responses of CSTMD1 (n = 8) to bars of varying height (size orthogonal to stimulus motion). CSTMD1 exhibits height tuning optimal for targets of ~1-3°. ***E,*** *Target height tuning as in (D) for BSTMD2 (n = 6). BSTMD2 is tuned for small targets measuring ~0.5-4° with a moderate response to elongated features measuring >10° (target width of 1.5°).* ***F)*** *Plot showing the mean and standard error of stimulus onset (left, CSTMD1, n=8; BSTMD2, n=6) and stimulus offset (right, CSTMD1, n=21; BSTMD2, n=10). BSTMD2 exhibits slow onset and offset-time courses (100s of milliseconds to fully cease response to the target) compared to the faster kinetics of CSTMD1.*

Both STMDs are height tuned (size orthogonal to motion) for moving targets (width of 1.5° parallel to motion), with CSTMD1 preferring 1-3° (Figure 37D) and BSTMD2 preferring targets that are 1-4° (Figure 37E). CSTMD1 responses are weak to high contrast targets taller than 5°. In comparison, BSTMD2 is responsive to these larger targets, though does not return to the peak levels measured for smaller targets. In addition to their differences in size tuning and receptive field shape, CSTMD1 and BSTMD2 have an interesting difference in their response kinetics. Even accounting for maximum activity, CSTMD1 has less latency than BSTMD2 (Figure 37F). BSTMD2 has a sluggish offset compared to CSTMD1, that lasts for several hundred milliseconds (Figure 37G).

BSTMD2 is a centrifugal neuron, with two main arborizations, one in the outer part of the primary lobula and the other in the medial lobula with several other regions of input or output in the lateral midbrain (Figure 38A). BSTMD2 is named for its very large binocular receptive field, with highly direction-selective responses to small targets.

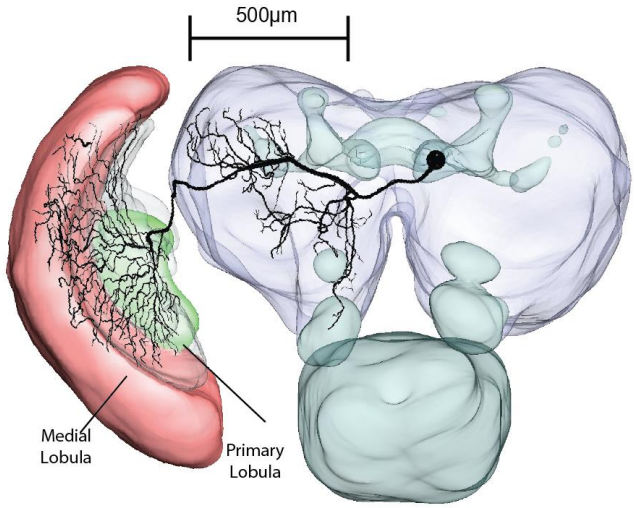

Figure 38: Intracellular labelling of BSTMD2 reveals a centrifugal neuron arborizing widely in the medial and primary lobula.

5.4.2 Selective Attention Across Hemifields in CSTMD1

It has been previously described that CSTMD1 exhibits 'selective attention', responding exclusively to one target, even when presented in a pair (Wiederman & O'Carroll 2013). These experiments tested for selective attention to paried targets within the excitatory hemifield (contralateral) of CSTMD1 (Figure 37C). To investigate the impact of background clutter on CSTMD1 responses requires analyzing how selective attention may operate across the two hemifields. Previous research (Bolzon et al., 2009) found that CSTMD1's possesses an inhibitory hemifield which, though suppressing neuronal activity, has a similar shape and mirror-symmetric direction selectivity to the excitatory hemifield. Bolzon et al., 2009 found that when two targets are presented simultaneously, one in each hemifield, average responses of CSTMD1 are strongly suppressed.

Is it possible that this inhibition is a result of selective attention? To test this, we drifted two targets (1.5° x 1.5°) vertically through CSTMD1's receptive field (Figure 39A), one in the inhibitory hemifield (T_1) and one in the excitatory hemifield (T_2). Each target was presented at a horizontal separation of 50°, well outside any region of binocular overlap (~10°). When presented alone T_1 generates strong inhibition and T_2 strong excitation (Figure 39B), as also evident in the receptive fields shown in Figure 37D. To examine the variability in responses between trials and dragonflies, we presented individual raster plots (Figure 39C). Each point represents an individual spike and the colour changes between red and blue indicate a change in animal (n = 8). The stimulus conditions are shown in the pictograms to the left of the raster plots. The stimulus timings of each target are shown below the rasters (T_1 green, T_2 orange).

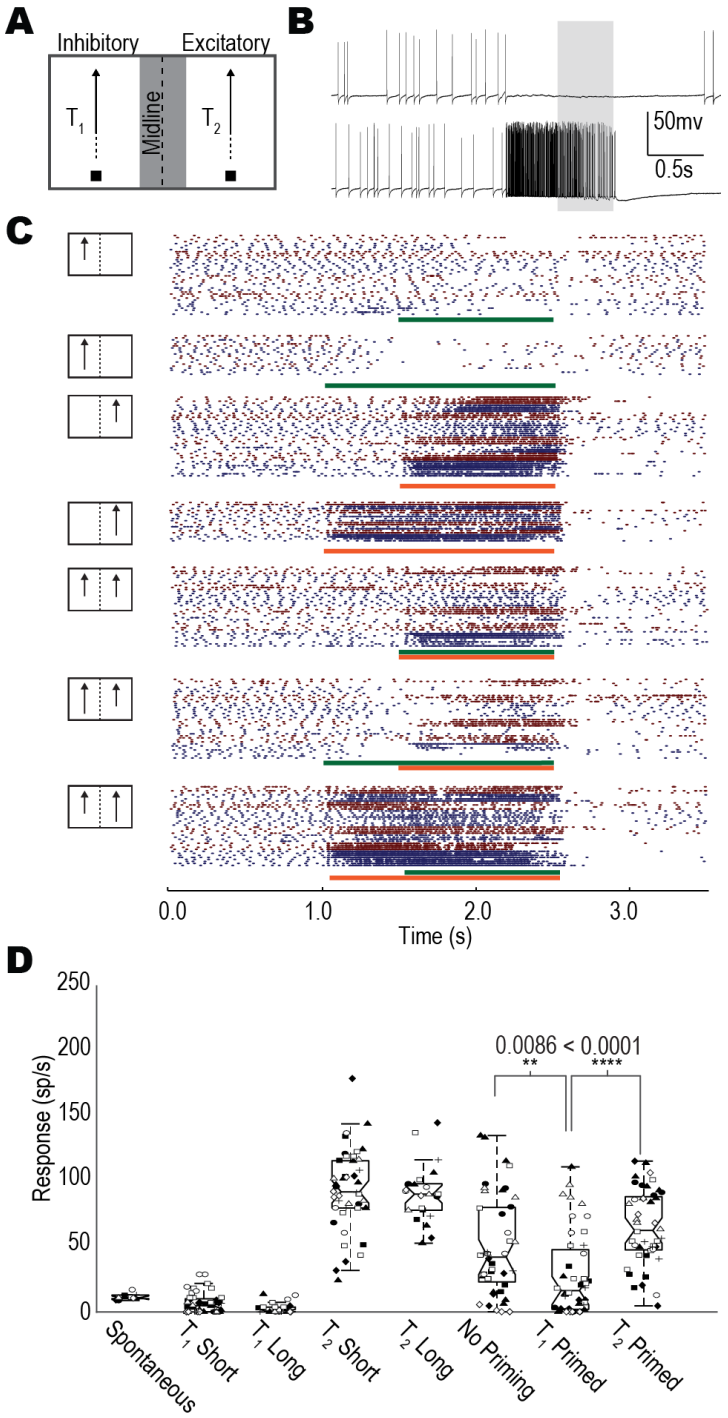

Figure 39: One target is selected when a pair of targets is presented on either side of the visual midline within CSTMD1's receptive field. ***A,*** *Targets are drifted vertically up the stimulus display, either individually or as a pair. T_1 travels through the inhibitory region of CSTMD1's receptive field. T_2 moves through the excitatory receptive field. Targets are separated by 50° to avoid the region of binocular overlap (grey region).* ***B,*** *Example neuronal traces of CSTMD1 responses to T_1 (top) and T_2 (bottom). T_1 generates inhibition, T_2 generates excitation.* ***C,*** *Raster plots over time (each point is a single spike) for single-target and paired target trials of CSTMD1. Green (T_1) and orange (T_2) bars indicat stimulus duration. Pictograms illustrate stimulus locations and length of the target trajectory (short, from 1.5 s to 2.5 s; long, from 1 s to 2.5 s). CSTMD1 trials from different dragonflies are separated by colour changes (n=8). T_1-alone trials are inhibitory, reducing spike rate to*

below spontaneous levels. T_2-alone trials are excitatory, facilitating responses over longer trajectories. CSTMD1 responses to the paired target trials are either T_1-like or T_2-like, indicative of selective attention. This selection was biased by the presence of a preceding 'primer' target (inhibition when T_1 primed, excitation when T_2 primed). ***D,*** *Each point represents the mean neuronal response over a 500 ms window (grey shaded region in B) prior to stimulus cessation (2-2.5s) for an individual trial (excluding spontaneous category). Different dragonflies (neurons) are marked with different symbols. Distributions of this data are presented in boxplots. In all paired target trials there is significant variability with some trials exhibiting below-spontaneous activity (inhibition) and others strong activity (excitation). Priming T_1 results in a shift of the mean response from the no-priming and T_2-priming cases ($p = 0.0086$, $p < 0.0001$, $n = 8$, Kruskal Wallis multiple comparisons with Dunn post-hoc test).*

We presented targets on short (1 s) or long (1.5 s) paths, either in the excitatory or inhibitory hemifield. Targets were either presented individually or as a pair. Either no target was primed (T_1 & T_2 had 1 s duration) or one of the targets preceded the introduction of the other using a 0.5 s primer (T_1 and T_2 respectively), followed by 1 s of the paired stimulus.

In T_1-only trials, CSTMD1 is inhibited by targets presented in the ipsilateral hemifield, which is more pronounced for the longer trajectories (i.e. inhibtion is self-facilitated). T_2-only trials generate excitation, which is stronger with the longer trajectory (i.e. the facilitatory effects observed in predictive gain modulation). Different dragonflies show different overall levels of spiking activity, likely due to slight difference of location within the inhomogeneos receptive field. In paired target trials with no priming, we observed trials that produced strong excitation and others which produced strong inhibition. For the unprimed condition, individual dragonflies preferred either T_1 or T_2, though we found two examples of both excitatory and inhibitory trials in a single animal (Figure 39C). For these unprimed, paired targets, the inhibitory target appeared to win more often than not, similar to the long-range inhibition previously observed in average responses (Bolzon et al., 2009).

When the inhibitory hemifield was primed with a preceding target (T_1 primed), responses during the paired stimuli were more consistently suppressed. Vice versa, paired responses were more strongly excited when primed with T_2. From the rasters plots it is apparent that priming influences which target is selected when simultaneously presented in both hemifields.

To quantify these observations, we counted spikes in a late 0.5 s window commencing 0.5 s after the appearance of the second target (grey region in Figure 39B). This provided sufficient time for each of the pair of targets to facilitate (Nordstrom et al., 2011). We aggregated this response activity across trials and neurons for each of the conditions (Figure 39D), however

we also show individual trials (with varying symbols for the different dragonfly neurons). As expected, the T_1 short and long trajectories alone are inhibitory, whilst the T_2 short and long trajectories elicit excitation. Long T_2 responses are similar to the short, possibly because in this late analysis window a hyperpolarising, inhibitory current is produced by the sustained excitation (a form of motion-after-effect).

The three paired, priming conditions show large variation in their trial-to-trial responses including those similar to an excitatory or inhibitory target alone. Trials in individual neurons can be variable, for example in the no priming condition there are two neurons (▲, □) with both high and low response activity (i.e. selecting either the inhibitory or excitatory target). In many trials, the T_1 priming established strong inhibition, however the biasing was not always effective. In some trials, following strong inhibition to the primer target, responses in the late window switched to the excitatory target, in either all trials (△) or in a subset of trials (□, ○).

However, on aggregate there is a significant inhibitory effect of the T_1 priming compared to the unprimed conditon ($p = 0.0086$, Kruskal-Wallis with multiple comparisons). Individualy, the responses (either excitation or inhibition) are of a similar strength to what would be expected in the unpaired condition, indicative of their selection. CSTMD1 responses to the T_2 priming condition are significantly different than the T_1 priming condition ($p < 0.0001$, Kruskal-Wallis with multiple comparisons), revealing that the target selection can be biased by a preceding primer target. However, although many T_2 primed trials produce excitation, there are again exceptions, with some neurons exhibiting trial-by-trial variability between inhibition and excitation (○, ⧫, □). Therefore, the aggregate inter-hemispheric inhibition observed in previous studies (Bolzon et al., 2009) was likely the average result of more complex inter-trial variability induced by a selective attention network.

5.4.3 Selective Attention across directions in BSTMD2

In BSTMD2, we tested for selective attention with a similar set of conditions as used in CSTMD1, this time using directionality, rather than spatial location as the inhibitory or excitatory drive. (Figure 40A). Here we drifted one target in the inhibitory direction, left to right, (T_1) and one target in the excitatory direction, right to left (T_2). Figure 40B shows individual examples of BSTMD2 responses to targets drifted in both directions. The targets were separated vertically by 10° and the positions pseudo-randomly swapped (i.e. either T_1 or T_2 at the top) to avoid bias from the inhomogenous receptive field. In longer recordings, sets

of trials were located in different parts of the receptive field to avoid local effects of habituation.

The resultant raster plots are shown in Figure 40C, with T_1-alone trials generating inhibition and T_2-alone trials excitation (for both short and long paths). BSTMD2 has an inhomogeneous excitatory receptive field (Figure 37), with stronger responses when a target is presented in the ipsilateral hemifield. Therefore, dependent on the location of the target experiments, the amount of excitation will vary between dragonfly neurons. The unprimed paired target trials resulted in either inhibition or excitation, at strengths comparable to the target alone cases. In some trials, responses changed from the inhibitory target to the excitatory target midway through the trial, correlated with a transition across the visual midline from one eye to the other. During these trials the inhibitory T_1 won against the weaker, excitatory T_2 until T_2 crossed over the midline to the stronger region of the receptive field. Priming with an inhibitory T_1 only induced a small number of additional trials (11%) to be inhibited (no additional excitation trials). Priming with T_2 resulted in more excitatory trials, with effects lasting well beyond the slugish onset dynamics of BSTMD2 (Figure 37F). Alternatively, these changes in response could be due to surround interaction effects, as they often occur near the 2.0 s time point when the paired targets are at their closest point (10°). Like CSTMD1, individual BSTMD2 neurons tended to exhibit consistent preferences for one target or the other, however this was not always the case.

Examining responses over the late 500 ms analysis window (Figure 40D) showed the clear directionality of BSTMD2 single target responses, either inhibited by T_1 or excited by T_2. For the longer paths, the inhibition was enhanced for T_1, and the excitation was enhanced for longer T_2 (i.e. unlike CSTMD1, there is no motion after effect). For the unprimed condition, the distribution of points is not as bimodally distributed as CSTMD1 (the distribution expected for selection), however this is due to the midline switching. For this condition, five individual neurons (⧫, ▲, ■, ●, △) exhibited inter-trial variability, either excitation or inhibition. In the T_1-primed and T_2-primed trials four (⧫, ▲, ●, △) and five (⧫, ▲, ■, ●, △) neurons respectively also exhibited large response variability. Even though hidden in the aggregate data, there was still a significant difference between priming T_1 and priming T_2 ($p = 0.047$, $n = 6$, Kruskal-Wallis with multiple comparisons test), indicative that selective processes were amenable to priming by a preceding target.

The ability to select a target, whilst ignoring the distracter, is therefore not unique to CSTMD1. Interestingly, the mechanism can select inhibitory stimuli, established by either spatial location CSTMD1's inhibitory hemifield) or directionality (BSTMD2's inhibition to rightwards targets). While individual neurons exhibit a bias towards either the excitatory or inhibitory targets (less so in BSTMD2), some also reveal strong inter-trial variability between this selection. These observations provide important insight into how STMD processing may operate in cluttered, visual environments where 'false-targets' may compete for selection.

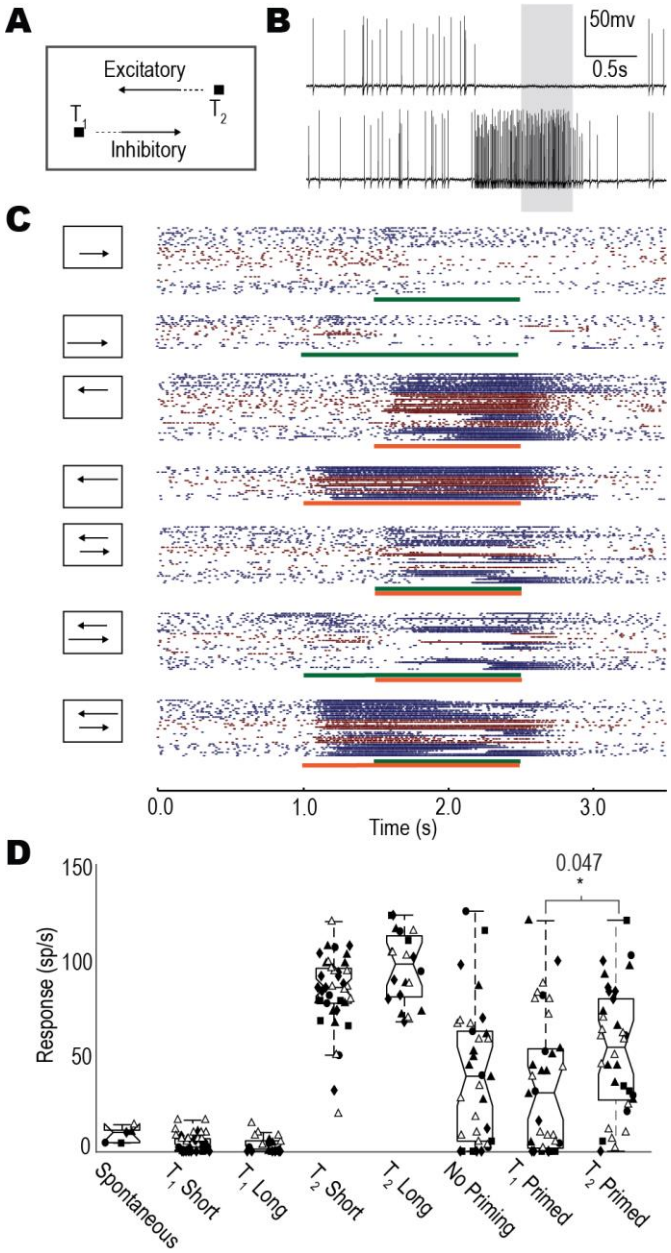

Figure 40: One target is selected when a pair of targets is presented moving in opposing directions within BSTMD2's receptive field ***A,*** *Targets are drifted horizontally across the stimulus display, either individually or as a pair. T_1 travels in the inhibitory direction of BSTMD2's receptive field. T_2 moves in the excitatory direction. Targets are separated by 10°*

to limit surround antagonistic interactions. ***B,*** *Example neuronal traces of BSTMD2 to T_1 (top) and T_2 (bottom). T_1 generates inhibition, T_2 generates excitation.* ***C,*** *Raster plots (each point is a single spike) for single-target and paired target trials of BSTMD2 (blue and red bars indicated stimulus time, T_1 & T_2 respectively, pictograms indicate stimulus locations, different animals are separated by colour, n=6). Paired target trials exhbit both T_1-like and T_2-like responses indicative of a selective attention mechanism. Primed trials shift these responses to the primed target (inhibition T1 primed, excitation T2 primed).* ***D,*** *Boxplots showing different conditions shown in (C). Each point represents the mean response over a window of 500ms (grey shaded region in B) just prior to stimulus cessation (i.e. 2.0-2.5s) for an individual trial (excluding spontaneous estimates which are averaged across neurons). Different animals are marked with different symbols. No statistical difference seen between primed and unprimed trials. T1 primed and T2 primed trials are significantly different ($p = 0.019$, Mann-Whitney U-test) indicating priming is effective in BSTMD2.*

5.4.4 Minimizing Contrast Variation in Natural Images

STMDs are primarily responsive to dark targets moving against light backgrounds (Wiederman et al. 2013). However, a dark target varies in contrast as it crosses light and dark regions of natural scenes. This variation will modulate STMD responses over time due to the neuron's sensitivity to contrast (O'Carroll & Wiederman 2014). Such changes will not only be due to 'local' intensity variation in both space and time (i.e. contrast), but also from longer-term adaptive (Wiederman et al., 2008) and self-facilitatory effects from prior history (Wiederman, Fabian et al., 2017). In addition, STMDs are influenced by spatial inhibitory interactions that underlie their size selectivity. To investigate the effect of background clutter on STMD responses we developed a stimulus that limited these local effects (Figure 41). The inclusion of a grey, horizontal strip along the direction of target travel ensured that 1) local target contrast was constant 2) there was no texture varying the degree of local, temporal adaptation and 3) minimized spatial inhibitory interactions underlying size selectivity.

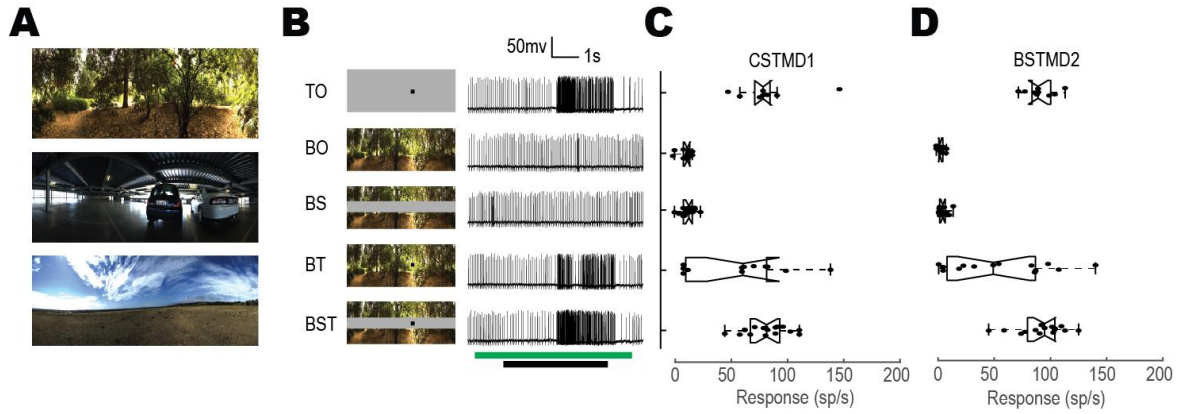

Figure 41: Background effects of contrast variation and size selecitivty on CSTMD1 and BSTMD2 responses are limited by the inclusion of a grey strip. ***A,*** *Three stationary natural images were presented to the dragonfly, chosen for their differering degrees of clutter.* ***B,*** *Each stimulus set was composed of variants, target-alone (TO), background-alone (BO),*

background with strip (BS), background with target (BT) and background with target and strip (BST). Example CSTMD1 responses are shown to each variant of one image in the set. Moving targets (during black stimulus bar) elicit robust reponses compared to the presence of the static background (green stimulus bar). Variation induced by a target moving across the background (BT) is decreased by the presence of the grey strip (BST). ***C,*** *Comparison of the five conditions in CSTMD1 (n=2). Each point represents the mean spike-rate in a 500 ms analysis window (corresponding to the strongest part of the receptive field), with aggregate data shown in boxplots. Variation on the target response induced by the variable background contrast (BT) are reduced by the addition of the grey strip (BST) with responses shifted towards those observed with the target alone (TO) condition.* ***D,*** *This method of deconfounding local contrast and surround antagonism is also effective in BSTMD2 (n=3).*

In a small sample of neurons, we tested the efficacy of the grey strip on three natural images (Figure 41A), with a strip height of 11° (Figure 41B). We varied the alpha channel of the strip to allow it to 'fade' over a further 2° either side, to prevent the introduction of sharp edge artifacts. Here we are not testing the effect of moving backgrounds on target responses, rather the influences of local contrast variation. Therefore, in these experiments the target moves across stationary images.

We presented three images, each with a set of 5 conditions: moving target-alone (TO), stationary background-alone (BO), stationary background with grey strip (BS), stationary background with moving target (BT) and stationary background with grey strip and moving target (BST). The results in both CSMTD1 (Figure 41C) and BSTMD2 (Figure 41D) show that the additional contrast variation introduced by the background caused BT trials to have lower and more varied responses than TO trials for both neurons. However, the inclusion of the grey strip (BST) shifted the responses to those observed in the target-alone condition, revealing the effectiveness of the grey strip. To test this statistically, we first calculated the TO mean response across all trials. Then for each image separately, we segmented the mean TO, individual BTs and individual BSTs into 5 ms bins. We determined the bin errors between mean TO and BT trials, as well as mean TO and BST trials. We performed a one sample t-test to determine whether the errors were non-zero. We found that the error between TO and BT were non-zero for both CSTMD1 ($p = 0.011$) and BSTMD2 ($p = 0.0054$). However, there was no difference between the TO and BST for either CSTMD1 or BSTMD2 ($p > 0.4$). Thus, the grey strip paradigm was sufficient to remove contrast variations and inhibitory effects induced by a stationary background.

5.4.5 Target Tracking with Background Motion

It was previously shown that CSTMD1 can respond robustly to small targets when they move at the same velocity as the background (Wiederman and O'Carroll, 2011). However, as well

as embeding the targets we purposely limited the spatial extent of these backgrounds to reduce the complexity of the analysis. How do STMD neurons respond to large, cluttered environments where relative motion is present?

We tested the same three panoramic images, moving either left to right, or right to left (Figure 42A) at 15°/s with the background commencing at two different phases (i.e. horizontal positions) of the 360° panoramic images (0° and 180°). The target (1.5x1.5°) moved at 25°/s, always in the preferred direction (rightwards for CSTMD1 and leftwards for BSTMD2). The target and background speeds were chosen to elicit robust repsonses, whilst avoiding effects of output saturation. We used five different stimulus conditions (Figure 42B): moving target alone (TO), leftward moving background with strip (BSL), rightward moving background with strip (BSR), moving target with leftward moving background and strip (BSLT) and moving target with rightward moving background and strip (BSRT). Background direction and timing is indicated by green arrows and green stimulus bars. Target direction and timing is indicated by black arrows and black stimulus bars (Figure 42B).

Individual examples of CSTMD1 responses to these five stimulus conditions are shown in Figure 42B. Here we observed a strong response to a target presented alone (TO) and limited responses to a background presented alone in either left or right directions (BSL, BSR). When background and target were presented together, the overall response to the target was reduced. This reduction was larger when the background moved in a rightward direction (BSRT) compared to the leftward direction (BSLT). That is, the reduction in spiking acitvity was larger when the background moved in CSTMD1's preferred direction. Note that in CSTMD1, the panoramic images are simultaneously within both the inhibitory and excitatory regions of the receptive field.

For the individual example of BSTMD2 responses, targets evoke robust activity (TO). However, features in the background alone also induce both excitatory and inhibitory activity dependent on whether the background is moved in the preferred (BSL) or antipreferred direction (BSR). Background features are expected to elicit more responses, because of BSTMD2's broader size tuning (Figure 37E). When both target and background are moved in the preferred direction (BSLT), we observe responses to both the target and background features. Interestingly, when background and target are moved in opposing directions (BSRT), neuronal responses are inhibited for a period of time. However, when the target moves into a more sensitive region of the receptive field, responses rapidly switch to strong excitatory spiking activity.

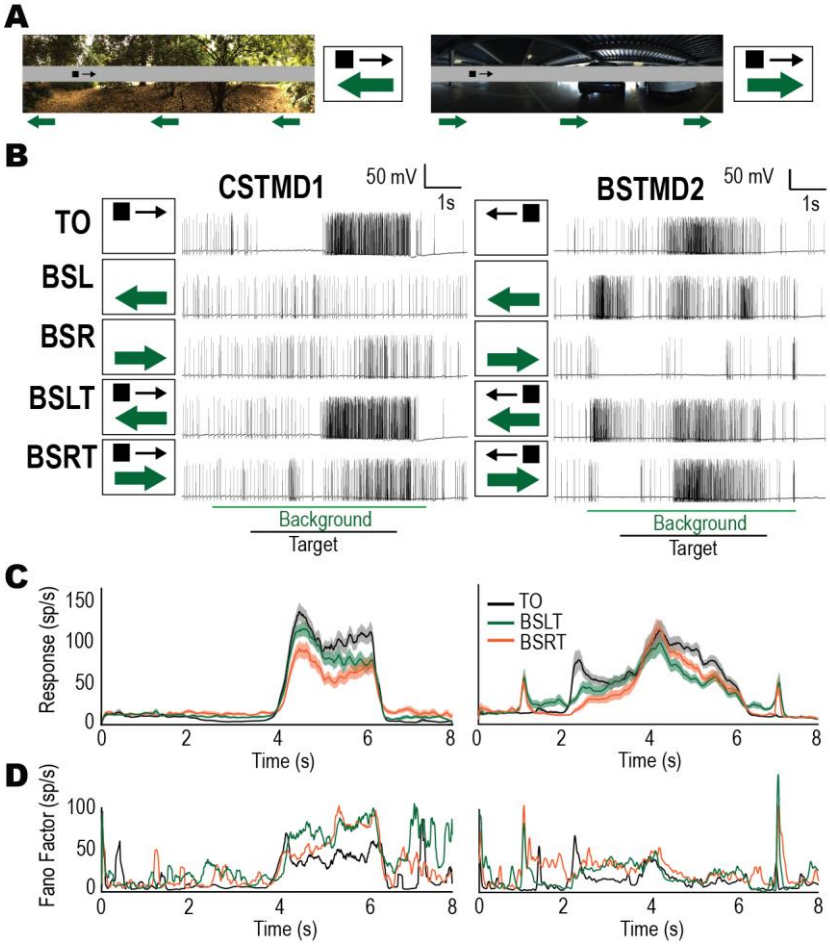

Figure 42: Responses of both CSTMD1 (n=12) and BSTMD2 (n=6) are reduced by backgrounds moving in either direction. ***A,*** *Example of stimuli with either a leftwards or rightwards moving background (green arrows) and rightwards moving target (black arrow). Corresponding pictograms are illustrated.* ***B,*** *Example raw traces from CSTMD1 and BSTMD2 when presented with: target alone (TO), leftward moving background alone (BSL), rightward moving background alone (BSR), target with leftward background (BSLT) and target with rightward background (BSRT). Background features can both elicit activity as well as decrease responses.* ***C,*** *Mean spike rate and standard error (shaded regions) of neurons (CSTMD1 left, BSTMD2 right) to TO, BSLT and BSRT. Targets with leftward or rightward backgrounds result in robust, but reduced spike rates.* ***D,*** *The Fano-factor (variance/mean) calculated from plots in (C). In CSTMD1 and BSTMD2 the Fano factor increases for the two moving backgrounds, revealing that propotionally, variance increases with respect to the mean spike acitivty, a result incongruent with a constant inhibitory drive from a wide field motion pathway.*

For our aggregated analysis, we separated each trial into 5 ms bins and calculated the mean and standard error across all trials (including images and image phase) for each of the five conditions. The three conditions that include the target are shown in Figure 42C (CSTMD1 left, BSTMD2 right). In both background conditions (BSLT green line; BSRT orange line)

responses are reduced compared to the target only (TO, black line) condition. In CSTMD1, this is more pronounced with the BSRT condition having a lower overall response compared to the BSLT condition. Background features will be at different time points across the three images (each with two image phases), therefore average responses should trend towards spontaneous activity.

Are the reductions in mean response to the target when in a background (BSLT, BSRT) due to a consistent inhibitory effect from the background? To test this, we calculated the a metric of variation (Figure 42D), defined as the variance divided by the mean response. This provides a metric for the variation in data relative to its response, for either target alone (black line), target with leftwards background (BSLT, green line), or target with rightwards background (BSRT, orange line). We observe that in both CSTMD1 and BSTMD2 there is an increase in Fano factor with the introduction of the background. Like other insects, the dragonfly widefield system exhibits sustained, velocity dependent responses to translated natural images. If the reductions observed were the result of an inhibitory feedback from the widefield system, we would expect the variation in response to reduce proportionally to the overall response, causing the variation metric to remain roughly constant. Instead, we see a marked increase in variation metric indicating that the presence of the background is increasing the variation. While this does not rule out any interaction with the widefield system, it does demonstrate that another mechanism may be involved in the changes in response.

If instead the target system is selecting either from background features (weakly excitatory) or the target (strongly excitatory), we would expect a larger variation overall, even with a modest reduction in mean response. Thus, the increase in variation metric indicates that selection is likely to be underlying neuronal responses in both CSTMD1 and BSTMD2 .

To illustrate these inter-trial variations, we generated raster plots for each of the individual trials separated by condition. For CSTMD1, we observed robust responses when the target was presented alone (Figure 43A). A target in the left hemifield generates inhibition (2-4 s) and in the right hemifield generates excitation (4-6 s). When presented alone, the background-only stimuli (BSL, BSR) evoked weak and intermittent responses to background features (Figure 43B). When both target and background were presented together (BSLT, BSRT), the response to the target was often still present, but in many cases was shorter in duration, weaker or less consistent (Figure 43C). Interestingly, this reduction is more pronounced when the background was moving in the same direction as the target, which

might be expected if background features are vying for selection in the neuron's weakly excitatory (i.e. non-preferred) direction. The strong inhibition seen in the inhibitory hemifield (2-4s) was also less consistent than in the TO condition with many instances of responses to background features during this period.

Because the inhibitory hemifield of CSTMD1 (ipsilateral) is direction selective for leftward moving targets, we may have expected the BSLT trials to show more consistent inhibition of the target's response compared to the rightwards background (BSRT). Instead, it implies that CSTMD1 may have selected for weakly excitatory features of the background (in the contralateral hemifield) in preference to the target. This is surprising given that the target properties were more salient, with a faster velocity, higher contrast and optimal size. Thus it appears that in background scenery, CSTMD1 sometimes selects weakly excitatory features in preference to the moving target.

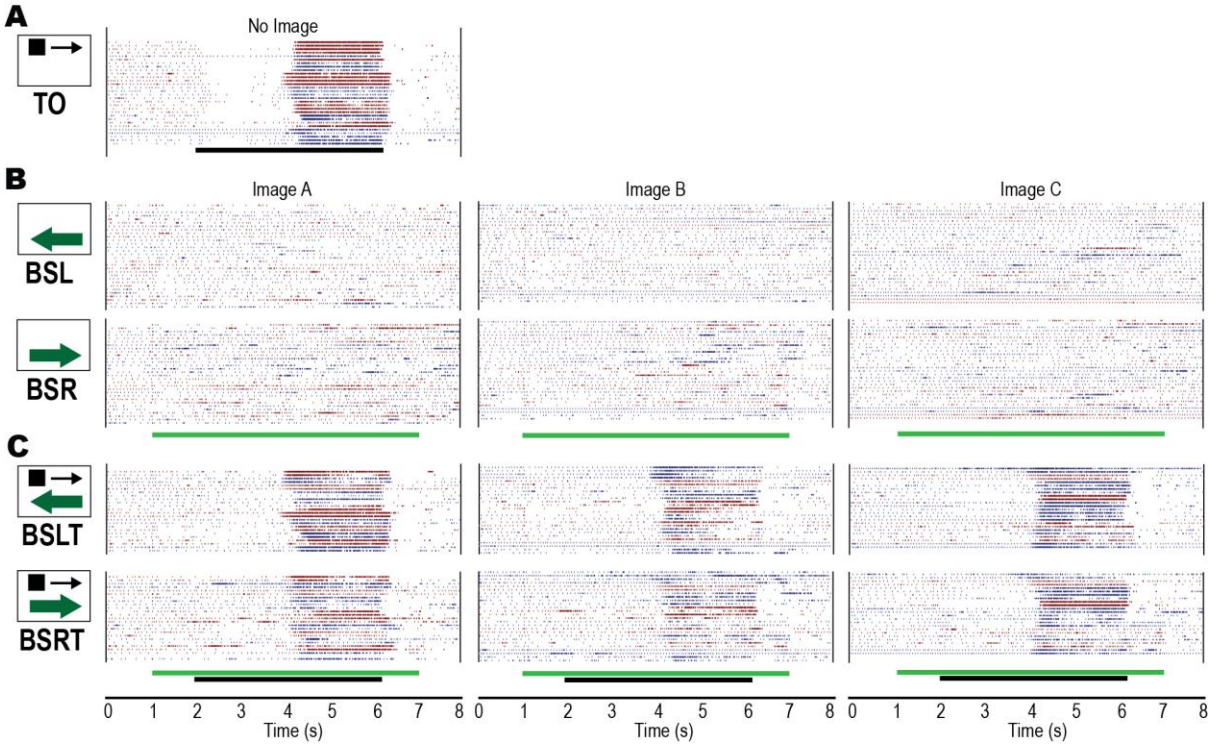

Figure 43: CSTMD1 raster plots reveal that moving backgrounds cause target responses to become less sustained and consistent. ***A,*** *Target only responses (TO). Black stimulus bar indicates time target is moved on the display* ***B,*** *Responses to leftward moving backgrounds (BSL) and rightward moving backgrounds (BSR). Green stimulus bar indicates time background is moved on the display. The moving backgrounds result in intermittent, weak responses.* ***C,*** *Backgrounds moving in either direction, with a target moving in the preferred (BSLT, BSRT), cause a reduction in target response compared to target alone responses.*

We recorded from BSTMD2 less often than CSTMD1 and were not always able to complete the full image set. Figure 44 shows raster plots of the BSTMD2 trials in response to the three

natural scenes, under different stimulus conditions. BSTMD2 responses to the target alone (TO) can last most of the stimulus duration (black bar), indicative of its large receptive field (Figure 44A). For all three images, when backgrounds are presented alone (Figure 44B), BSTMD2 responds to background features directionally, either with excitation (BSL) or inhibition (BSR). Figure 44C shows that responses to the moving background with target trials (BSLT, BSRT) are much less consistent than the TO trials (Figure 44A), which we propose is indicative of the selective processes. For example, there are both excitatory and inhibitory responses from 1 to 2 s (when no target is present), as well as responses to the background feature (e.g. inhibition), when the target is present and should be providing a robust excitatory response. That is, in some trials the neuron responds in an excitatory manner to the presence of the target moving in the preferred direction and in other trials selects a background feature moving in the anti-preferred direction that suppresses responses.

As observed in CSTMD1 responses, when the target and background are moving in the same direction, there is a reduction in BSTMD2 responsiveness to the target. This is despite leftward motion of the background (BSL) increasing activity. In this case, it is perhaps the selection of weakly excitatory features that could account for the reduction in response. One explanation of this selection of the weak over the strong is that the background features can appear earlier than the target. It has been previously shown that switching is an uncommon phenomenon in selection (Wiederman and O'Carroll 2013) and thus the earlier selection of a background feature may allow it to win against the more highly salient target.

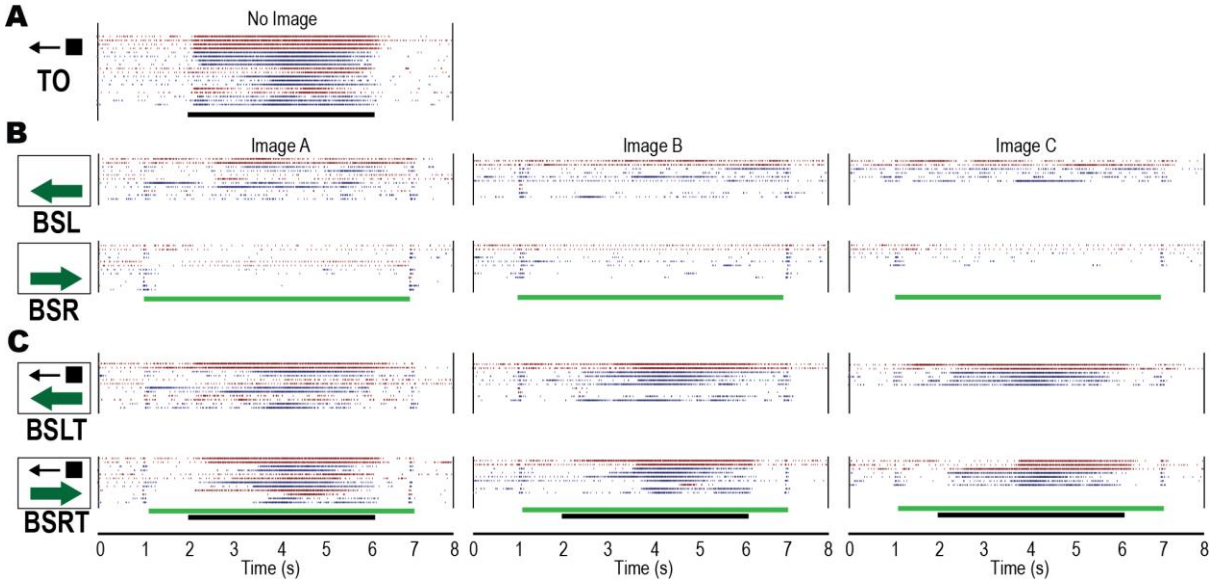

Figure 44: As observed in CSTMD1, BSTMD2 raster plots reveal that moving backgrounds cause target responses to become less sustained and consistent. ***A,*** *Target only responses are robust in the preferred direction.* ***B,*** *Backgrounds moving in the preferred direction produce*

many more 'false-positive' responses than CSTMD1. Background features moved in the anti-preferred direction induce inhibition. ***C,*** *Background and target moved in the preferred direction result in spike activity from different parts of the receptive field. Background and target moved in opposite directions result in distinct target or background responses.*

Previous experiments and modelling suggest that most natural images contain features that can elicit responses in STMD neurons, though not as robustly as embedded targets (Wiederman & O'Carroll 2011). Here strong responses to background features were rare, though this was tested with a slow background velocity (15°/s). Despite this we found several examples where background features robustly stimulated the neuron (Figure 45). Traces of these individual examples are either background alone (green line), target alone (red line) or background with target (black line). The background direction is shown as the green arrow as either right or left. In all these examples the only difference is the presence or absence of the target and background with all other parameters (such as image and image phase) the same. Throughout these examples, an asterisk (*) marks when background with target response is similar to responses to background-only trials. Figure 45A shows an example where introduction of the background completely suppresses CSTMD1's response to the target (compare black and red lines). Here the neuron was responding to background features (*) both before and after the target, indicative of selecting weaker features of the background. Figure 45B and Figure 45C show examples where both the target and a background feature generate strong excitatory responses indicating the system can respond to both. Figure 45D shows an example where the response to both background and target (black) is better matched to the background only (green) than the target only (red) even when the target is in the highly sensitive part of the receptive field. This example clearly illustrates how a neuronal response may be reduced overall (background with target is weaker than target alone) by a feature (from the background), which is excitatory when presented alone. This is again indicative of a selective process. Figure 45E and Figure 45F show responses from the same BSTMD2, with a difference in the background phase. In both trials, responses have transitioned from background features to the target mid-trial. Interestingly, in both trials the target response in the background (black line) is slightly weaker than the target alone (red line). This is unlike our previous descriptions of selective attention in CSTMD1, where once selected, responses are as though any distracters do not exist. Therefore, in BSTMD2 there may be an additional inhibitory effect within the processing.

Figure 45G-J show several examples from BSTMD2 where the background is moving opposite to the target. In these trials the background-alone response is inhibitory (green line)

but wins over the excitatory target response (red line) when the background is presented with a target (black line). Double asterisks (**) mark times when an inhibitory background feature (green line) is selected over an excitatory target (red line) when the neuron is presented with a background and target together (black line).

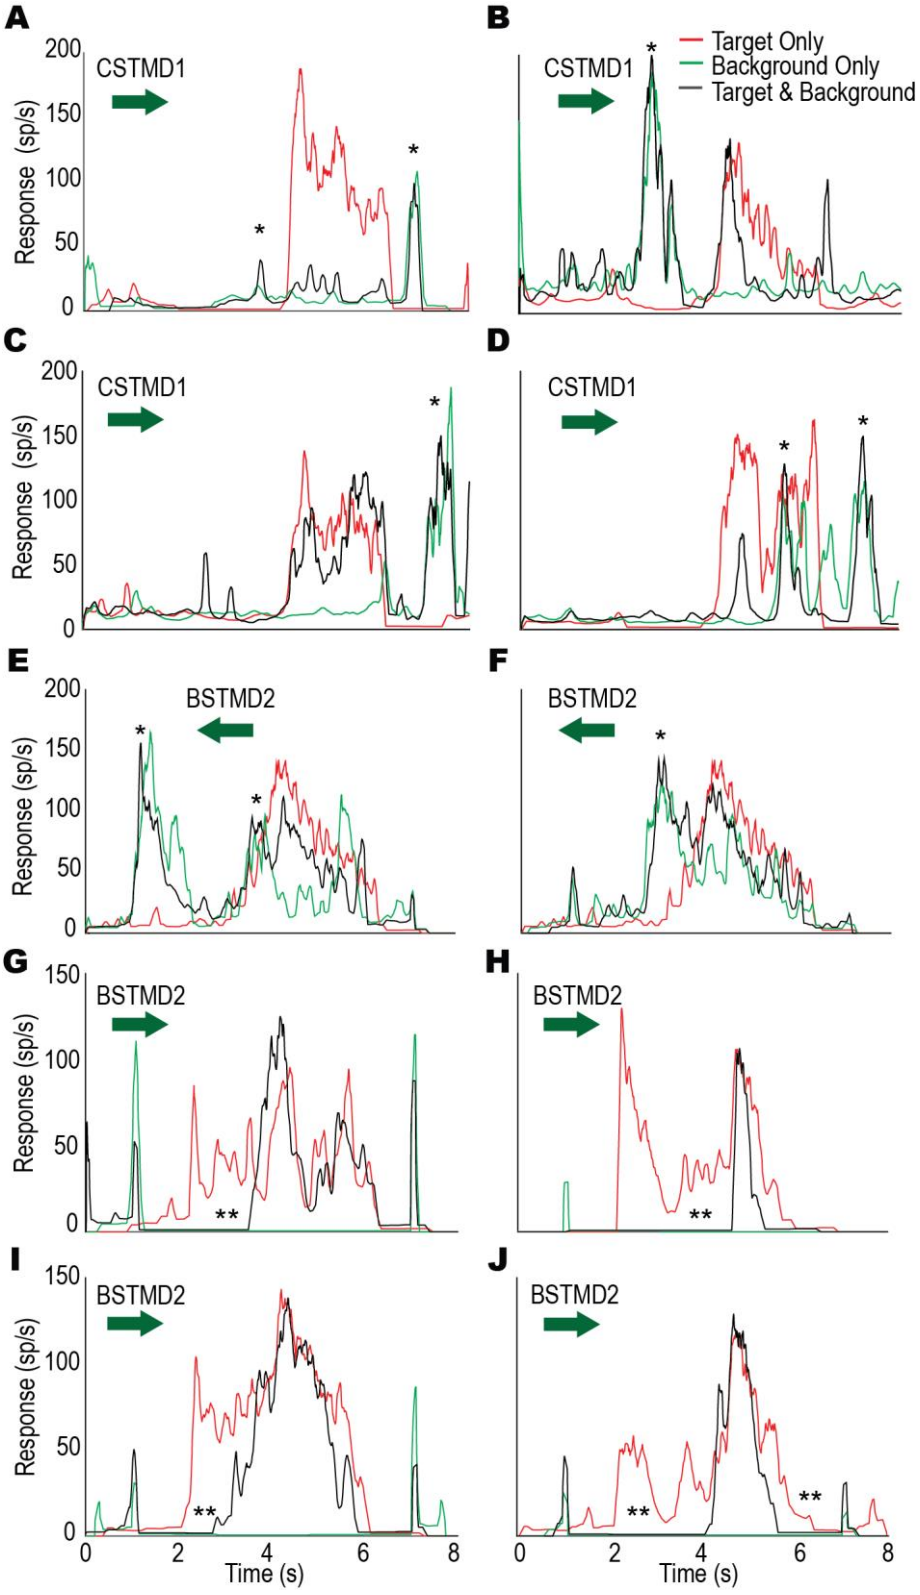

Figure 45: Individual examples comparing target-only, background-only and target and background trials. ***A)*** *An example in CSTMD1 where when the background and target are presented, the responses match the background-alone trial which reveals that sometimes a background feature can win the selection process.* ***B-D)*** *Examples from three cells where the background generates strong responses and causes some suppression of the target response.* ***E-F)*** *Examples from BSTMD2 where the neuron switches between background and target features.* ***G-J)*** *BSTMD2 trials exhibit numerous examples of the neuron being inhibited by background motion despite target features, including switches between background and foreground features.*

5.4.6 Modelling Selective Attention

We applied our modelling efforts to capture the selection between targets and background features (Figure 45) observed in the individual examples in CSTMD1 and BSTMD2. We developed five models of neuronal interactions to compare with the physiological data, a similar approach to our previous investigation of selective attention within the excitatory receptive field (Wiederman & O'Carroll 2013). The models were designed to predict the background with target (BSLT, BSRT) from the combination of the target-alone (TO) and background-alone (BSL, BSR) physiological data (Figure 46A). For each animal (CSTMD1, n=12; BSTMD2, n = 6), we found the mean response of the TO trials. For the background-alone input, we separated background responses for each image, phase and direction (3 images, 2 directions and 2 phases resulting in 12 background-alone averages). Both the mean target-alone and set of mean background-alone were segmented into 5 ms bins and then combined according to the rules of each model variant (6 variants). We calculated an error between this model output and the physiolgical background and target data (matched to the corresponding image, phase and direction).

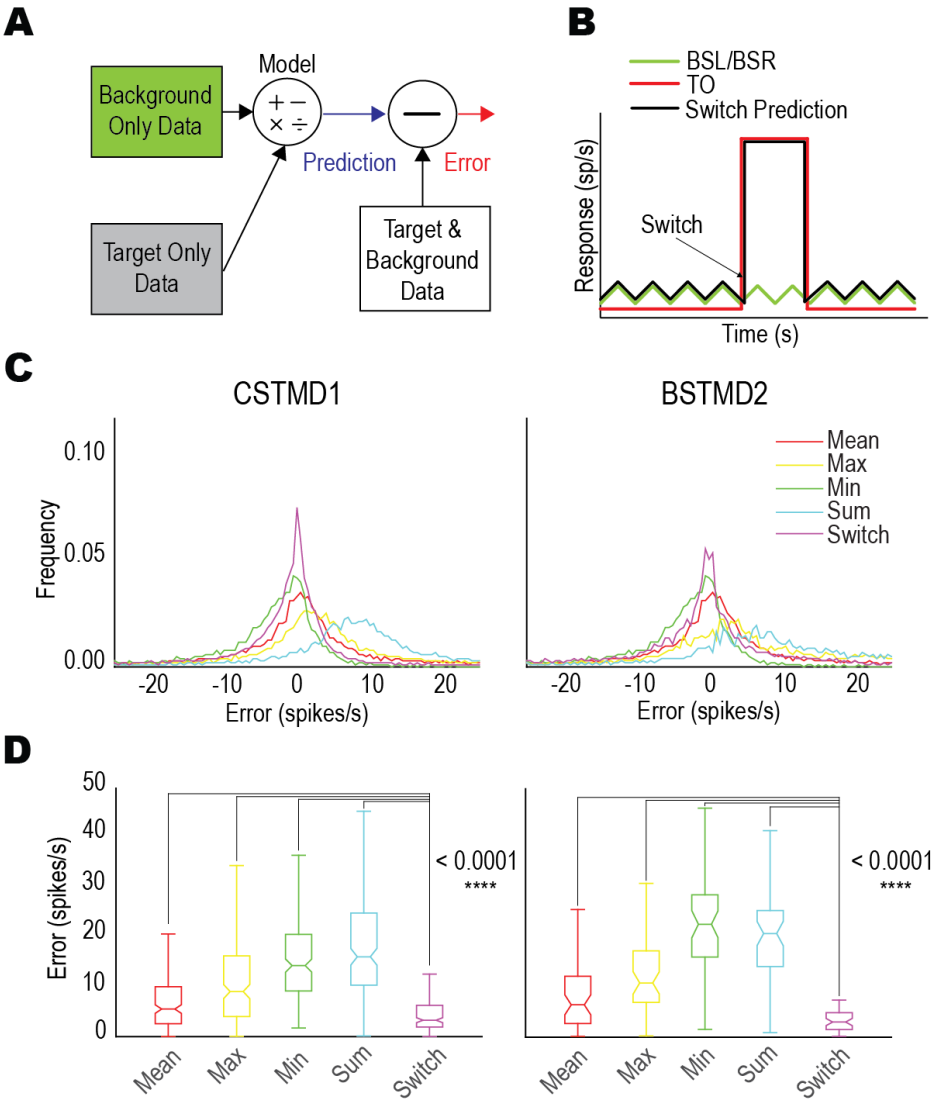

Figure 46: A selective attention model that includes the abilty to switch, outperforms the other model variants. ***A,*** *A schematic of the modelling approach. The target-only (TO) and background-only (BSL, BSR) data are used to predict a background with target response and this is then compared to the background with target data (mean error).* ***B,*** *Illustration of a model variant; Switching (black line) generated from the corresponding TO (red line) and background-alone (green line) responses. The error is the difference between the model output and the physiological background and target responses (not shown).* ***C,*** *Histogram of errors for the model variants Models include Mean (average of background-only and target-only responses), Max (maximum of background-only and target-only responses), Min (minimum of background-only and target-only responses), Sum (combination of background-only and target-only response and Switch (smaller error between target-only and background-only responses).* ***D,*** *Boxplots showing mean error across each trial for each model variant. A smaller value represents a better match between model and data. The Switch model is the best match for both CSTMD1 and BSTMD2 ($p < 0.0001$, t-Test, Bonferroni correction).*

The mean target-alone and mean background-alone responses were combined between each bin over time with the following models; 1) Mean: the average of each bin 2) Max: the maximum value at each bin 3) Min: the minimum value at each bin 4) Sum: the addition of

the values at each bin 5) Switch: the smallest error between the target-alone and background-alone when each was subtracted from background with target (Switch).

A histogram of the errors for each 5 ms bin, across all model variants is shown in Figure 46C. The Switch model resulted in the smallest errors indicated by the high peaks at zero. The Sum and Max models generated large over-estimates of the response while the Min model produced underestimates. We determined the absolute mean error (preventing cancelling of positive and negative errors) for each trial (averaged across 5 ms bins) and treated each trial as an independent error (Figure 46D). The Switch model produced the smallest mean error of the six variants tested. The Switch model was significantly better than all other variants ($p < 0.0001$ T-Test with Boneferroni correction, in all five comparisons). This further supports that background interactions with these two large-field STMD neurons functions through a selective attention mechanism.

5.4.7 Velocity Effects

The previous experiments were conducted at target and background velocities that would induce STMD responses well above threshold, however, would not result in interactions hidden by neuronal saturation. How might a fast-moving background, as would be induced by quick turns during pursuits, affect neuronal responses to a fast or slow-moving target? To investigate this, we chose a single, cluttered background image (Background A, Figure 40A) and recorded responses to targets moving at one of three velocities (15, 35 and 90°/s preferred direction only) whilst the background moved at one of three different velocities (15, 35, and 90°/s, preferred direction only) in CSTMD1. Because faster target velocities evoke responses over a shorter duration of time (Figure 47A) we changed our analysis window to correspond to a constant region of space (50°).

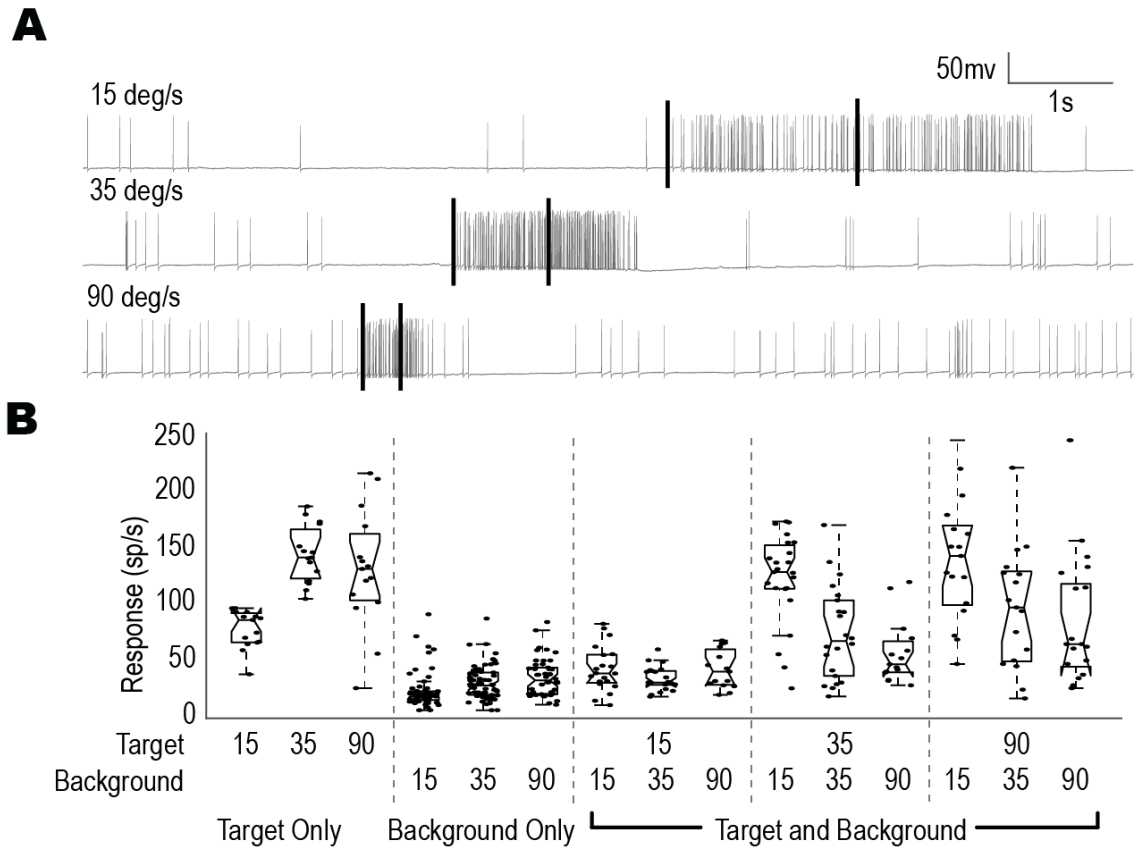

Figure 47: Velocity influences the selective of moving targets. ***A,*** *Three raw traces taken from targets of varying velocities showing the difference in response and time-course. Black bars indicate an anlysis window of equal size (50°).* ***B,*** *Box plot showing CSTMD1 responses over the analysis window, for the image conditions (target only, background only, target and background) across combinations of target and background velocities. This includes three target only velocities, three background only velocities and nine target speed and background speed combinations (all in the preferred direction). Target response is more robust when the target moves as fast, or faster than the background.*

The results in Figure 47B reveals the spiking activity within the analysis window for each background and target speed combination. Target only trials exhibited the expected velocity tuning as we have previously described (Dunbier et al., 2012). Background-only trials produced low responses at all speeds within the analysis window (the same analysis window used in the target-only trials).

The background with target trials showed two important interactions. Firstly, the neuronal response decreased as the background speed increased and secondly the neuronal response increased as the target speed increased. Both these interactions had a statistically significant effect on neuronal response (two-way ANOVA, n=11; background speed, $p < 0.0001$; target speed, $p < 0.0001$). The latter of these two relationships is expected due to the velocity-tuning of CSTMD1. To test whether the target speed had an effect on response independent of velocity tuning, we normalized the target with background trial responses by dividing their

response by the mean response of the corresponding target-only trials (for example all trials with the target moving 90°/s had their responses divided by the mean of the 90°/s target-only trials). This transformed the background with target responses into unitless proportions of the target-only data (between zero and one). Even with this normalization, the two relationships were maintained (two-way ANOVA, n=11; background speed, $p < 0.0001$; target speed, $p = 0.0004$), indicating this interaction was not merely velocity tuning.

How does this interaction occur? If our hypothesis that selective attention explains the interactions between background and target this data would indicate that faster targets or slower background result in more selection for the target compared to the background features, rather than a constant offset in the population response. In all the background with target conditions there are responses that match the corresponding target-only distribution and others which are matched to the background-only distribution with relatively few points half-way in between. This selective trend is particularly visible in the background with target trials with a background speed of 90°/s (3rd box-plot of each set, Figure 47B). All three show a marked division between responses around 50 spikes/s and response greater than 100 spikes/s. In the slow target case (15°/s) all responses are clustered around 50 spikes/s. In the medium target case (35°/s) most trials match the slow target, but with two outliers greater than 100 spikes/s. In the fast target case (90°/s) there is still a cluster of responses centred around 50 spikes/s but now there are many more outliers above 100 spikes/s (7). Not all conditions exhibited as strongly a bimodal distribution as the 90°/s-background cases, however the presence of switching seen in Figure 45 may help explain these intermediate responses.

The background with target trials demonstrate that both target speed and background speed have an important effect on the selection process of CSTMD1. Both relationships indicate that velocity is a key determinant in selection with faster features (whether strongly salient targets or weakly salient distracters) being preferentially selected over slower features independently of the stronger response to faster targets (i.e. velocity tuning).

5.4.8 Detection Error Trade-Off Performance

In behavioural contexts, eventually all detection resolves down to a simple binary question: initiate pursuit or not. How (and whether) STMDs are involved in this behaviour remains unknown, but assuming they play a role, how might the performance of STMDs in clutter affect this kind of decision making? While relating the spike-rate of these neurons to these behaviours remains infeasible it is still possible to examine how clutter might affect the

detection capability of STMDs towards these ends. At what point is a salient target indistinguishable from a moving background feature?

To examine this question, we used an extended set of our moving background data (Figure 42) in which the target contrast was also varied. There are several logical inferences which can be made from trials. For example, in background-only trials, there is no target and thus any strong response should be interpreted as a false-positive (FP). Likewise when a salient target is presented alone with no neuronal response, this represents a false negative (FN). When both target and background are presented it is impossible to categorically distinguish between target and background responses and thus positive responses remain ambiguous, however the lack of a response when a target is present is also a FN. A summary of this logic is shown in Figure 48A. A useful metric that can be derived from this data is the Detection Error Trade-Off (DET), which compares FP and FN events and is a measure of tracking performance. The advantage of the DET, is that it does not assume a fixed threshold of detection but instead shows how changes in threshold affect the prevalence of FP and FN events.

To measure the DET rate, we first had to establish the definitions of FP and FN events. For this analysis we separated target-only trials and background with target trials into two separate categories and performed the following analysis on both datasets. To calculate the FN rate, we used a 500ms window for each trial taken from the most sensitive region of the receptive field (adjacent to the midline). Thus the FN would depend on the chosen threshold, with neuronal responses below threshold being recorded as FNs.

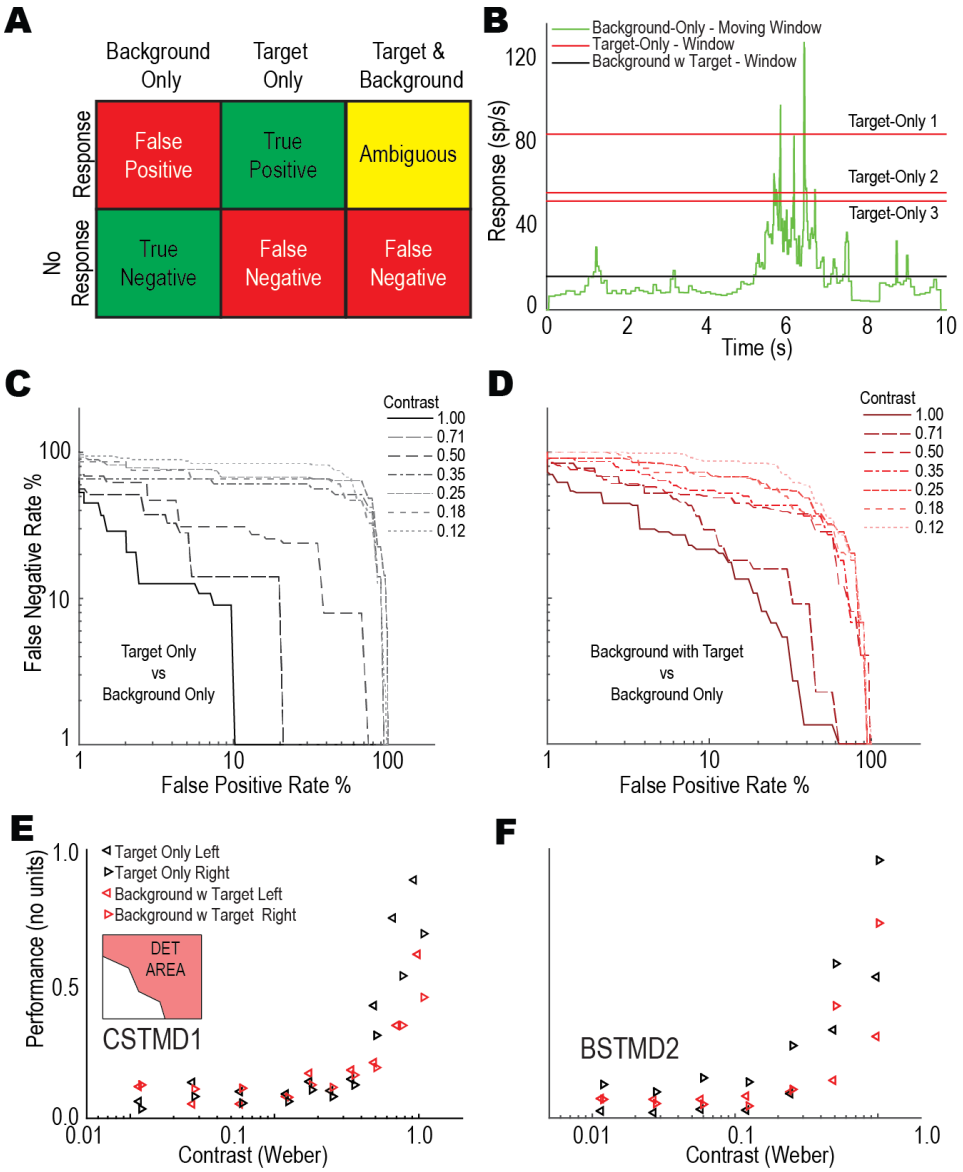

Figure 48: Performance of STMD neurons reduced in clutter regardless of detection threshold selection. ***A,*** *Table showing deductions that can be made based on a threshold-based response mechanism. Strong responses in target with background trials have an ambiguous interpretation.* ***B,*** *Example of a rolling-window measuring background response over the course of an entire background-only trial (green line) with the windowed responses of target-only controls and background with target responses. Regardless of chosen response threshold, it is impossible to achieve zero FP and zero FN simultaneously.* ***C,*** *Detection Error Tradeoff curves for CSTMD1 for all measured contrasts to target-only trials.* ***D,*** *Detection Error Tradeoff curves for CSTMD1 for all measured contrasts to background with target trials. Background with target trials show worse performance than target-only trials regardless of threshold selection.* ***E, F,*** *Area above logarithmic Detection Error Tradeoff curve for all contrasts separating left and rightwards moving backgrounds for CSTMD1 (E) and BSTMD2 (F). Both neurons show more reliable performance when the background moves against the preferred direction of the STMD.*

To generate a FP rate we aggregated the background-only trials. As these trials contained no target, the choice of which section of the trial to examine was arbitrary (i.e. not limited to

where the target would be present in other conditions). As such we calculated windows along the length of the entire trial. Using a sliding window method (window of 500ms) we generated an instantaneous snapshot of the background trials at each time point (Figure 48B – green line). Using this aggregated data, it was now possible to calculate a false-positive rate for any given threshold by simply calculating what percentage of each trial was above the given threshold.

To calculate the DET curve for target-only trials, we varied the threshold calculating the total number of target-only trials which failed to meet threshold while simultaneously calculating the FP rate from the background-only data. An example of how this process works is shown in Figure 48B. Here there are three target-only controls and a single background with target example. If one changes the theoretical threshold from zero spikes/s up through 120 spikes/s, the total number of FPs (where the green line exceeds the threshold) declines while the FN rate would increase as the threshold exceeds the target-only controls. In this example, it is impossible to eliminate all FNs and FPs simultaneously for either the target-only data or background with target data.

The results of these calculations for target-only trials are shown in Figure 48C and for background with target responses in Figure 48D (curves generated by aggregating all images and background directions – CSTMD1). This representation gives a more nuanced reflection of the performance reduction as the contrast of the target is reduced (movement towards the top right, i.e. large probabilities of FP and FN events). It also demonstrates that the introduction of clutter (Figure 48D) increases the rate of FP and FN events regardless of the chosen threshold and across different target contrasts. Interestingly, at very low contrasts, the target with background trials performed slightly better (top right corner of both plots). This is likely due to FPs from background clutter reducing the FN rate.

We then sought to simplify this data further by calculating the area above the curve and normalizing it such that a value of 1 represented 0% FP/FN and 0 represented 100% FP/FN. The resulting DET Areas (separated by background direction) are plotted in Figure 48E (CSTMD1) and Figure 48F (BSTMD2).

The DET Area demonstrates that apart from the contribution of contrast (lower contrast reducing performance), the presence of clutter always reduces the DET (red is uniformly lower than black at high contrast).

As the TO trials were calculated without clutter present, this difference indicates a target-clutter inhibitory interaction. Additionally, there is an apparent directional difference in both

cell types (both showing poorer performance when the background moves in the preferred direction). This is best explained by a much higher false-positive rate introduced by preferred-motion features in the background. This indicates that both CSTMD1 and BSTMD2 are more reliable target detectors when there is opposing motion within a scene, despite previous work showing relative motion is not required (Wiederman & O'Carroll 2011).

5.5 Discussion

It is a natural intuition that the presence of cluttered scenery should reduce the performance of any target-tracking system simply because of changes in local contrast and potential distracting features such a background inevitably introduces. When this contrast variation is accounted for, one might expect a return to robust responses. Alternatively, it might be possible to replace the strong target responses with equally strong false-positive responses (or inhibitory responses) generated by background features. However, what we have shown is a combination of reduced responses tracking weakly excitatory background features (rather than an inhibitory interaction), inhibitory responses (likewise from background features) or completely unchanged strong responses to the target.

Previous research has shown that robust responses to background features are rare (Wiederman & O'Carroll 2011), indicating that few background features strongly match the finely tuned selectivity of STMD neurons. However, matching the fine-tuning of an STMD isn't a prerequisite for selection. Here we have shown numerous individual examples of STMDs preferring weakly excitatory features with poorly matched velocity and size tuning trumping targets crafted to give strong responses. This results in a general reduction in successful target detection across a range of background and target velocities. We have further shown that being an excitatory stimulus is not necessary to the selection mechanism and that inhibitory stimuli are also competing for attention with their excitatory counterparts.

The fact that background features and inhibitory features can both be selected also indicates that the selective attention mechanism may be partially detached from the parameters of STMD neurons. There is no evidence yet to suggest that small (strongly salient) targets are any more likely to win attention than large (weakly salient) targets. It is entirely possible that large objects will win the selective attention game more commonly than small ones. We have however shown that velocity remains a key factor in determining the winner of the selection process in natural image contexts. Intuitively, faster moving features are more likely to win attention even when their higher general salience in accounted for.

Switches have been shown to be rare (Wiederman & O'Carroll 2013) which may help explain how background features can ever win neuronal attention. Processes such as facilitation (Wiederman & Fabian et al., 2017) result in stronger target responses on predictable forward trajectories while causing inhibition in other parts of the receptive field. As many background features have the advantage of appearing prior to the presented target in our experiment, it is possible that facilitation effects help suppress the response to the target sufficiently to allow a lower-salience background feature to remain the most salient feature and win the selective attention mechanism. However, the numerous examples (Figure 45) where the target response rapidly drops away indicating a switch onto background features argues against this as the only explanation.

It was previously shown that the response to targets was strong despite a lack of relative motion (Wiederman & O'Carroll 2011). Here we have shown that relative motion isn't beneficial to detection in terms of net spike-rate when the background and target moved in opposing directions. However, in BSTMD2, while the absolute response reduces, the relative salience may have been improved in many cases, where the response to anti-parallel background features induced inhibition while the target produced only strong excitation. How this might be used as part of a target tracking network remains unclear. While it is usually easier to track objects with differing velocity, STMDs appear only to suffer from extraneous background motion. Even in cases where neurons responded to the target, there were no noticeable increases in performance indicating that STMD neurons function almost entirely agnostic to relative motion cues. This suggests that pop-out effects might be purely due to inhibitory interactions (such as caused by facilitation). It is also possible that during closed-loop flight efference copies (Kim et al., 2015) may affect the responses to STMDs negating some of the performance degradation observed here.

Our two-target experiments also provide evidence that selective attention works across hemispheres. Whether each hemisphere generates its own selection system remains unclear, but the two systems certainly interact further clarifying long-range inhibitory effects previously described (Bolzon et al., 2009). They also clarify the underlying mechanism of selective attention. As CSTMD1/BSTMD2 are capable of preferentially tracking inhibitory stimuli, the selective attention mechanism must go beyond mere inhibitory modulation. While this inhibitory modulation may play a role in selection, inhibitory responses cannot be used to power an inhibitory network simultaneous to excitatory responses.

Finally, in a behavioural context, our findings raise a few questions. Firstly, in prey pursuit, dragonflies have been shown to minimize the relative slip of the target (Mischiati & Lin 2014), essentially aiming for zero velocity on the eye. If velocity is a key factor in selective attention, it indicates that this cannot be in play (or must function differently) during pursuit as all clutter features would have a higher retinal velocity than the target. This indicates that STMD neurons may have a more important role in target detection rather than tracking. These findings also confirm the standard behaviour of dragonflies, which tend to hawk (or perch) to minimize background motion while choosing locations where regions of clear sky are visible to maximize the contrast of potential prey.

Of course, dragonflies do not always choose their engagements. Conspecific encounters over territory can often commence with the competitor hidden in front of cluttered background. In these circumstances, and the complex fighting flights that follow, STMDs are still capable of generating robust responses to moving targets assuming good contrast.

5.6 References

1) Barnett PD, Nordstrom K & O'Carroll DC, 2007, Retinotopic organization of Small-Field-Target-Detecting Neurons in the Insect Visual System, Current Biology (vol. 17, pp. 569-578), https://doi.org/10.1016/j.cub.2007.02.039
2) Brinkworth RSA & O'Carroll DC, 2009, *Robust Models for Optic Flow Coding in Natural Scenes Inspired by Insect Biology*, PloS Computational Biology (vol. 5, no. 11, e1000555), https://doi.org/10.1371/journal.pcbi.1000555
3) Collett TS, 1971, Visual Neurones for Tracking Moving Targets, Nature (vol. 232, pp. 127-130), https://doi.org/10.1038/232127a0
4) Collett TS & Land MF, 1978, How Hoverflies Compute Interception Courses, Journal of Comparative Physiology, (vol. 125, pp. 191-204).
5) Dunbier JR, Wiederman SD, Shoemaker PA & O'Carroll DC, 2012, Facilitation of dragonfly target-detecting neurons by slow moving features on continuous paths, Frontiers in Neural Circuits (vol. 6, 79) https://doi.org/10.3389/fncir.2012.00079
6) Farrow K, Haag J & Borst A (2003) Input organization of multifunctional motion-sensitive neurons in the blowfly J. Neuroscience (vol. 23, pp. 9805-9811)
7) Field DJ, 1987, Relations between the statistics of natural images and the response properties of cortical cells, Journal of the Optical Society of America A (vol. 4, no. 12, pp. 2379-2394) https://doi.org/10.1364/JOSAA.4.002379

8) Geurten BRH, Nordstrom K, Sprayberry JDH, Bolzon DM & O'Carroll DC, 2007, Neural mechanisms underlying target detection in a dragonfly centrifugal neuron, The Journal of Experimental Biology (vol. 210, pp. 3277-3284)

9) Kim AJ, Fitzgerald JK & Maimon G, 2015, *Cellular evidence for efference copy in Drosophila visuomotor processing*, Nature Neuroscience (vol. 18, no. 9, pp. 1247-1255), https://doi.org/10.1038/nn.4083

10) Mischiati M, Lin HT, Herold P, Imler E, Olberg R & Leonardo A, 2015, *Internal models direct dragonfly interception steering*, Nature (vol. 517, pp. 333-338), https://doi.org10.1038/nature14045

11) Nordstrom K, Barnett PD & O'Carroll DC, 2006, Insect Detection of Small Targets Moving in Visual Clutter, PLoS Biology (vol. 4, no. 3, e54)

12) O'Carroll DC, 1993, Feature-detecting neurons in dragonflies, Nature (vol. 362, pp. 541-543)

13) O'Carroll DC & Wiederman SD, 2014, Contrast sensitivity and the detection of moving patterns and features, Philosophical Transactions of The Royal Society B (369: 20130043) https://doi.org/10.1098/rstb.2013.0043

14) Warzecha AK, Egelaaf M & Borst A (1993) Neural circuit tuning in fly visual interneurons to motion of small objects. I. Dissection of the circuit by pharmacological and photoinactivation techniques. J Neurophysiology (vol. 69, pp. 329-339)

15) Wiederman SD, Fabian JM, Dunbier JR & O'Carroll DC, 2017, A predictive focus of gain modulation encodes target trajectories in insect vision, eLife, (vol. 6, e26478),

16) Wiederman SD & O'Carroll DC, 2011, Discrimination of Features in Natural Scenes by a Dragonfly Neuron, Journal of Neuroscience (vol. 31, no. 19, pp. 7141-7144)

17) Wiederman SD & O'Carroll DC, 2013, Selective Attention in an Insect Visual Neuron, Current Biology (vol. 23, pp. 156-161), http://dx.doi.org/10.1016/j.cub.2012.11.048

18) Wiederman SD, Shoemaker PA & O'Carroll DC, 2008, *A model for the detection of moving targets in visual clutter inspired by insect physiology*. PLoS One (3:e2784), https://doi.org/10.1371/journal.pone.0002784

19) Wiederman SD, Shoemaker PA & O'Carroll DC, 2013, *Correlation between OFF and ON Channels Underlies Dark Target Selectivity in an Insect Visual Systems*,

The Journal of Neuroscience (vol. 33, no. 32, pp. 13225-13232)
https://10.1523/JNEUROSCI.1277-13.2013

6 Insect Target Tracking Neuron Locks on to Attended Targets

6.1 Preamble

As in many laboratories, where one project finishes and another starts can be quite arbitrary. Often these boundaries are heavily blurred. So is the case with dragonfly vision in natural images and the broader question of selective attention. With another student formally investigating selective attention within our laboratory, I was somewhat constrained in the kinds of experiments I could investigate to test the effects of selective attention in natural scenes.

My original formal contribution to this proto-paper was slight. I had successfully convinced a wide-field neuron to respond to two modestly sized targets to demonstrate that selective attention was a 'special' property of STMDs. Though this style of experiment was conducted in numerous neurons, only a single example deigned to produce sufficient response for proper analysis. However, it would be incorrect to call this my only contribution.

Since the inception of this project, I was informally involved, both as author of the code to generate the frequency-tagging paradigm used in this paper and in the proposed analysis. For a lengthy portion of this project, I had a hand in the analysis and presentation (though perhaps informally). It was only when one of the experiments proposed for the paper didn't reach the 'publishability' threshold that my involvement spiked. To supplement the work, I produced a modest modelling effort to describe the high-level behaviour of selective attention to supplement my wide-field neuron contribution. Having attained authorship, I then went on to become more tightly involved in the presentation and analysis, replicating the analysis results and advising on presentation and analysis matters throughout the paper. A curious quirk of the challenges of electrophysiology also results in much informal 'sharing of data', when one researcher has a particularly stable cell, many researchers run experiments on the same animal. As such, I have also been involved in much of the data collection. I have also acted as author/reviewer for much of the paper including parts I did not explicitly author. I would however, like to point out that this remains largely the work of Benjamin Lancer who takes much of the credit for experimental design, analysis and certainly the extensive background literature.

As this proto-paper also directly relates to experiments conducted in my previous proto-paper (particularly priming experiments), which were conducted around the same time, it seems a sensible inclusion in this thesis.

Statement of Authorship

Title of Paper	Insect Target Tracking Neuron Locks on to Attended Targets
Publication Status	Submitted to BioRxiv
Publication Details	Lancer BH, Evans BJE, Fabian JM, O'Carroll DC & Wiederman SD, unpublished

Co-Author

Name of Principal Author	Bernard Evans		
Contribution to the Paper	Experiment conceptualisation, animal collection, data collection, data analysis, data interpretation, figure generation, modelling, manuscript authorship.		
Overall percentage (%)	25%		
Certification	This paper reports on original research I conducted during the period of my Higher Degree by Research candidature and is not subject to any obligations or contractual agreements with a third party that would constrain its inclusion in this thesis. I am the second author of this paper.		
Signature		Date	14/12/18

Co-Author Contributions

By signing the Statement of Authorship, each author certifies that:

i. the candidate's stated contribution to the publication is accurate (as detailed above);
ii. permission is granted for the candidate in include the publication in the thesis; and
iii. the sum of all co-author contributions is equal to 100% less the candidate's stated contribution.

Name of Principle Author	Benjamin H Lancer		
Contribution to the Paper	Experiment conceptualisation, animal collection, data collection, data analysis, data interpretation, figure generation, manuscript authorship.		
Signature		Date	14/12/18

Name of Co-Author	Joseph M Fabian		
Contribution to the Paper	Experiment conceptualisation, data collection, data analysis, figure generation, manuscript evaluation		
Signature		Date	28/11/18

Name of Co-Author	David C O'Carroll		
Contribution to the Paper	Experiment conceptualisation, data interpretation, manuscript evaluation		
Signature		Date	4/12/2018

Name of Co-Author	Steven D Wiederman		
Contribution to the Paper	Experiment conceptualisation, data interpretation, manuscript evaluation		
Signature		Date	14/12/18

6.2 Introduction

The visual world contains a wealth of information about the environment and surroundings, but even the most sophisticated visual systems lack the capacity to encode all the information contained in a scene over time. Instead, animals must parse a scene for behaviourally relevant information and discard the remaining clutter. One solution to this problem is *selective attention*, the ability to selectively respond to one stimulus amongst multiple alternatives. Selective attention is observed across species, from humans and other primates (Treue, 2001), to 'simple' insects, including the fruit fly (De Bivort & van Swinderen, 2016; Nityananda 2016). Selective attention is particularly important in predatory animals that hunt among swarms containing potentially hundreds of prey and conspecifics, such as the dragonfly (Edman & Haeger, 1974; Baird & May, 1997). Many predators hunting in these conditions are susceptible to the 'confusion effect', a reduced success rate due to difficulty tracking a single target amidst the swarm (Landeau & Terborgh, 1986, Jeschke & Tollrian, 2007). Some dragonflies, however, show particularly good performance hunting among swarms across all stages of life (Jeschke & Tollrian, 2007; Combes et al, 2012).

Successful prey capture relies on the ability to filter irrelevant information, such as background clutter and conspecifics, whilst selecting and tracking prey amongst equally valuable alternatives. Indeed, the confusion effect is diminished where predators are able to visually identify and track individual prey (Landeau & Terborgh, 1986). In order to achieve this, the underlying neuronal system must be able to 'lock-on' to and track an individual, noisy target, while simultaneously flexible enough to switch targets when this would increase the chance of success.

Wiederman and O'Carroll (2013) previously identified a visual neuron in the dragonfly optic lobe that exhibits a 'winner-takes-all' selective attention (Wiederman & O'Carroll, 2013), named 'Centrifugal Small-Target Motion Detector' (CSTMD1). CSTMD1 is tuned for the movement of small (1°-3°) dark targets against a bright background (O'Carroll 1993; Geurten et al, 2007), matching the demands of an ethologically relevant target-detection system (Labhart & Nilsson, 1994; Olberg et al, 2005; Olberg et al., 2007). When presented with two such targets, CSTMD1 encodes the *absolute* strength of the selected target without interference from distracters (Wiedermen & O'Carroll, 2013). In contrast, typical findings in primates (eg. Recanzone, Wurtz & Schwarz, 1997; Treue & Maunsell, 1999), Owls (Asadollahi, Mysore & Knudsen, 2010) and other insects (Tang & Juusola, 2010; van Swinderen, 2012) show a response that is modulated by the presence of non-selected

distracters. Encoding an absolute representation of a selected target (i.e. ignoring the distracter) has been observed in the auditory system of crickets (Pollack, 1998) and in primate neurons in MT (Harrison et al., 2013). The analogue of CSTMD1 processing in human psychophysics is 'inattentional blindness', whereby an object in the visual field is ignored while attention is focused elsewhere (Simons & Chabris, 1999).

Previously, it has been shown that CSTMD1 exhibits properties important for a prey-tracking system. Firstly, the rare observation that selection could switch between targets mid-way through a trial (Wiederman & O'Carroll, 2013). This raised the intriguing possibility that an ongoing competitive mechanism drives target selection, even after an initial target has been selected, and that this mechanism can direct switches at opportune moments. Secondly, CSTMD1 exhibits 'predictive gain modulation' whereby a local faciltatory 'spotlight' of increased gain spreads forward along the predicted trajectory of a target (even accounting for occlusions), with inhibition elsewhere in the receptive field surround (Dunbier et al, 2012; Wiederman, Fabian et al., 2017). This facilitation may represent a mechanism for 'locking-on' to a selected target, for example, a chosen fruitfly in a swarm.

Here, we have developed a technique to frequency-tag targets by exploiting their contrast dependant response (O'Carroll & Wiederman, 2014), thus permitting us to determine which target has been selected at any moment with improved certainty. We show that CSTMD1 is both able to dynamically switch selected targets mid-trail and lock-on to selected targets, even in the presence of a higher contrast distracter. We therefore describe a neuronal system more complex than the traditionally modelled winner-takes-all framework. This provides important insight into how selective behaviours are implemented by underlying neuronal processing.

6.3 Materials & Methods

6.3.1 Experimental Preparation

We recorded from a total of 26 male, wild-caught *Hemicordulia tau* dragonflies. Dragonflies were stored at 7°c for up to 7 days before experimentation. Dragonflies were warmed and then immobilized to an articulating magnetic stand with a 50/50 wax-rosin mixture. The head was tilted forwards to allow access to the back of the head, and a small hole was dissected in the rear of the head capsule adjacent to the oesophagus to allow visual and physical access to the brain.

We pulled aluminosilicate electrodes (Harvard Apparatus) using a Sutter Instruments P-97 electrode puller, which were filled with a 2M KCl solution. Electrodes were then inserted into the lobula complex using a piezo-electric stepper with a typical resistance of 40-140 MΩ. Intracellular responses were digitised at 5 kHz for offline analysis with MATLAB.

6.3.2 Visual Stimuli

We presented stimuli on high-definition LCD computer monitors (120 – 165 Hz) using a custom-built presentation and data acquisition suite based on MATLAB (RRID: SCR_001622) and PsychToolBox (RRID: SCR_002881. Available: http://psychtoolbox.org/). The animal was placed 20 cm away from the monitor and centred on the visual midline, thus minimizing off-axis artefacts. Stimuli consisted of a single or pair (~20° separation) of 1.5° by 1.5° squares of modulated contrast ascending the receptive field at a speed of 40°/s.

We applied to our intracellular recordings a frequency-tagging technique inspired by human electroencephalography research (Norcia et al., 2015) and local field potential research in insects (van Swinderen, 2012). We presented two competing, flickering targets each with varying contrast at two different frequencies. As neuronal responses are themselves modulated by the contrast, the response frequency permits identification of the selected target. We presented non-harmonic frequency-pairs of either 8 Hz (F_1) and 12 (F_2) Hz, or 11 Hz (F_1) & 15(F_2) Hz. The two frequency pairs tested the robustness of the technique as well as ensuring that there were no artefacts induced from interactions with the display refresh rates. That is, the frequencies were not multiples of one another and were divisible by the monitor refresh rate thus ensuring the full range of intensities were presented within each period. We tested with both sinusoidal and square wave flicker. These results were pooled because there was no difference in their power to successfully identify which target the neuron was responding to.

Frequency tagged targets flickered between a minimum Weber contrast of 0.06 and maximum of 1 (mean contrast of 0.51 and a white background of 337 Cd/m^2). In single target trials, one target contrast varied at either F_1, F_2, or 0 Hz (Non-flickering control at maximum contrast) and was presented moving vertically up the display at one of two spatial locations, T_1 or T_2 (locations separated 20° horizontally within CSTMD1's receptive field). In paired target trials, two flickering targets were presented at T_1 and T_2 locations. The choice whether the spatial location T_1 or T_2 was either F_1 or F_2 (e.g. 8 Hz or 12 Hz), was pseudo-randomized to control for any preferred frequency response.

6.3.3 Experimental design and statistical analysis

For testing hypotheses about trial-by-trial selection processes, any given trial is an independent event and cannot be averaged as a technical replicate. However, to ensure robustness of the result we repeated experiments across a number of dragonflies. Here we use 'n' to denote the number of trials and additionally report across how many dragonflies. We visualise all trial data points and describe similarities or differences across animals.

We report exact P except when smaller than 0.001. All tests are nonparametric, two-tailed and corrected for multiple comparisons (Bonferroni-Holm correction). Box & Whisker plots indicate median, interquartile and minimum/maximum range. Unless otherwise stated outliers are indicated with crosses.

All data analysis was conducted in MATLAB 2017a (RRID: SCR_001622), including the Wavelet Toolbox. Complete Wavelet Transforms (CWT's) used an analytic Morse wavelet with gamma = 3.

6.4 Results

To test the validity of the frequency tagging technique, we presented a single flickering target moving vertically up the display within the dragonfly's field of view (Figure 49A). The target drifted at 40°/s within the excitatory, contralateral region of CSTMD1's receptive field (Wiederman and O'Carroll, 2013; Wiederman, Fabian et al, 2017). We use the term 'Frequency tagging' to refer to the modulation of Weber contrast: $(Intensity_{target} - Intensity_{background}) / I_{background}$, over time at a set frequency (in Hertz). Since CSTMD is responsive to dark targets (Wiederman, Shoemaker & O'Carroll, 2013), we flickered a black-to-grey target against a white background (Figure 49B). An example of an individual data trace in response to a 15 Hz target shows the spike activity during the stimulus presentation (Figure 49C, dark bar). To extract any frequency-tagged response modulation, we first determine spike times and calculate the instantaneous spike rate (Inverse Inter-Spike Interval) over time (Figure 49D). We then apply one of two mathematical transforms to this data. The application of a Fast Fourier Transform (square root to provide amplitude) reveals a peak in the frequency domain at 15 Hz (Figure 49E), equivalent to the target contrast modulation (a response at 0 Hz is due to the non-zero mean over time). We repeated this process for a series of different frequencies (averaged across neurons) to determine the most appropriate for further experiments (Figure 49F). This data shows that from 7 to 19 Hz the frequency content of the stimuli is well preserved in the intracellular response of single neurons. However, it

was previously shown that CSTMD1 can ‘switch’ selection mid-trial (Wiederman & O’Carroll, 2013). In this circumstance, power in an FFT would be distributed between the two target frequencies, corresponding to the total time each target was selected. Therefore, Fourier analysis cannot distinguish when: (1) trials where modulation was genuinely shared between T_1 and T_2 (indicative of a lack of competitive selection, such as neuronal summation) or (2) selection switched from T_1 to T_2 or T_2 to T_1 mid-way through the trial. To account for possible switches, we instead applied Continuous Wavelet Transforms (CWTs) which provide an approximate power across pseudo-frequencies *over time*. Averaging this wavelet analysis across time is similar to an FFT though reveals a broader peak in the frequency domain centred at 15 Hz (Figure 49G). The broader shape observed in the CWT is inherent to the wavelet analysis, and is the cost of providing information of how frequency components might vary over time. Although in the frequency domain CWT responses are blurred in comparison to their FFT counterparts, there are statistically significant differences for any two frequencies separated by at least 2 Hz ($P < .001$). Thus we were able to analyse all further data using CWTs to derive the benefit of examining the response evolution over time of the individual trials.

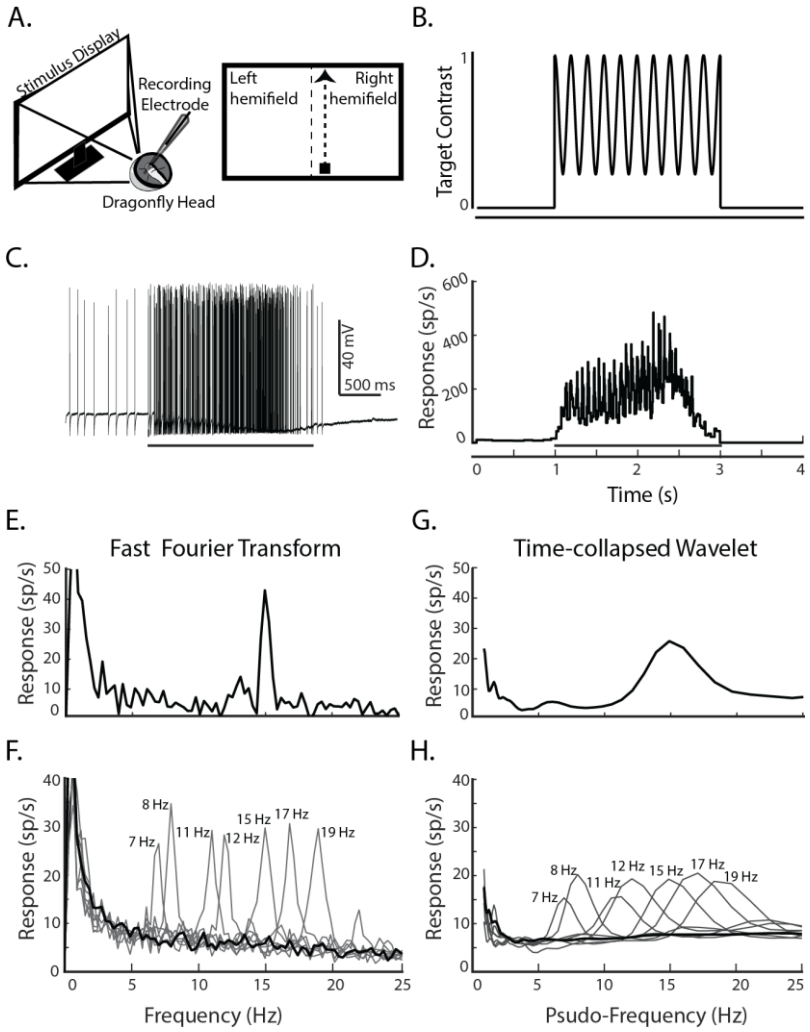

Figure 49: The frequency of the tagged target is preserved in the intracellular responses of CSTMD1. ***A)*** *Left: intracellular in vivo electrophysiology involves inserting an electrode into the intact brain to record single-cell responses to stimuli presented on a computer screen. Right: stimulus pictogram, a single small target ascends CSTMD1's excitatory receptive field.* ***B)*** *Frequency-tagging involves modulating the contrast of the stimulus over time at a specific frequency (5 Hz in this illustrative pictogram).* ***C)*** *An example spike train in response to a stimulus modulated at 15 Hz, presented at 1 s for a duration of 2 s (stimulus bar).* ***D)*** *The inverse inter-spike interval is calculated to determine the spike rate over time. This calculation provides a continuous signal that is amenable to frequency-domain analysis.* ***E)*** *A Fast Fourier Transform of the signal in D reveals a distinctive peak at 15 Hz, corresponding to the frequency-tagged stimulus.* ***F)*** *Averaged FFT of responses to trials of varying frequency (n = 119 across 4 dragonflies)* ***G)*** *The output of the wavelet analysis (collapsed across time to be comparable to the results shown in E.) provides an alternative analysis that can preserve time-domain information (in later, non-collapsed analysis).* ***H)*** *Averaged time-collapsed continuous wavelet transform for the same data presented in F, which although less peaked, still reveals statistically distinctive humps at the relevant frequencies.*

Can the frequency tagging technique replicate the selective attention result (Wiederman and O'Carroll, 2013) when two targets flickering at different frequencies are used? To test this, we presented either single targets (pseudo-randomly at either f_1 or f_2) at either spatial location

T_1 or T_2 (both within CSTMD1's excitatory receptive field). Randomly interleaved with the single target trials (Figure 50A), we also presented paired targets (simultaneously at target locations T_1 and T_2) which were frequency-modulated at the two *different* frequencies (pseudo-randomly between T_1=f_1, T_2=f_2 and T_1=f_2 and T_2=f_1). As our interest is in the chosen target (T_1 or T_2), rather than the frequency of the 'identifier', we pooled across the frequency-pairs.

In single target trials (Figure 50B, T_1 dark dots and T_2 light dots), we observed modulation at the frequency of the presented target and low modulation at the other frequency (i.e. a frequency that does not exist in the stimulus). However, in some individual trials there was insufficient modulation in the frequency domain to enable accurate identification of the selected targets. This is likely to result from two factors: (1) neuronal habituation in the receptive field diminishing the strength of the modulation (2) neuronal saturation from a highly responsive cell limiting the possible strength of the modulation. To analyse trials free of these effects, we used single-target responses to determine a threshold for data inclusion. For each location, T_1 and T_2, we calculated the average magnitude at the frequency *not presented*, which provides a value of the noise inherent in the frequency domain. This floor was defined as the mean power at the non-presented frequency plus twice the standard deviation. This provided an objective level of the modulation noise at the other frequency. That is, the expected, non-zero modulation at f_2 when the neuron has selected a target modulated at f_1, and vice-versa (Figure 50B – dashed lines). Trials in the bottom-left corner of Figure 50B thus fail the acceptable signal-to-noise threshold for both frequencies. 172 trials were rejected from further analysis for this reason and our frequency-target technique thus worked for 71.4% of the total trials presented. There was no significant difference in the amount of trials excluded via this metric between any of the three conditions (X^2-test, $P > 1$, *Bonferroni-holm correction*). Therefore, we propose that the lack of modulation was due to ether habituation or saturation, rather than related to the test under consideration – the presence of selective attention. In the successful trials, signals were above threshold at either f_1 or f_2, indicating significant response to either one, or both, of the targets. Qualitatively, we observe that the responses to paired targets (Figure 50B, crosses) were mostly either modulated at the frequency of the target at location T_1 or T_2 (but not both – crosses within the 'Shared or Switch' region).

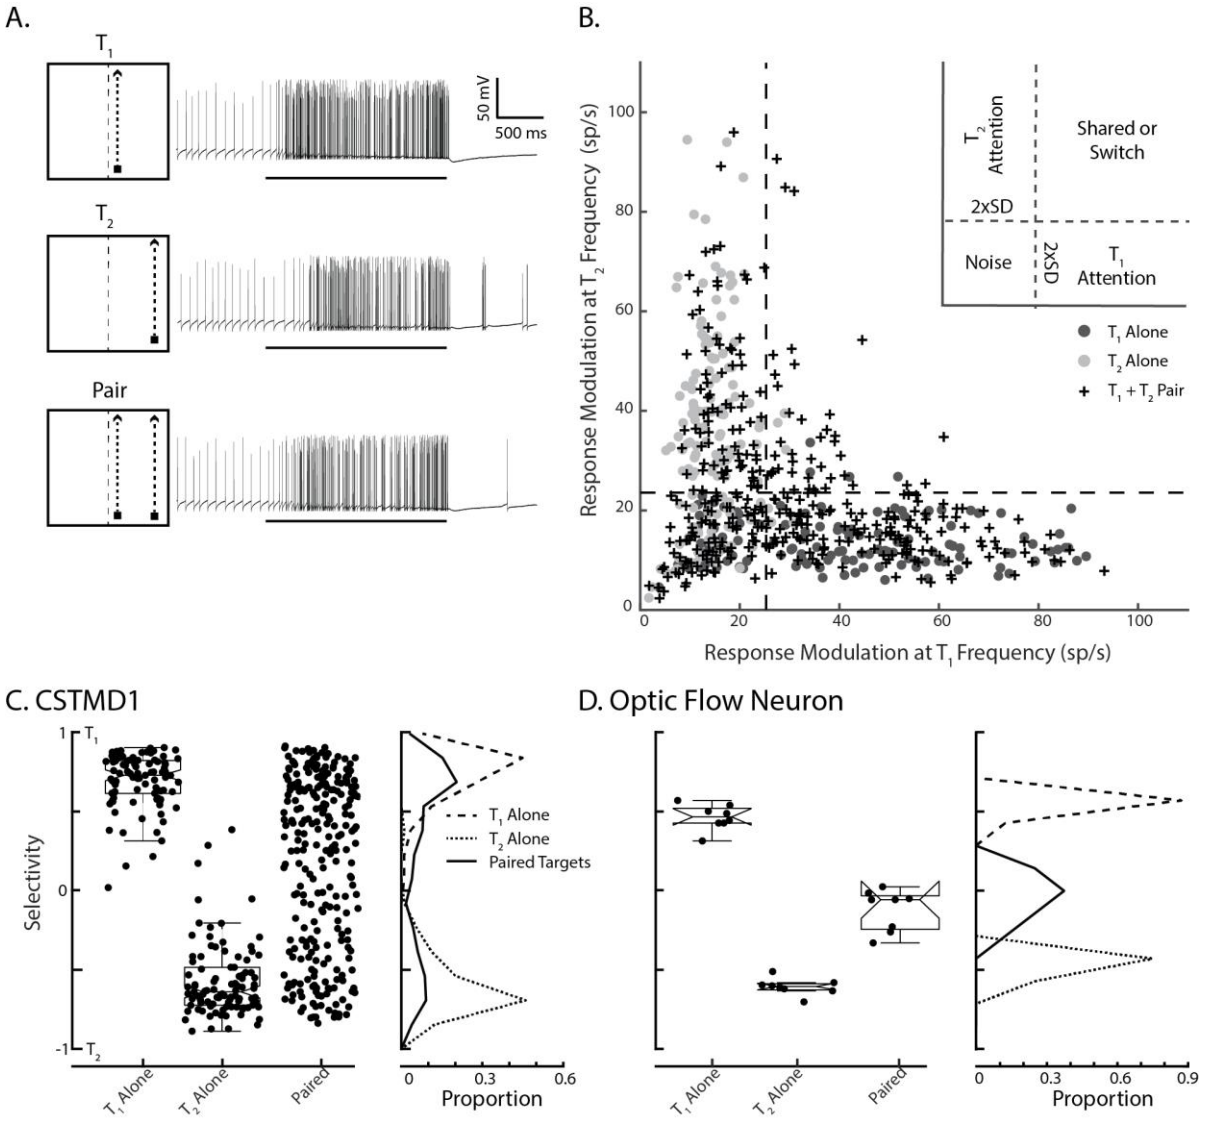

Figure 50: Frequency-tagging identifies the selected target in a paired-target trial ***A)*** *Illustrative pictograms and corresponding electrophysiological responses for the 3 stimulus conditions. From top-to-bottom: T1 Alone; T2 Alone; Paired Targets.* ***B)*** *The response modulation at the T_2 frequency plotted against response modulation at the T_1 frequency. Responses are plotted to a single target at the T_1 location (dark dots) or at the T_2 location (light dots) when presented alone. Crosses represent responses to the Paired stimulus (n = 447 trials across 13 dragonflies). Dashed lines indicate a noise threshold. Most of the responses to paired targets elicit responses at either one or other of the target flicker frequencies (not both together), indicative of a selection process* ***C)*** *The Selectivity index represents the degree to which the response favours one of the frequency tagged stimuli over the other. Values around zero indicate that both frequencies are equal components of the response. Frequency polygons illustrate the relative proportion of these points, with the bimodal distribution to the paired stimulus clearly revealing the selection of one target or the other.* ***D)*** *In contrast, results from an optic-flow neuron in the dragonfly show no selective attention (n = 8 trials in 1 dragonfly), with a unimodal distribution around zero to the paired stimulus, indicative of shared modulation to both target frequencies.*

The absolute modulation above this noise threshold (i.e. the distance of the data points along the abscissa or ordinate in Figure 50B) is related to the trial-by-trial sensitivity, rather than to the degree of the selective attention to one or either of the targets. To quantify our data, we therefore defined a *Selectivity Index* (Figure 50C), which measured the degree of target selection, independent of the strength of response modulation (though above the noise threshold previously described). For each data point, we calculated:

$$Selectivity\ Index = {T_1}\Big/{\sqrt{{T_1}^2 + {T_2}^2}} - {T_2}\Big/{\sqrt{{T_1}^2 + {T_2}^2}}$$

Equation 1

T_1 and T_2 values are averages of the pseudo-frequency amplitude (known as 'scale') over the trial duration (i.e. collapsed across time from the CWTs), for each of the corresponding target frequency-tagging modulations. The selectivity index ranges between +1 and -1 and represents the selection of T_1 (+1) and T_2 (-1), respectively. Here 'selectivity' is referred to in the original definition of 'selective attention' as selection of one from multiple competing stimuli, as would be expected in a winner-takes-all network. A value of 0 would occur if the response magnitude at f_1 and f_2 were equal (irrespective of the absolute distance from the origin), indicating either shared (co-varying) selection across the trial, or a switch in selection during the trial.

In Figure 50C, we observe significant differences in the Selectivity Index distribution between paired and both T_1-alone and T_2-alone conditions ($P < 0.001$). In single-target conditions, the Selectivity Index is narrowly distributed (T_1 $\mu = 0.68$, $\sigma = 0.17$; T_2 $\mu = -0.58$, $\sigma = 0.23$), whereas in paired-target trials the Selectivity Index is non-normal ($p > 0.001$, one-tailed Kolmogrov-Smirnov test) with peaks at approximately 0.65 and -0.55. The bimodal distribution of responses to paired targets reveals the selection of either T_1 or T_2. For comparison to a potential 'null' hypothesis (i.e. no selective attention), Figure 50D shows results from a single 'optic flow' neuron in the dragonfly. This neuron generates robust responses using spatial summation in order to encode wide-field optic flow, analogous to Lobula Plate Tangential Cells in Diptera (Hausen, 1982). We presented the same experimental paradigm, though with larger targets (1.5° x 10°) to elicit a response. In contrast to the results observed in CSTMD1, the optic-flow neuron had a Selectivity Index around 0 (modulation at both frequencies of the paired targets) indicative of neuronal spatial summation.

Not all of the paired-target trials were solely modulated by one of the target frequencies (Figure 50B, shared zone). If CSTMD1 is only selectively attending one of the presented targets, what could account for this apparent shared modulation? There are two possible explanations: firstly, the neuron is excited by both stimuli at their respective frequencies and not selecting a single target. That is, spatial summation similar to what is observed in the Optic Flow neuron (Figure 50D) and in primate V4 (Ghose & Maunsell, 2008). Alternately, a switch mid-way through the trial could result in significant modulation at both frequencies, as both targets are selected during the trial, though at discrete times.

To test this possibility, we first simulated a switch in response from f_1 to f_2 by presenting a single-target that changed frequency in the middle of the trial (Figure 51A). An example of the intracellular response to such a pseudo-switch stimulus is presented in Figure 51B. We subtracted the two wavelet magnitude 'slices' from one another (Figure 51C, dashed lines) derived from the CWT analysis, thus producing a difference in magnitude between the two pseudo-frequencies over time (Figure 51D). This difference in magnitude highlights the degree to which one frequency is selected over the other throughout the time course. For example, capturing the frequency change in the simulated switch with a peak and trough dissociated in time. This technique therefore provides a 'read-out' through time of when the response was governed by which frequency. In a case where power was shared, we would expect a flat line representing little variation in the frequency magnitudes across time.

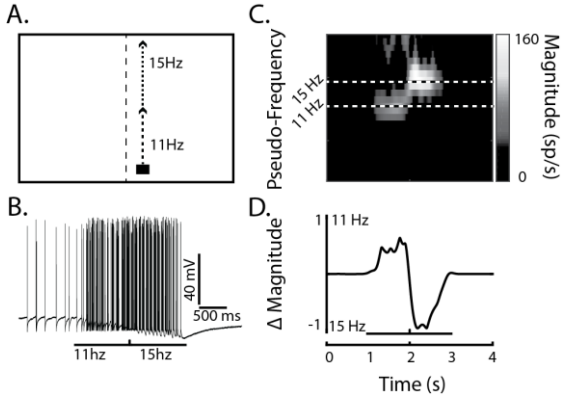

Figure 51: Simulation of an 'attentional switch' produces a characteristic result ***A)*** *Illustrated pictogram of a single target that changed frequency halfway through the trial, simulating an attentional switch from a target of one frequency to the other.* ***B)*** *An example of a CSTMD1 response to this switching stimulus.* ***C)*** *The CWT analysis of the inverse ISI of the trial in B, reveals the switch that occurs halfway through the trial. The black-and-white dashed lines indicate the 11 Hz and 15 Hz frequency slices.* ***D)*** *A 'difference slice' (delta magnitude) is calculated by taking the difference between the wavelet slices at 11 and 15Hz across time.*

We applied this 'read-out' analysis to determine whether the paired target responses with modulation at both frequencies (Figure 50B, shared or switch region) were due to spatial summation or switching. Figure 52A shows examples from six such trials, all of which exhibit discrete peaks and troughs in time. The traces indicate that these CSTMD1 responses are switching between targets over time, rather than being modulated by both target frequencies simultaneously.

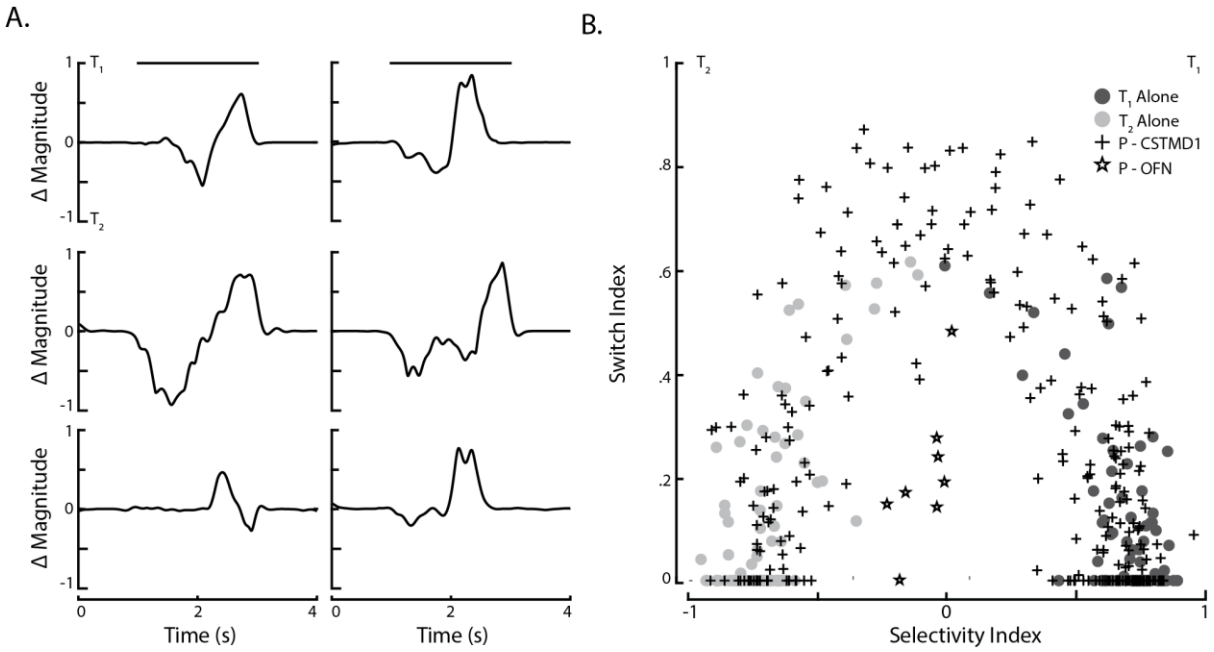

Figure 52: Shared modulation results from switches in selection ***A)*** *Individual trial examples of the 'difference slice' from the wavelet analysis of paired-target trials showing high modulation for both targets at different epochs of time.* ***B)*** *The 'Switch index' and 'Selectivity index' for all single target (dark and light points) and paired target (crosses) trials. When selectivity for paired targets is low (middle abscissa, close to zero) then the Switch Index is high, indicating that responses switched between targets. Instead, the optic flow neuron (stars) has low Selectivity and a low Switch, indicative of response summation (modulation at both frequencies across points in time).*

To compare aggregate data, we calculated a 'Switch Index' for each trial (Figure 52B). This index was calculated by first determining the proportion of time the system selected either T_1 or T_2. To ensure that these selections were robust we only considered a selection valid when either target was significantly stronger (> 5 spikes/s) than its counterpart (i.e. when $T_1 > T_2 + 5$ or vice-versa). Having established how long each target was selected, we then multiplied these two values together. This has the effect that when one of the two targets was not selected at all, the Switch Index is zero, while it is maximized when both targets are selected for 50% of the trial. This value was normalized between 0 and 1. In single-target trials (dark or light dots), the Switch Index is overall low, however in paired-target trials the Switch

Index is distributed between high and low values. In trials with a Selectivity Index around 0, the Switch Index is uniformly high, indicating that the low selectivity is almost entirely due to switches. In contrast, paired-target trials in the dragonfly optic flow neuron show both a low Selectivity Index and low Switch Index, indicating genuine modulation at both frequencies over time due to the spatial summation used in optic flow computations (Figure 52B, stars).

6.4.1 Biasing selection with priming

Next, we tested the ability of a priming stimulus to *bias* the selection of a spatially-associated target in a paired-target condition. In this experiment, a lone untagged primer was first presented for one second moving towards the trajectory of either spatial location T_1 or T_2 (Figure 53A). Note that here the frequency-tagged T_1 and T_2 pathways commence midway up the stimulus display, immediately after the single 'primer' target has moved along its trajectory. From our previous work, we expect CSTMD1 to 'lock-on' and predictively facilitate responses in front of the target's prior path (Nordstrom 2011; Dunbier et al, 2012; Wiederman & Fabian et al., 2017). In this experiment, we introduced a frequency-tagged distracter midway through the receptive field (horizontally offset by 20°) paired with a frequency-tagged target that continued along the previous 'primer' targets' trajectory (Figure 53A). We calculated the Selectivity Index across the entire second where both targets are presented together and reveal a statistically significant ($P < 0.001$) bias for selection towards the target that continues along the primed trajectory (Figure 53B). This selection may be attributed to the previously described predictive gain modulation, whereby a 'local spotlight' of enhanced gain is generated ahead of a moving target, with suppression in the surround (Wiederman & Fabian et al., 2017). In our experiment, the continuing target is within the spotlight created by the preceding target, but the second target appears within the supressed surround.

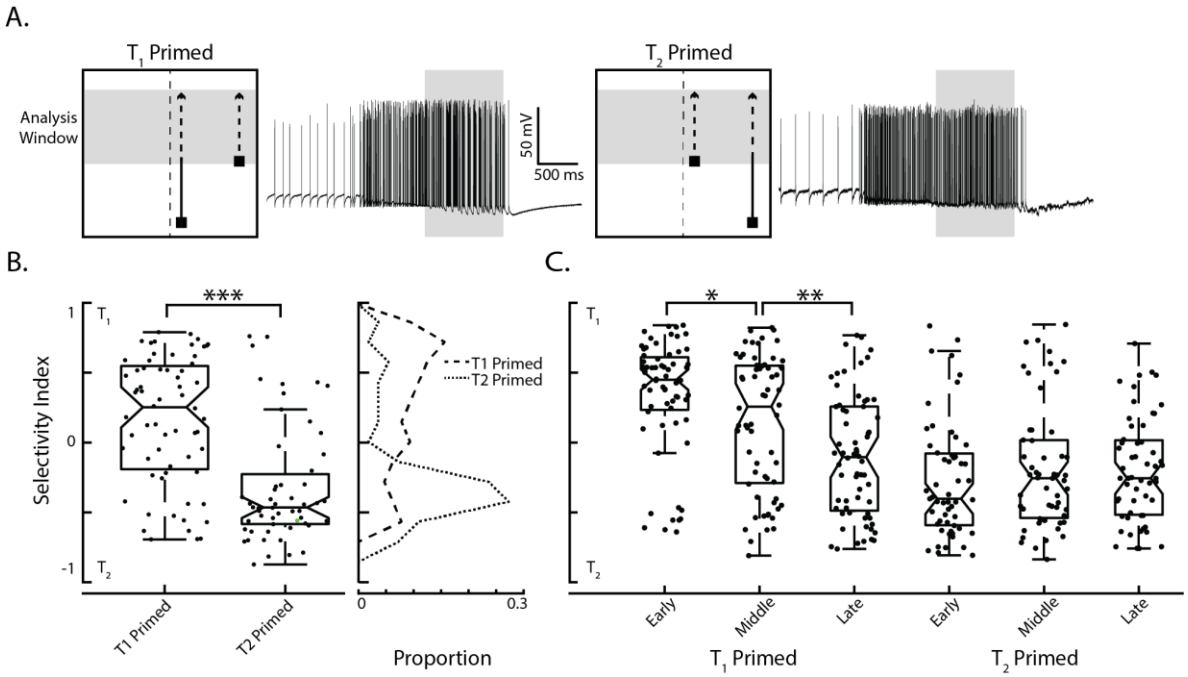

Figure 53: Priming with a preceding target biases selection towards the continuing trajectory ***A)*** *Pictograms illustrate the biasing stimulus towards either spatial location T_1 or T_2, next to individual example of CSTMD1 responses to the stimuli. The short-path target (distracter) appears at 1 second, when the preceding target reaches midway up the screen (the analysis window indicated with the grey shade).* ***B)*** *There is a significant difference between the Selectivity Index when T_1 was primed with the preceding target compared to when T_2 was primed (n = 295 across 7 dragonflies). Frequency polygons reveal the distributions of the Selectivity Index for T_1 primed (dashed line) and T_2 primed (dotted line).* ***C)*** *In order to assess the impact of attentional capture we split the paired target period into 3 windows which were analysed separately.*

In the human psychophysics literature, attentional capture is an effect whereby the presentation of an abrupt-onset stimulus (Yantis & Jonides, 1984) or a novel object (Franconeri, Hollingworth & Simons, 2005) involuntarily captures attention (Remington, Johnston & Yantis, 1992), even when task-irrelevant. In order to test for a capture of CSTMD1's selection, we analysed the previous biased paired-target responses (Figure 53B) separated into three 400 ms periods (Early, Middle and Late). We included 100 ms overlap between these periods because this duration was required for meaningful CWT analysis. If CSTMD1 responses displayed attentional capture, we hypothesise that the early period would be dominated by responses to the distracter stimulus, returning to the original path at later periods of time (as the distracter is assessed and ignored). Our results reveal the opposite effect (Figure 53C), with the early window exhibiting the strongest effect of the biasing which dissipates over time. This reveals that the selection is not automatically captured by the abrupt-onset novel stimuli presented within CSTMD1's receptive field, rather responses are locked on to the preceding target. Here we observed asymmetry in results from the T_1

compared to T_2 priming, which reflects the broader (noisier) distribution of values in the T_1 Primed condition when analysed over the entire analysis duration (Figure 53B). When primed to T_1 (the target closer to the dragonflies' midline), the Early window (Figure 53C) reflects this biasing to the continued path trajectory (though note the several clear exceptions). However, in some cases over time (Middle and Late windows) selection changes towards the distracter location at T_2, resulting in significant changes in the Selectivity Index between these periods ($P < 0.001$). Note that visual inspection of the CWT analysis reveals that these are switches that occur at specific points in time in the individual trials. In the T_2 priming condition (the target located in the more peripheral location), the selection has locked on to the preceding target and maintains this selection throughout the rest of the trial, with no significant difference between the Early, Middle and Late periods (again with several clear exceptions).

In a traditional winner-takes-all network, the introduction of a high contrast distracter during the presentation of a lower contract target would result in a switch to the one with higher salience. However, how would the dragonfly feed in a swarm if often distracted by a novel, transiently more salient target? To determine whether CSTMD1 locks-on to the lower-salience stimuli, we presented primers of varying contrasts followed by a paired frequency tagged distracter. In this experiment, we wanted the lower contrast target to retain its lower saliency throughout the course of the trial, even during the period when the frequency-tagged distracter was present. Because the original target (of varying contrast) is never frequency-tagged, it is the measure of modulation that is the indicator of distracter selection. Primers were presented at constant Low (0.06), Medium (0.15) or High (0.51) contrast, pseudo-randomly located at spatial locations T_1 or T_2 (Figure 54A, primer at T_2 location shown). The High contrast primer was set at 0.51 to be equiluminant with the *average* contrast (over time) of the frequency-tagged, high contrast distracter. Figure 54A shows example responses of an individual CSTMD1 to these stimulus conditions, both to when the primer retains selection and when the distracter takes over. This shows that there can be trial-by-trial variability in which of the targets was selected (primer or distracter).

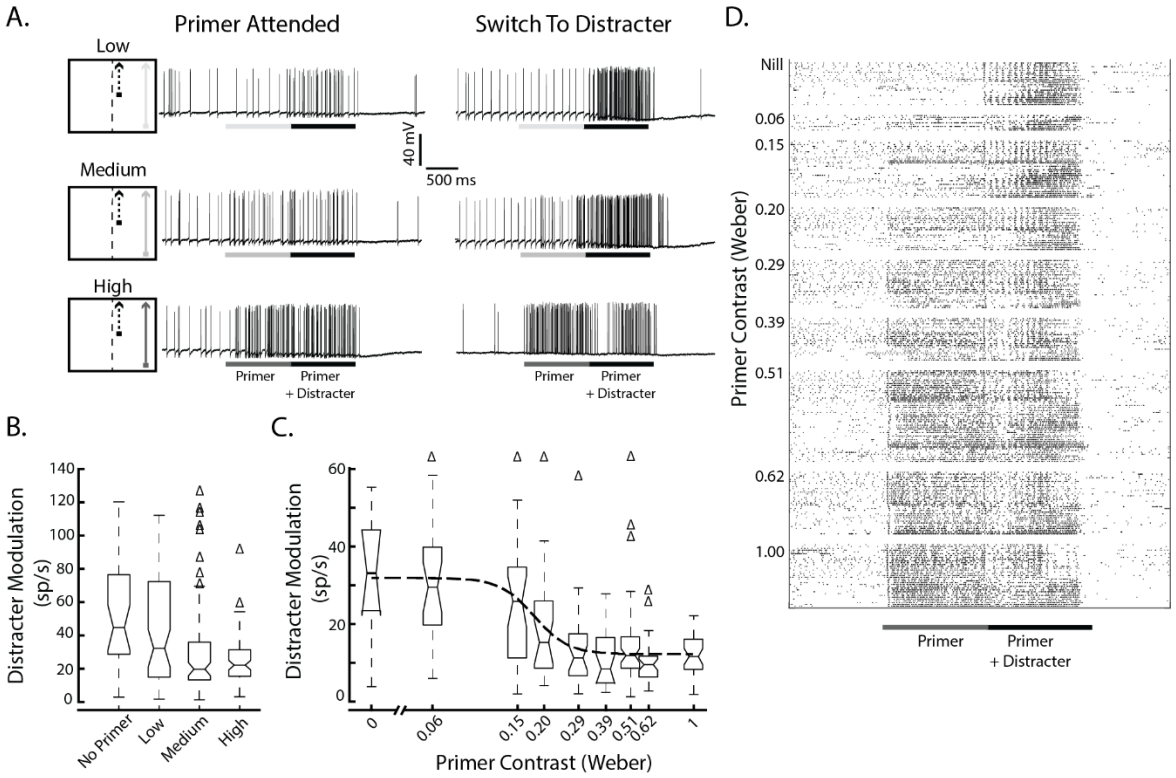

Figure 54: Selective attention in CSTMD1 can 'lock on' to a lower contrast target, ignoring the introduction of a high contrast distracter (frequency-tagged) ***A)*** *Stimulus pictograms and example traces from the same CSTMD1 for Low, Medium, and High contrast primer conditions.* ***B)*** *Boxplots showing the modulation at the distracter frequency across four primer conditions (n = 204 across 5 dragonflies). Δ indicate outliers. The distracter only condition (No Primer) shows the expected distribution of modulation values if the priming target is never selected. Even in the Low contrast condition, there is a significant difference, indicative that in some trials the Low contrast target is selected during the period when a high contrast distracter is present. As the salience of the primer is increased, the number of trails where the distracter is selected decreases* ***C)*** *In an individual CSTMD1 recording, we assayed across a larger range of primer contrasts, revealing a sigmoidal contrast sensitivity function (n = 212 across 1 dragonfly). Δ indicate outliers. Off-axis outliers are indicated.* ***D)*** *Spike rasters organized by primer contrast from the individual CSTMD1 data presented in C. with trials within each contrast are ordered by presentation time.*

Figure 54B shows a significant reduction in distractor modulation as primer contrast rises (Kruskal-Wallis One-way ANOVA, df = 3, X^2 = 21.32, p < 0.001), indicating that CSTMD1 can lock-on to Low contrast targets even in the presence of a high contrast distracter. There is still trial-by-trial variability, however as primer contrast increases a higher proportion of trials do not exhibit distracter modulation, thus have selected the primer. Due to the previously described biasing effect of a preceding primer (Figure 53), we would expect less distracter modulation at the equiluminant contrast (higher than 0.51). We observed a significant reduction in distracter modulation in the medium (P = .006) and high (P > .001) contrast

group, but not the low-contrast group (P = .755), compared to the no primer group (Figure 54B), suggesting that CSTMD1 was indeed able to lock on to primed targets of .15 and higher contrast. However, outliers observed in both the Medium and High contrast conditions indicate that CSTMD1 is still able to switch selection to the distracter on some trials, consistent with the previously observed "rare" switching which occurred with two equally salient targets (Wiederman & O'Carroll, 2013).

From a single CSTMD1 recording, we were able to assay across more primer contrasts (Figure 54C). In this neuron, the primer often locked-on to the much lower primer contrasts (as low as 0.2), ignoring the simultaneously presented distracter target. Therefore the mechanism underling this neuronal selective attention cannot be a 'simple' winner-takes-all network. Raster plots of spikes throughout the trial reveal some interesting attributes of the response (Figure 54D). Even in trials where the distracter was not selected, the onset was marked with a reliable spike, perhaps indicative of a transient breakthrough in the underlying network. Additionally, in conditions when the primer contrast is higher, there is often a transient suppression of response (timed with the distracter onset) before responses continue with their selection. Further experiments are required to elucidate how this phenomenon relates to the suppressive surround observed in the predictive gain modulation during single target trajectories (Wiederman, Fabian et al. 2017)

Overall, these data clearly reveal that CSTMD1 is able to lock on to low-contrast targets and select them even in the presence of a high contrast distracter. Intriguingly, response to these continued primer trajectories are not associated with an increase in spike rate as would be expected by models of attention where low-contrast stimuli are attended by neuronally boosting contrast (Reynolds & Desimone, 2003). Instead, even when responding to low-contrast stimuli, CSTMD1 encodes the absolute strength of the attended target as if the distracter was not present (Figure 54A). This could be critically important in behaviour where a target is selected for pursuit amidst a swarm, where absolute rather than relative activity might ber required to drive the closed-loop control system controlling the animal's flight trajectory.

6.4.2 Modelling

What mechanism best explains the measured data? To test this, we developed six algorithmic models. The six models included two models that assumed shared attention (including one with saturation), two models that applied selection and two models which applied selection with switching. For input to these models we collected the single target trial (i.e. T_1-only or

T_2-only) response modulation amplitude from the wavelet analysis (Figure 50). From this we produced four lists (T_1f_1, T_1f_2, T_2f_1, T_2f_2) representing the response modulation amplitude at the target's flicker frequency and at the comparison frequency (i.e. no modulation). We binned these responses and fit a log-normal distribution to each target and frequency pair (T2 examples are shown in Figure 55A). We were then able to infinitely sample from these model distributions to generate an arbitrary number of synthetic target responses.

Simulating switches requires a time-course of the response modulation amplitudes over time. To simulate this, we generated a 1s time course for testing all models (even non-switching models). These time-courses represent the instantaneous response modulation amplitude over time equivalent to taking a 1-dimensional slice from the CWT analysis (as in Figure 51). It is unrealistic that the response modulation amplitude at a given frequency would be constant over a 1s period. To ensure a more realistic response, we added noise (white noise with a 5mV max width). We then smoothed the data using a 0.2s average filter which produced waveforms qualitatively similar to those observed from taking a single-frequency slice of a CWT. For switching models the smoothing was done after calculating the switch. For example, if a switch from T_1 to T_2 occurred at 0.4s, the first 0.4s would use both the T_1f_1 and T_1f_2 (each with noise added) while the subsequent 0.6s would use T_2f_1 and T_2f_2 (again with noise added). This process regularly produced 'step-like' responses to which we applied smoothing (as mentioned above) to generate the smooth transitions we saw in the CWT data.

We sampled from the distribution 1000 times for each pairing (T_1f_1, T_1f_2, T_2f_1, T_2f_2). Each model used some combination of these to generate output responses for both f_1 and f_2. Basic Summation (BS) assumed that the output power at both f_1 and f_2 were the corresponding powers of the input target (i.e. T_1f_1 & T_2f_2). Saturating Summation (SS) summed like BS, but applied a soft saturation to reduce the overall modulation power evenly between f_1 and f_2 to a maximum potential power of 100 spikes/s. Random Selection (RSe) randomly selected either T_1 or T_2 and used that target's corresponding power for f_1 and f_2 (i.e. if T_1 was selected the frequency responses would be T_1f_1 and T_1f_2). Winner Selection (WSe) selected the target with the greatest modulated power (using the tacit assumption that the modulation was proportional to the target response) and used the winner's frequency response solely. Thus if $T_1f_1 > T_2f_2$, T_1 would be selected and vice versa. Random Switching (RSw), randomly selected an initial target (as per RSe) but assumed that a switch occurred in a percentage of trials at some point during the trial's duration (specific values determined after optimization, see below). Multiple Switching (MSw) assumed a more sophisticated switching rate,

allowing the system to switch multiple times. The switch probability was defined by the following formula:

$$P(switch) = S - \tau e^{-t/\tau}$$

Equation 2

S represents the probability that a switch never occurs and τ represents the rate of increase of switching over time (Figure 55B). The values of S and τ chosen were determined after an optimization step (see below).

The generated outputs of all six models are shown in Figure 55C. The summation model (BS) populates all four quadrants (including in the 'Shared or Switch' zone of Figure 50B). This combination of taking power from both targets together does not match the electrophysiological results (Figure 50B). Both selection models (RSe & WSe) adhere far closer to the distribution seen in Figure 50 except that the shared zone is very sparsely populated, especially in the WSe model. The switching models are good qualitative matches for the real data with a bias to T_1/T_2 only responses (the L shape) but with a reasonable number of shared zone responses indicative of switching.

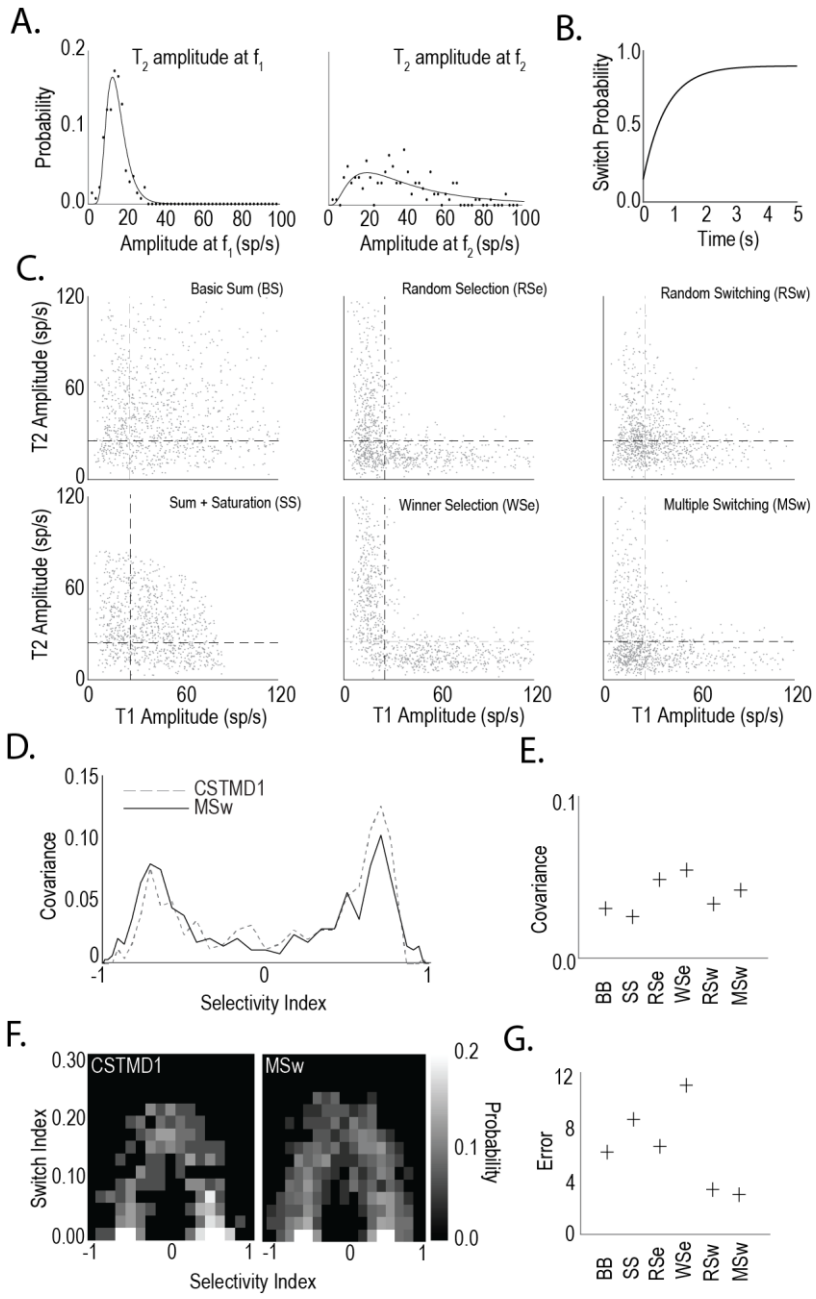

Figure 55: Simple Switching Models Matches Selective Attention Data ***A)*** *Power distributions for frequency responses from T_2 at f_1 (left) and f_2 (right) calculated from recorded trials. Modelled trial data were randomly selected from these power distributions representing the power contribution of each target.* ***B)*** *Switch probability as time progresses for model 6 (Multiple Switching). Initially the likelihood of switching is low before rising to 90%. After a switch, the switch probability resets allowing multiple switches to occur.* ***C)*** *Example scatter plots (as per Figure 50B) for each of the six models tested. Summation (top left), summation with saturation (bottom left), random selection (top middle), higher power always wins selection (bottom middle), random switching (top right), and multiple switch model (bottom right).* ***D)*** *Histogram of model selectivity for recorded data and model output. Error calculation as covariance curve.* ***E)*** *Results of six models against recorded data from histogram analysis shown in (D). Higher covariance is indicative of a more representative model.* ***F)*** *Two-dimensional histogram (selectivity/switch index) for recorded data (left) and model data (right). Error calculated as RMS deviations from recorded data.* ***G)*** *Results of six*

models against recorded data using 2D histogram data (F). Low values indicate representative model.

To assess each model quantitatively, we generated the frequency polygon (Figure 50, Figure 52) of the selectivity index values calculated from the model outputs. An example of the response of the MSw model (grey line) compared to the electrophysiological data (dotted line) is shown in Figure 7D. We compared each model's frequency polygon with frequency polygon from Figure 50C via cross-correlation.

Via this metric, both selection models (RSe, MSe) provided the best match to the recorded data (Figure 55E). However, the selection metric ignores the switching behaviour inherent in the model. To test whether pure selection was sufficient to explain the data, we used the model outputs to calculate the 'Switch Index' (Figure 52) for each model's responses. We then binned this data to generate a 2-dimensional histogram (Figure 55F). We repeated this process for the electrophysiological data and calculated the RMS error between the two. As both switching models had free parameters (i.e. probability of switching) we optimized both these models against this RMS error. The RSw model was most successful with a 100% probability of a switch at a random time during the trial. The MSw model was optimal with a 90% switch probability and 0.75s time constant. The remaining models (Summation and Selection) did not have any parameters to effectively optimize.

Figure 55G shows the results of the Switch Index comparison for the six models. Both switching models have lower RMS than the other models with the multi-switch model performing the best overall. It is clear from both the qualitative and quantitative aspects that the switching models produce the best outcomes. This is in line with our expectations. As mentioned previously, the Summation models (BS, SS) generate too many responses in the shared zone by effectively increasing the overall power, while the Selection models (RSe, WSe) go the opposite direction effectively eliminating most of the responses from the shared zone. The Switching models (RSw, MSw) provide a suitable compromise, with a general shift towards the upper-left and lower-right quadrants while maintaining some responses in the shared zone, However, while there were numerous responses in this region they have lower overall power indicative of temporal sharing (rather than summation). It is clear that the best explanation for the results seen is a model that selects a single target but is capable of switching one or more times during a trial.

6.5 Discussion

Frequency-tagging techniques have previously been used during higher-order brain measurements (e.g. EEG) or in extracellular recordings measuring local field potentials (LFP) in insects (van Swinderen, 2012). However, it is not yet known whether frequency components within the frequency-tagged LFPs originate at the level of single neurons, or are an emergent property of a neuronal population code. To our knowledge, here we present the first application of this identification technique at the intracellular level. We thus demonstrate that the frequency component of the stimulus is preserved in the response of an individual neuron.

Frequency tagging allows us to verify previous findings of selective attention in CSTMD1 (Wiederman and O'Carroll, 2013) and for the first time robustly identify which of a pair of targets was selected at any moment in time. However, it is clear that frequency tagging is not always robust. In approximately 25% of trials, regardless of stimulus conditions, levels of frequency modulation were below-noise despite a reasonable spiking response (Figure 50B; bottom-left corner). These trials were excluded as identification of the selected target could not be achieved. Difficulty in choosing the correct stimulus waveform may underlie this problem: Firstly, flickering targets located within the strongest parts of the receptive field may saturate, resulting in a lack of headroom for significant modulation. Conversely, frequency-tagged targets presented in less sensitive regions of the receptive field may not elicit responses strong enough to carry modulation over the underlying signal. Both factors, saturation and sensitivity, can vary dynamically as overall CSTMD1 responsiveness may change over time, location or between animals. These effects could be minimized by changes to the stimulus waveform, with a lower mean level of intensity accounting for saturation and a higher amplitude of contrast modulation for sensitivity. However, as these exclusions did not affect our hypothesis testing (were distributed equally across all experimental conditions), we kept the amplitude and mean level consistent across all experiments.

Although frequency tagging was used as an identifier, could the frequency itself interact with facilitatory or selective processing? Such a factor can play a role in other animal models, with honeybees preferencing 20-25 Hz and avoiding 2-4 Hz visual flicker (Van De Poll et al., 2015). Even a single luminance change is enough to break inattentional blindness in humans (Palmer et al, 2018). To minimise this possibility, we distributed the two tagging frequencies across the two spatial locations (T_1 and T_2) as well as testing our entire experimental paradigm at two different frequency-tagged pairs. Throughout these experiments, we did not

observe any effect of the frequency-tagging beyond our intended purpose as an identification technique.

Attention is a limited resource (Alvarez & Franconeri, 2007), therefore animals across species are motivated to guide the deployment of attention in an ethologically meaningful and efficient way. One guide is spatial or temporal cueing, often through inhibitory neural mechanisms (Romer et al., 2002; Ruthruff & Gaspelin, 2018). For example, Drosophila are more likely to orient towards cued locations of the receptive field when subsequently presented with multiple targets (Sareen et al., 2011). Female crickets prefer leading male auditory signals to signals arriving later (Snedden & Greenfield, 1998; Romer et al., 2002), suggesting an inherent bias towards 'locking on' to the first stimulus and ignoring those subsequent. This is similar to what we have observed in CSTMD1, with the priming by a preceding target biasing selection to those that continue along the projected trajectory.

In CSTMD1, the effect of spatiotemporal cueing was so strong that even targets of lower visual salience can win over the simultaneously presented distracter. In attentional networks, saliency is a prominent attribute for guiding selection and seems to innately capture attention. This leads to a conundrum; if the most salient objects were to capture attention moment-to-moment, then the system might too often be distracted from any given task. For example, will the dragonfly ever feed if the prey of varying contrast (i.e. moving against a cluttered background) becomes less salient than others in the swarm? Conversely, the onset of a novel salient stimulus may signal the necessity to attend to a new event or abandon the current task completely in favour of survival behaviour (e.g. an approaching bird).

In human psychophysics, both abrupt-onset (Yantis & Jonides, 1984) and perceptually new objects (Franconeri, Hollingworth & Simons, 2005) provoke attentional capture, a phenomena where attention is automatically and involuntarily directed at a particular, often task irrelevant, feature (Remington, Johnston & Yantis, 1992). The signal suppression hypothesis by Sawaki and Luck (2010) proposes that all stimuli automatically generate a saliency signal, but this signal can be supressed by top-down attentional mechanisms. In our CSTMD1 recordings, we found no evidence for attentional capture. Instead, the earliest period of the paired targets revealed the strongest bias to the previous primer trajectory, with the possibility of switching to the novel distracter at a later time. Thus rather than attending to a novel distracter, this system is locking on to the expected target trajectory. These results may be attributed to the previously observed effect that CSTMD1 predicts future target location following an occlusion (Wiederman, Fabian et al, 2017) with an enhancement for the

prior path and suppression in the surround. During the initial window, the continuing target is fully facilitated by the preceding target and continuously moving into its self-generated 'spotlight' of gain enhancement. However, the distracter appears within the supressed surround and therefore will not elicit attentional capture (Ruthruff & Gaspelin, 2018) in agreement with the signal suppression hypothesis (Sawaki & Luck, 2010). Shortly after its appearance, the distracter may have self-facilitated, enabling a more even competition for target selection and thus increasing the probability of a switch. Whether this self-facilitation occurs at both target locations before selection, or only at the single selected location is currently under investigation.

These results bear resemblance to behavioural results in *Drosophila* (Koenig, Wolf & Deisenberg, 2016). Tethered flies in an arena were presented with a pair of vertical lines equally offset from the flies' midline. Flies made a decision to respond to either one line or the other by turning to bring it into the midline. In subsequent trials, these flies displayed a bias for turning towards the originally selected stimulus and ignoring the alternative. However, over time this bias was lost. The mean 'attention span' (time before the bias was lost) was 4 seconds in wild-type flies, but reduced to 1 second in mutants defective in selective attention. Active switching between competing stimuli may be indicative of endogenous drive by top-down control mechanisms (Miller, Ngo & van Swinderen, 2012). Van Swinderen (2007) found that, in Drosophila, a minimum amount of time must pass between the original selection of a target and switching to a new stimulus, and switching at all was reliant on short-term memory genes.

The possibility that non-selected stimuli also generate a 'spotlight' of neuronal gain modulation is in agreement with proposed mechanisms underlying attention in primates (Reynolds & Desimone, 2003). Primate cortical cells are thought to be 'hard-wired' to respond to the highest contrast stimulus, a property that can be exploited by attentional systems in V4 (Schiller & Lee, 1991**;** DeWeerd et al., 1999).Here the representation of stimuli is modulated by enhancing the effective contrast of the focus of attention (Martinez-Trujillo & Treue, 2002; Reynolds & Desimone, 2003). Through this enhancement, less salient and even non-preferred stimuli can come to dominate the response of neurons in V4 (Reynolds & Desimone, 2003), MT, and MST (Recanzone, Wurtz & Schwarz, 1997; Treue & Maunsell, 1999).

This neuronal enhancement observed in primates may be mechanistically similar to the facilitation observed in CSTMD1 when in response to a single target gain is increased ahead

of the prior path and suppressed in the surround. However, in primates it is the presence of distracters that establishes this attentional enhancement (Treue & Maunsell, 1999; Reynolds, Pasternak & Desimone, 2000; Treue, 2001; Reynolds & Dismone, 2003). In CSTMD1, the single target is enhanced and in the presence of distracters the absolute strength of the selected target is retained as if the distracter did not even exist.

The ability of a neuron to respond with the same strength to a target presented alone, or when selected from a pair, may underlie the dragonflies' exceptional ability to hunt in swarms (Combes et al, 2012). Such neuronal processing may have evolved to overcome the confusion effect by singling-out targeted prey amidst a swarm (Landeau & Terborgh, 1986). Behavioural studies in some dragonfly species, *Libellula* adults (Combes et al, 2012) and nymphs (Jeschke & Tollrian, 2007) show that they are adept at hunting in swarms throughout life. Although not tested in *Hemicordulia,* this hawking dragonfly is also likely to benefit from neuronal processing that reduces the confusion effect via selective attention, as they spend most of their adult life hunting and patrolling territory on the wing in environments which include swarms of prey and conspecifics.

6.6 References

1) Alvarez GA, Franconeri SL (2007) How many objects can you track?: Evidence for a resource-limited attentive tracking mechanism. Journal of Vision 7(13): 14
2) Asadollahi A, Mysore SP, Knudsen EI (2010) *Stimulus-driven competition in a cholinergic midbrain nucleus*. Nature neuroscience 13(7): 889
3) Baird JM & May ML, 1977, *Foraging behaviour of Pachydiplax longipennis (Odonata: Libellulidae)*. Journal of Insect Behavour 10(5): 655-678
4) Combes SA, Rundle DE, Iwasaki JM & Crall JD, 2012, *Linking biomechanics and ecology through predator-prey interactions: flight performance of dragonflies and their prey*. J. Experimental Biology 215:903-913
5) De Bivort BL & van Swinderen B, 2016, *Evidence for selective attention in the insect brain*. Current Opinion in Insect Science 15: 9-15; doi: 10.1016/j.cois.2016.02.007
6) DeWeerd P, Peralta III MR, Desimone R & Ungerleider LG, 1999, *Loss of attentional stimulus selection after extrastriate cortical lesions in macaques*. Nature Neuroscience, 2(8), 753

7) Dunbier JR, Wiederman SD & Shoemaker P, 2012, *Facilitation of dragonfly target-detecting neurons by slow moving features on continuous paths*. Frontiers in Neural Circuits 6:

8) Edman JD & Haeger JS, 1974, *Dragonflies Attracted to and Selectively Feeding on Concentrations of Mosquitoes*. The Florida Entomologist 57(4): 408

9) Franconeri SL, Hollingworth A & Simons DJ, 2005, *Do New Objects Capture Attention?* Psychological Science 16(4): 275-281

10) Geurten BH, Nordstrom K, Sprayberry JDH, Bolzon DM & O'Carroll DC, 2007, *Neural Mechanisms underlying target detection in a dragonfly centrifugal neuron*. J. Experimental Biology 210(18): 3277-3284

11) Ghose GM & Maunsell JHR, 2008, *Spatial summation can explain the attentional modulation of neuronal responses to multiple stimuli in area V4*. Journal of neuroscience 28(19): 5115-5126

12) Harrison IT, Weiner KF & Ghose GM, 2013, *Inattention Blindness to motion in middle temporal area*. Journal of neuroscience 33(19): 8396-8410

13) Hausen K, 1982, *Motion sensitive interneurons in the optomotor system of the fly. II. The Horizontal Cells: Receptive field organization and response characteristics*. Biological Cybernetics 46(67-79)

14) Jeschke JM & Tollrian R, 2007, *Prey swarming: which predators become confused and why?* Animal Behaviour 74(3): 387-393

15) Labhart T & Nilsson DE, 1995, *the dorsal eye of the dragonfly Sympetrum: specializations for prey detection against the blue sky*. Journal of comparative physiology A 176(4): 437-453

16) Landeau L & Terborgh J, 1986, *Oddity and the 'confusion effect' in predation*. Animal Behavour 34(5): 1372-1380

17) MATLAB, RRID: SCR_001622

18) Martinez-Trujillo JC &Treue S, 2002, *Attentional modulation strength in cortical area MT depends on stimulus contrast*. Neuron 35(2): 365-370

19) Nityananda V, (2016, *Attention-like processes in insects*. Proc. R. Soc. B 283(1842); 20161986

20) Norcia AM, Appelbaum LG, Ales JM, Cottereau BR & Rossion B, 2015, *The steady-state visual evoked potential in vision research: A review*. Journal of Vision 15(6); doi: 10.1167/15.6.4

21) Nordstrom K, Bolzon D & O'Carroll, 2011, *Spatial facilitation by a high-performance dragonfly target detecting neuron.* Biology Letters 7(4): 588-592

22) Olberg RM, Seaman EC, Coats MI & Henry AF, 2007, *Eye movements and target fixation during dragonfly prey-interception flights.* Journal of comparative physiology A 193: 685-693

23) Olberg RM, Worthington AH, Fox JL, Bessette CE & Loosemore MP, 2005, *Prey size selection and distance estimation in foraging adult dragonflies*. Journal of comparative physiology A 191: 791-797

24) O'Carroll DC, 1993, *Feature-detecting neurons in dragonflies*. Nature 362(6420): 541

25) O'Carroll DC & Wiederman SD, 2014, *Contrast sensitivity and the detection of moving patterns and features.* Philosophical Transactions of the Royal Society B, 369(1636)

26) Palmer DB, Yamani Y, Bobrow TL, Karpinsky NC & Krusienski DJ, 2018, *Transient Signals and Inattentional Blindness in a Multi-object Tracking Task*. i-Perception 9.1; 2041669518754595

27) Pollack GS, 1988, *Selective attention in an insect auditory neuron*. Journal of Neuroscience 8(7): 2635-2639

28) Psychophysics Toolbox, RRID: SCR_002881, available: http://psychtoolbox.org/

29) Recanzone GH, Wurtz RH & Schwarz U, 1997, *Responses of MT and MST neurons to one and two moving objects in the receptive field*. Journal of Neurophysiology 78(6): 2904-2915

30) Remington RW, Johnston JC & Yantis S, 1992, *Involuntary attentional capture by abrupt onsets*. Perception & Psychophysics 51(3): 279-290

31) Reynolds JH & Desimone R, 2003, *Interacting Roles of Attention and Visual Salience in V4*. Neuron 37: 853-863

32) Reynolds JH, Pasternak T & Desimone R, 2000, *Attention increases Sensitivity of V4 Neurons*. Neuron 26(3)

33) Romer H, Hedwig B & Ott SR, 2002, *Contralateral inhibition as a sensory bias: the neural basis for a female preference in a synchronously calling bushcricket, Mecopoda elongata*. European Journal of Neuroscience 15(10): 1655-1662

34) Ruthruff E & Gaspelin N, 2018, *Immunity to attentional capture at ignored locations*. Attention, Perception & psychophysics 80(2): 325-336

35) Sareen P, Wolf R & Heisenberg M, 2011, *Attracting the Attention of a fly.* Proceedings of the National Academy of Sciences, 108(17): 7230-7235

36) Sawaki R & Luck SJ, 2010, *Capture versus suppression of attention by salient singletons: Electrophysiological evidence for an automatic attend-to-me signal.* Attention, Perception & Psychophysics 72(6): 1455-1470

37) Schiller PH & Lee K, 1991, *The role of the primate extrastriate area V4 in vision.* Science 251(4998): 1251-1253

38) Simons DJ & Chabris CF, 1999, *Gorillas in our midst: Sustained Inattentional blindness for dynamic events*. Perception 28(9): 1059-1074

39) Snedden WA & Greenfield MD, 1998, *Females prefer leading males: relative call timing and sexual selection in katydid choruses*. Animal Behavour (56): 1091-1098

40) Tang S & Juusola M, 2010, *Intrinsic Activity in the Fly Brain Gates Visual Information during Behavioural Choices*. PLOS ONE 12(5): e14455

41) Treue S, 2001, *Neural correlates of attention in primate visual cortex*. Trends in Neurosciences 24(5): 295-300; doi: 10.1016/S0166-2236(00)01814-2

42) Treue S & Maunsell JHR, 1999, *Effects of Attention on the Processing of Motion in Macaque Middle Temporal and Medial Superior Temporal Visual Cortical Areas.* The Journal of Neuroscience 19(17): 7591-7602

43) Van De Poll MN, Zajaczkowski EL, Taylor GJ, Srinivasan MV & van Swinderen B, 2015, *Using an abstract geometry in virtual reality to explore choice behaviour: visual flicker preferences in honeybees*. J. Experimental Biology 218:3448-3460; doi: 10.1242/jeb.125138

44) Van Swinderen B, 2012, *Competing visual flicker reveals attention-like rivalry in the fly brain*. Frontiers in integrative neuroscience 6: 96

45) Wiederman SD, Fabian JM, Dunbier JR & O'Carroll DC, 2017, *A predictive focus of gain modulation encodes target trajectories in insect vision*. eLife 6; e26478

46) Wiederman SD & O'Carroll DC, 2013, *Selective Attention in an Insect Visual Neuron*. Current Biology 23:156-161

47) Wiederman SD, Shoemaker PA & O'Carroll DC, 2013, *Correlation between OFF and ON channels underlies dark target selectivity in an insect visual system*. J. Neuroscience 33(32); 13225-13232

48) Yantis S & Jonides J, 1984, *Abrupt visual onsets and selective attention: Evidence from visual search*. Journal of Experimental Psychology: Human Perception and performance 10(5): 601-621

7 Salience Invariance with Divisive Normalisation in Higher-Order Insect Neurons

7.1 Preamble

As with many PhD projects, finding one's niche can take some time. Early in my PhD I explored an interesting quirk of the properties of gaussian functions observed by my supervisor. To this end I more thoroughly explored these properties below. While this is the only published output of this line of research, it has not been entirely abandoned and will likely represent research conducted after my PhD. Suffice to say, this remains an unanswered question, but as a computational technique for retrieving lost ambiguous information, it remains a viable approach (you'll have to read below).

The following represents my own work (with the aid of my supervisors).

Statement of Authorship

Title of Paper	Salience Invariance with Divisive Normalisation in Higher-Order Insect Neurons
Publication Status	Published
Publication Details	Evans BJE, O'Carroll DC & Wiederman SD, 2016, *Salience invariance with divisive normalization in higher-order insect neurons*, 6th European Workshop on Visual Information Processing

Principle Author

Name of Principal Author	Bernard Evans		
Contribution to the Paper	Experiment conceptualisation, model development, model analysis, data interpretation, figure generation, manuscript authorship.		
Overall percentage (%)	70%		
Certification	This paper reports on original research I conducted during the period of my Higher Degree by Research candidature and is not subject to any obligations or contractual agreements with a third party that would constrain its inclusion in this thesis. I am the primary author of this paper.		
Signature		Date	14/12/18

Co-Author Contributions

By signing the Statement of Authorship, each author certifies that:

i. the candidate's stated contribution to the publication is accurate (as detailed above);
ii. permission is granted for the candidate in include the publication in the thesis; and
iii. the sum of all co-author contributions is equal to 100% less the candidate's stated contribution.

Name of Co-Author	David C O'Carroll		
Contribution to the Paper	Experiment conceptualisation, data interpretation, manuscript evaluation		
Signature		Date	4/12/2018

Name of Co-Author	Steven D Wiederman		
Contribution to the Paper	Experiment conceptualisation, data interpretation, manuscript evaluation		
Signature		Date	14/12/18

7.2 Abstract

We present a biologically inspired model for estimating the position of a moving target that is invariant to the target's contrast. Our model produces a monotonic relationship between position and output activity using a divisive normalization between the 'receptive fields' of two overlapping, wide-field, small-target motion detector (STMD) neurons. These visual neurons found in flying insects, likely underlie the impressive ability to pursue prey within cluttered environments. Individual STMD responses confound the properties of target contrast, size, velocity and position. Inspired by results from STMD recordings we developed a model using a division operation to overcome the inherent positional ambiguities of integrative neurons. We used genetic algorithms to determine the plausibility of such an operation arising and existing over multiple generations. This method allows the lost information to be recovered without needing additional neuronal pathways.

7.3 Introduction

A common trait in biological sensory processing is the filtering and extraction of salient higher-order stimulus components such as features or local motion, using nonlinear operations that confound desired information with extraneous factors, such as feature contrast and angular size. A good example are the Small Target Motion Detector (STMD) neurons of dragonflies and other flying insects. STMDs show exquisite tuning to small moving targets (1-3°), but their response also depends on target velocity and contrast, creating ambiguities in the neural coding of any of these parameters (O'Carroll & Wiederman 2014). This ambiguity manifests as broad tuning functions where varying any individual parameter (for example target size) leads to a peak response at an optimum size and a diminishing response as the target becomes either larger or smaller. Distinguishing between a large, dim target and a small, bright target becomes impossible for single neurons. Despite these problems, predatory insects require precise positional information to successfully track and pursue prey. Indeed, they demonstrate an outstanding capacity to do so, with dragonflies boasting a 97% successful capture rate (Olberg et al., 2000).

This leads to the obvious question of how insects reconcile the spatially ambiguous detection systems with the precision required to successfully track targets. This represents a more fundamental problem of classification. As the number of sensors integrated to encode more complex information increases, spatial information is often lost in the process (Kouh & Poggio 2008). Recovering the lost spatial information requires either a separate parallel

pathway or a recombination of existing integrated information to re-extract the lost information.

One potential solution is to use additional sensor units, allowing for inter-unit comparisons. A simple example from image processing is using neighbouring pixels to interpret data from the chosen pixel via spatial filtering. This process helps reduce noise of various kinds. A high-pass filter can bring out local detail invariant to overall scene contrast while a low pass filter does the opposite. In the biological context, this would result in the integration of multiple parallel units. In dragonflies, there are numerous classes of STMD neurons, each covering different sized regions of the dragonfly's field of view. It is hypothesized that many of the large-field STMDs are integrated from smaller STMD fields, supporting a parallel unit comparison model. However, this does not answer the question of how the integration operates.

An alternative approach is to compare a sensor to itself over time. This is a common occurrence in biological systems, which often encode changes in a sensor value rather than the absolute value itself (Mante et al., 2008; Troy & Enroth-Cugell 1993). This adaptation effect allows the system to ignore persistent stimuli in favour of novelty. It is also possible to reconstruct spatial information by inducing periodic movements in the sensor, allowing for hyper-acuity between two broadly tuned sensors (Viollet & Franceschini 2010) by repurposing temporal information into spatial information.

Assuming the parallel sensor approach, there are still many alternatives to normalize the inputs against salience parameters such as contrast and size. Normalization has been shown to be used in many different neural systems such as odorant intensity normalization (Luo et al., 2010; Olsen et al., 2010), discriminability normalization, contrast invariance (Heeger 1992; Busse et al., 2009), invariance to ambient light levels (Mante et al., 2008), and invariance to size, location or occlusion (Kouh & Poggio 2008). Simple examples include taking absolute or relative ratio differences between the sensor and its population mean. The latter is an example of divisive normalization (Luo et al., 2010).

One advantage of systems that employ divisive normalization is that they can ignore changes in parameters that affect multiple sensors equally. For example, if the input from two sensors both double, when one is divided by the other, the ratio is maintained. Thus, such systems encode changes relative to the norm rather than the absolute values of the measurements. However, divisive normalization has an inherent mathematical risk when the divisor approaches zero. Any noise in either the numerator or denominator becomes magnified

significantly in this circumstance. This is concerning in neurons where very low spike-rates are not uncommon. Any system based on a division by zero cannot be robust. Thus, while mathematical models are useful tools, it is necessary that their explanatory power is based on feasible biology. Fortunately, it has been shown that divisive normalization is a feasible neuronal operation (Kouh & Poggio 2008).

7.4 Methods

When using divisive normalization on two closely placed Gaussian tuning functions, there is a function that resembles a single saw-tooth with a gently rising and sharply descending edge (Figure 56). The rising edge becomes, in-essence, a monotonic relationship between two parameters (in our case, position and neuronal output).

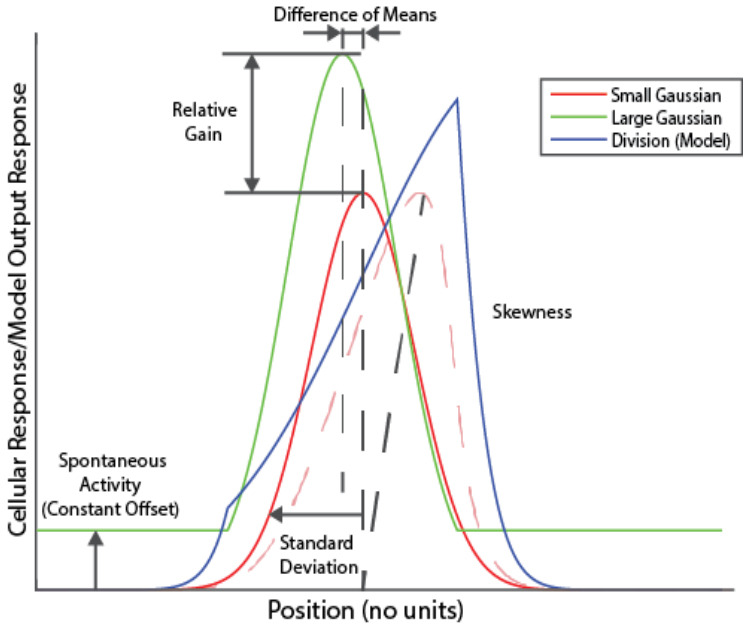

Figure 56: Theoretical model depicting divisive normalization on two one-dimensional receptive fields (small and large) and the resulting sawtooth output.

Using this as a base model, our approach was to generate a MATLAB model capable of extracting positional information from pairs of broadly tuned neurons. The model implemented one-dimensional parameterized versions of the receptive fields of each neuron (one small, one large). We used the divisive normalization process to generate the output signal (spike rate vs. position) of the system. Using this generic model, we then used evolutionary programming to optimize the resulting relationship based on a series of fitness criteria to establish the model's feasibility as a method for generating output invariant to target contrast.

7.4.1 Receptive field model

Receptive fields in visual neurons represent regions of physical space (defined by the relative angle to the insect's eye) in which presenting a stimulus will produce an excitatory or inhibitory response. Typically, these receptive fields have a region of strongest response in a nominal centre with responses weakening laterally until indistinguishable from spontaneous activity (Figure 57A). Insect STMD neurons exhibit a variety of different receptive field forms in terms of their size and shape, some subtending angles of more than 100 degrees while others are only a few degrees wide. STMD classes are also distinguished by their tuning for targets of different size, speed and direction of motion.

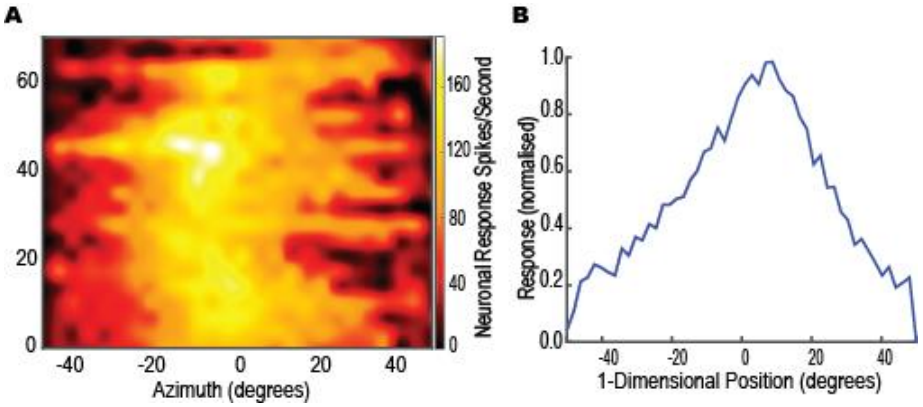

Figure 57: A) Receptive field map showing spike-rate response to a small black target being drifted through the dragonfly's field of view. B) 1-Dimensional horizontal cross-sections of dragonfly receptive field normalized and averaged.

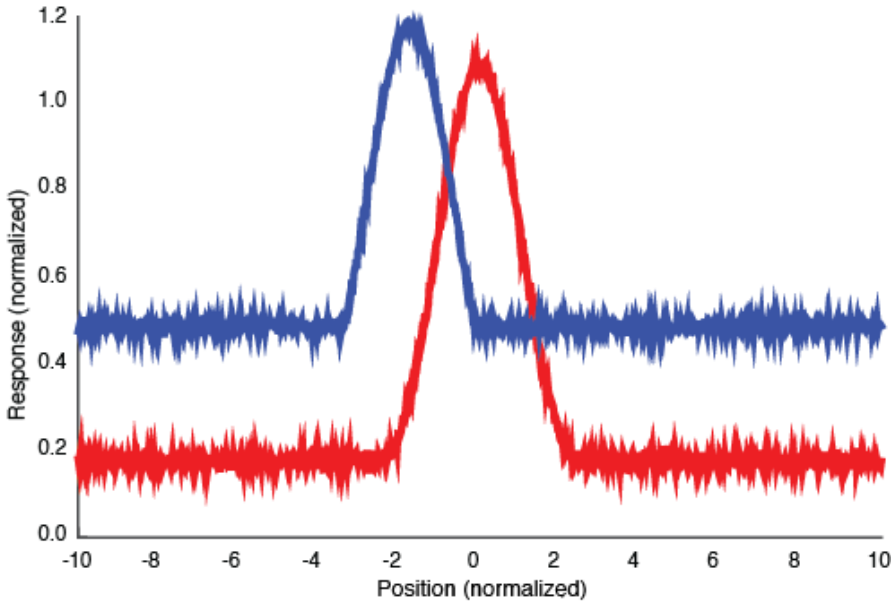

Figure 58: Receptive Fields modelled as 1-dimensional Gaussian functions with added white Gaussian noise.

We modelled each receptive field as a Gaussian distribution of position versus response (nominally spike-rate) representing a cross-section of a classical receptive field (Figure 57B). The receptive field models were allowed to vary based on their mean, standard deviation, gain and skewness.

In order to account for noise in the imperfect sensory system, the receptive fields were also subject to varying levels of Gaussian noise to demonstrate robustness of the system (Figure 58).

7.4.2 Divisive normalization

The model performed a simple divisive normalization process using the larger of the two receptive fields as the denominator and the smaller receptive field as the numerator. Individual points along each curve were divided and the resulting curve examined for fitness.

7.4.3 Saturating nonlinearities

One of the limitations of biological systems is that the maximum spike-rate of any individual neuron is limited by physiological constraints. Between individual neuronal spikes, there is a 'refractory period' preventing further spikes, which leads to a maximum spike rate. Unlike electronic signals, there are no sharp cut-offs at the saturation threshold for this kind of saturation (i.e. clipping) and instead, biological systems tend to encode information non-linearly and have 'rounded edges' with a gentle flattening of the response tuning in the vicinity of the peak (Figure 59B). These saturation points show variation between different neurons.

Our model implemented soft-saturating non-linearities using a hyperbolic tangent function (tanh) to set a theoretical maximum saturation point.

7.4.4 Spike threshold and spontaneous activity

Many neurons exhibit what is referred to as spontaneous activity. This indicates that the neuron in question spontaneously generates 'spikes' without any stimulus being presented (Figure 60), indicating that the threshold for spike generation is close to their resting potential. Spontaneous activity has a number of useful properties in neural circuits, of particular note; it allows neurons to encode negative (inhibitory) stimuli via a spike rate below the spontaneous level. This would be impossible for a neuron that has no spontaneous activity (negative spike-rates do not exist). It is common for cells with larger receptive fields to exhibit higher spike-rates than those with smaller receptive fields.

Modelling spontaneous activity however is not completely straightforward. There are two possible approaches, the first being purely additive (Figure 59C), simply adding a constant to the receptive field. The second option is to have spontaneous activity act a noise-floor that the receptive field falls below (Figure 59D). Alternatively, both methods could work in combination.

Our model uses both additive and noise-floor implementations for each receptive field used in the calculations.

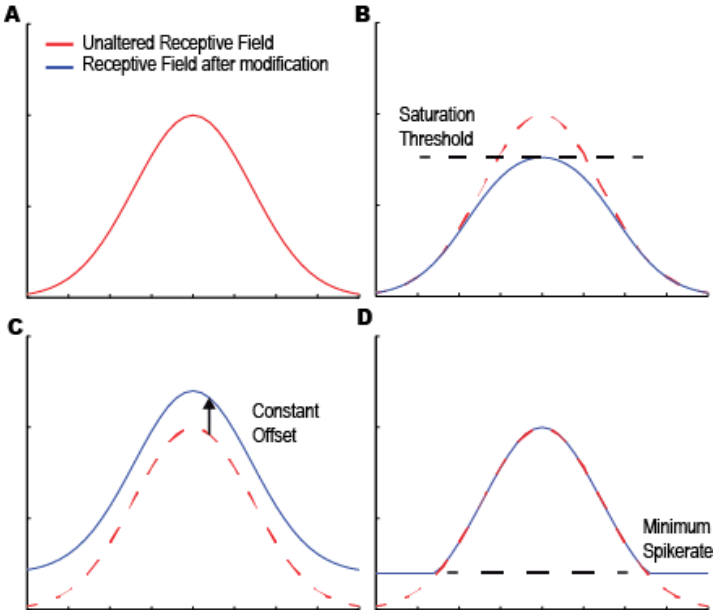

Figure 59: A) Unaltered Receptive Field. B) Receptive field subject to saturating non-linearity (hyperbolic tangent). C) Receptive field subject to constant offset. D) Receptive field subject to minimum spike rate

7.4.5 Evolutionary algorithm model

Evolutionary algorithms provide a good approach for solving optimization problems where the form of the final solution is not explicitly understood, but where there are clear ways of evaluating one solution over another. Moreover, in biological systems, any proposed model should be biologically plausible, making evolutionary algorithms particularly appealing (i.e. could this system evolve in nature?). As such, evolutionary algorithms were chosen to allow optimization of a number of parameters, while allowing the system to find its own viable solutions.

Our model implemented genetic programming, representing each of the candidate receptive fields as a series of parameters (or 'strain'). Each strain 'mutated' into four distinct child strains, which after divisive normalization were evaluated against fitness criteria (Figure 61). Child strains were 'culled' until the original number of strains was reached. We took the

approach of removing unfit children rather than selecting fit children to enable multiple fitness criteria to be considered easily.

Numerous parameters of each receptive field were subject to mutation, including: mean; standard deviation; difference of means; gain; skewness; saturation; minimum spike-rate, and spontaneous activity. In order to simplify the solutions, only the mean of the large receptive field was allowed to vary, with the smaller held constant (to prevent either wandering left or right). Likewise, the gain of the smaller receptive field was held constant with the relative gains allowed to vary via changes in the larger receptive field.

In order to determine the fitness of any particular strain the four following criteria were used: linearity, domain, signal compared to spontaneous behaviour and noisiness of the position estimate.

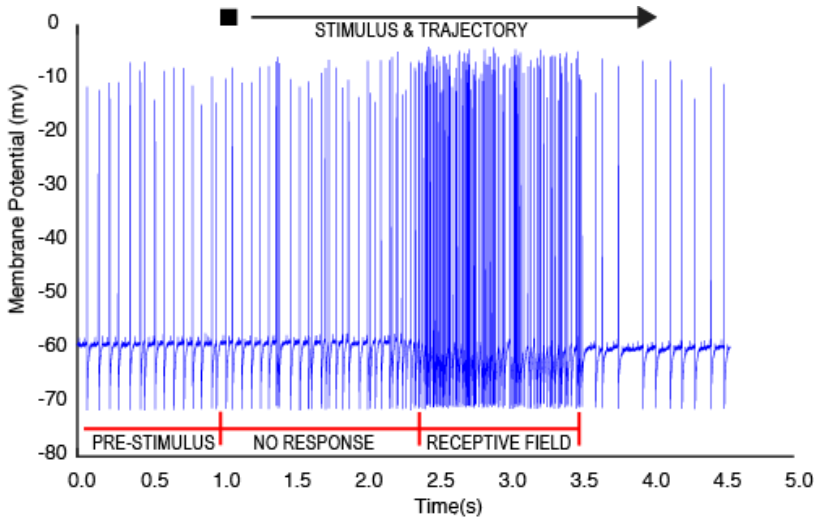

Figure 60: Example of Spontaneous Activity in a spiking neuron during the pre-stimulus interval.

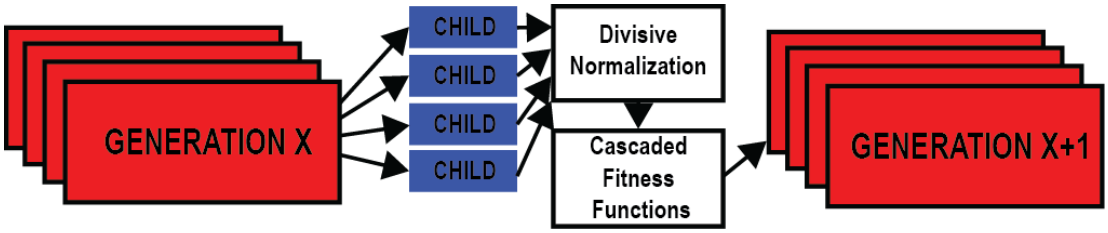

Figure 61: Model Process Chain

Many biological systems encode information non-linearly (Brinkworth et al., 2007), however, we chose to optimize our strains so that the relationship between output spike-rate and position was as close to a linear relationship as possible (Figure 62). This selection criterion

has the effect of selecting for monotonic functions and thus preventing the system from destroying any achieved position invariance through introduced ambiguity.

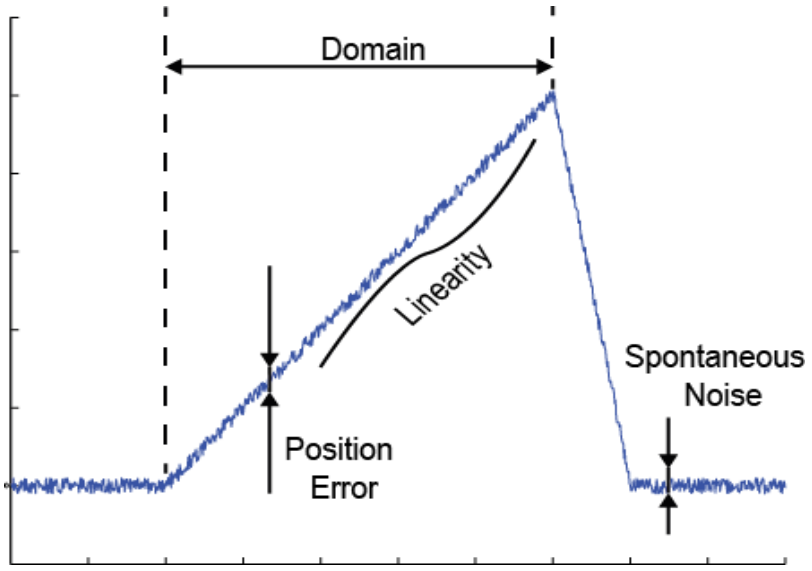

Figure 62: Example of fitness criteria used for child selection in evolutionary algorithm

The second main criterion was to maximize the domain over which the position invariance is useful (Figure 62). By maximizing the positional range over which the system calculates position, it minimizes the number of neurons required to perform the function over a given region of space. Likewise, there would be little purpose in a system than only accurately predicted position in very narrow bounds.

The third criterion was the signal-to-noise ratio as defined as the response difference between the start and end of the position-invariant region, divided by noise within the spontaneous activity (i.e. when neither receptive field was being stimulated). The purpose of this measurement is to ensure that the linear region of the resulting curve is easily distinguished from the surrounding regions (i.e. there is no target in the receptive field).

The final criterion was to minimize the noise in the position-invariant section of the output (Figure 62). Noise within this section of the output signal creates positional ambiguity and thus should be minimized where possible to give the most accurate estimate of target position.

Each simulation involved between ten and forty strains as an individual population. Each starting condition was tested using ten such populations. Dozens of different starting

conditions were tested to ascertain the convergence properties of the system. Each population was allowed to propagate for one hundred or more generations.

7.5 Results

7.5.1 Receptive fields

Results from the evolutionary model showed that it is indeed possible to generate a monotonic relationship between output spike-rate and target position from two overlapping receptive fields using divisive normalization with a simple fitness criteria. Moreover, as shown in Figure 63B, this relationship is maintained when the contrast of the target is varied (i.e. contrast invariance). Strictly speaking, the output generates an ambiguity on the right-hand side of the peak (the trailing edge of sawtooth); however, this demonstrates a significantly steeper slope and also falls outside the sensitive regions of either receptive field. As such, this section of the output could be relatively easily ignored via use of a threshold demanding a minimum response from either neuron before considering the output valid.

7.5.2 Convergence and robustness

Results from the model showed that the selected fitness criteria were well chosen for producing a robust end-point. Different populations starting from unique starting points tended to converge on very similar receptive field relationships, for example convergent means (Figure 63C) and standard deviations (Figure 63D).

An interesting consequence of using very broad fitness definitions is that while the solution strains did tend to converge towards very similar ‘ideal’ position detection systems, there is still quite a degree of variation in the end points. While for many other applications having such a broad spread of solutions may be considered a hindrance, in biological systems this instead represents another form of robustness, i.e. the system is not so finely tuned that small variations could destroy the relationship and render the system dysfunctional.

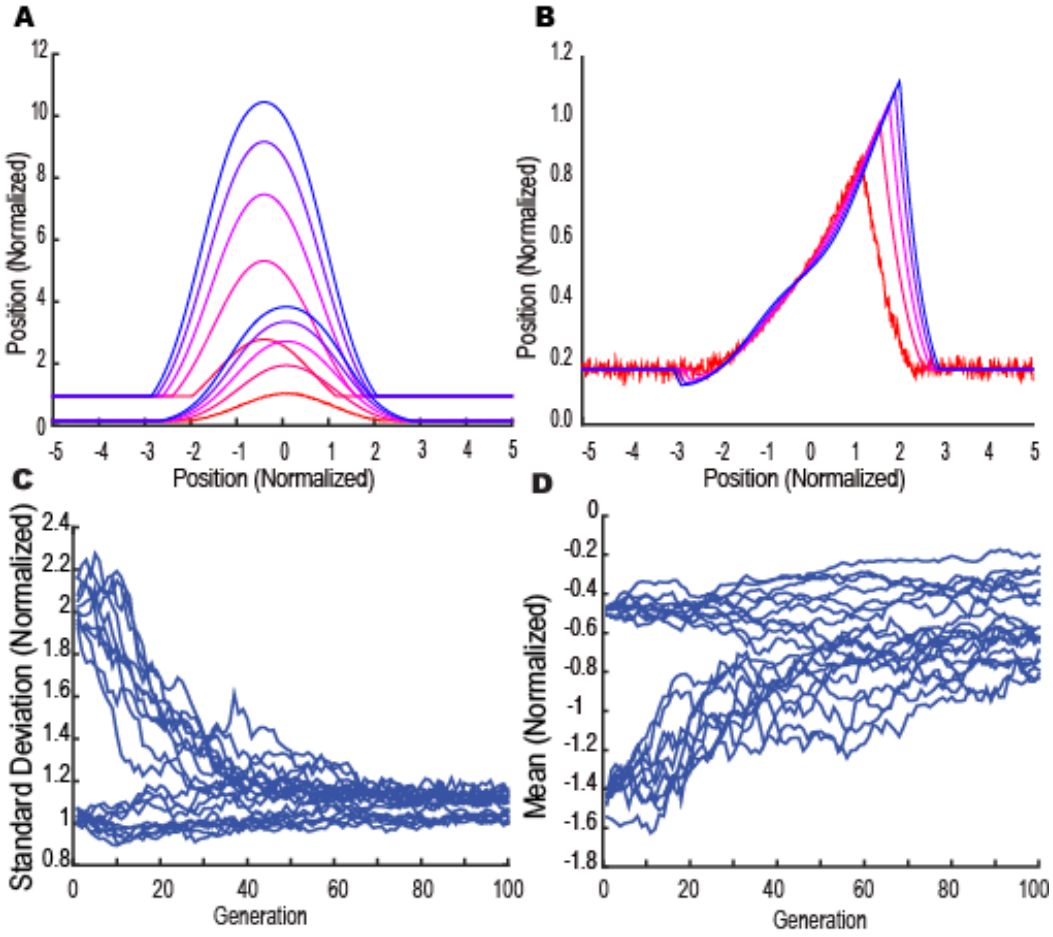

Figure 63: A) Example of final generation receptive fields. Variation in line colour represent varying contrast inputs. B) Example of final generation result of divisive normalization, invariant to the salience (e.g. contrast) of the input. C) Example of convergence of means between different initial starting conditions. D) Example of convergence of standard deviations between different starting conditions.

7.5.3 Spontaneous activity

An interesting result of the evolutionary model is that when noise in the output is minimized, the model tends to select for large receptive fields that exhibit higher spontaneous activity. This is intuitive mathematically, since a spontaneous spike-rate prevents the denominator from becoming near-zero and causing noise amplification. However, the fact that this mirrors the activity of many visual neurons with larger receptive fields gives weight to the feasibility of this model for contrast-invariant position calculation.

In order to have a strictly convergent spontaneous activity measure, it was necessary to first hold the gain of the larger receptive field constant (as they are inter-dependent). In this circumstance we found that there was a very strong convergence to a level well below the maximum output of the large receptive field (Figure 64A). This is again biologically

plausible, matching the behaviour observed during physiological recordings (Geurten et al., 2007).

Of the two forms of spontaneous activity, the model tends to select for large receptive fields that exhibit a minimum spike rate rather than simply adding a constant offset. In terms of the output shape, having a sharper corner at the base of the Gaussian curve leads to a sharper peak, as is seen in Figure 8B.

7.5.4 Saturating nonlinearities

If saturating non-linearities are allowed to vary freely, the model tends to select against having them (unsurprisingly). However, in the event that either receptive field has an enforced saturation, the other receptive field tends to attempt to match this value proportionally (Figure 64B). This again is intuitive from a mathematical perspective, but also emulates data captured during physiological recordings where STMD neurons with large receptive fields often exhibit higher maximum spike rates than their smaller counterparts (Geurten et al., 2007; Barnett et al., 2007).

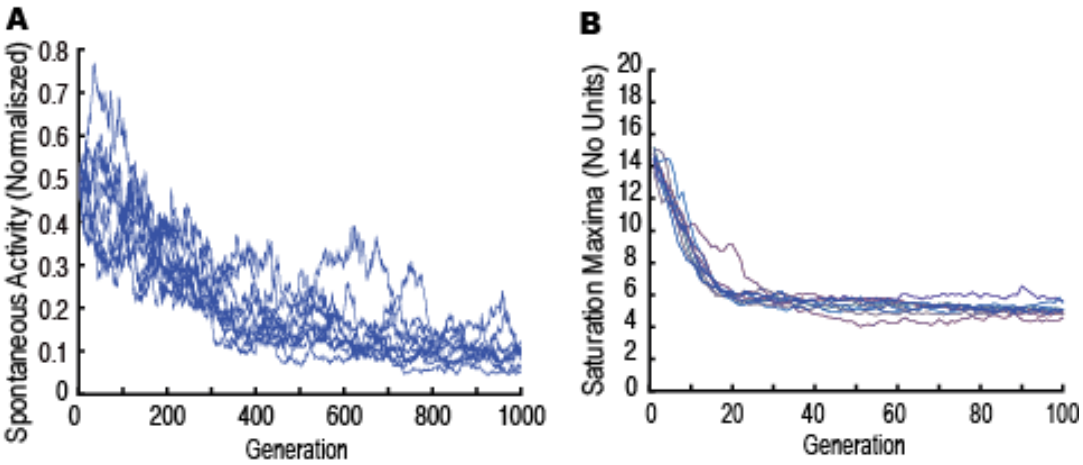

Figure 64: (A) Example of minimum spike rate selection in large receptive field, (B) Example of convergence of 'optimal' spontaneous activity point

7.5.5 Contrast invariance

In addition to producing a good model for position discrimination, the divisive normalization model also demonstrates a strong invariance to changes in target contrast. This makes the system particularly robust for natural pursuit conditions, where the relative contrast of a

target may change significantly due to background clutter (Wiederman et al., 2011, Bagheri et al., 2015).

We found that contrast invariance was maintained in all of our simulations. We tested our resulting receptive fields with a five-fold change in contrast input. Over the monotonic region of the output, this only resulted in deviations from the estimate of less than 3% (normalized against the range of discriminable output. For simulations than ran for longer periods, this dropped even lower to as small as a 1.5% error. This is a remarkable outcome given that only two receptive fields were used to generate this position signal and that contrast invariance wasn't an explicit selection criterion for the model. This is also comparable to the foveation (visual centring) behaviour found in dragonflies during pursuits (4-7° region of foveation) (Mischiati & Lin et al., 2015).

7.6 Discussion

7.6.1 Applicability for neuronal models

Results from the divisive normalization model show that it is indeed possible to establish a reliable relationship between neuronal response and position using two large-field STMD neurons. This model helps overcome some of the inherent difficulties of integrative neurons, which discard position information as they perform classification tasks. This model allows the lost information to be easily recovered without the need for additional neuronal pathways or complex integrative techniques. The divisive normalization step itself can be achieved relatively simply and is thus not metabolically expensive (Sterling & Laughlin 2015).

The model in question has further application in other domains beyond simple contrast invariance. Populations of cells that are tuned for different sizes, velocities or directions could also utilize similar techniques to generate precise estimates. In this way, a small population of integrated cells with broadly tuned behaviour can actually represent a large spectrum of different values, similarly to how human cone cells tuned to three colours allow a multitude of different colours to be experienced.

Moreover, these techniques could potentially be used in tracking applications that employ analogue sensors with broad tuning characteristics.

7.6.2 Future work

While the basic principle has been established by our work here, additional changes to the system will be needed to more specifically link the model to receptive fields measured from physiological recordings. Furthermore, our model addressed strictly one-dimensional cross-

sections of receptive fields. Further work will be required to model two-dimensional systems to determine whether the divisive normalization model continues to produce good results in these circumstance or whether further processing is required to allow for a two-dimensional model. Additional modelling will also be required to determine if timing considerations affect accurate calculation of the divisive normalization step including the effects of synaptic delays.

Furthermore, additional higher-level features of STMD neurons remain to be incorporated into this model, such as facilitation (Nordström et al., 2011) and selective attention (Wiederman et al., 2013).

7.7 References

20) Bagheri ZM, Wiederman SD, Cazzolato BS, Grainger S & O'Carroll DC, 2015, *Properties of neuronal facilitation that improve target tracking in natural pursuit simulations*, Interface (12: 20150083)

21) Barnett PD, Nordstrom K & O'Carroll DC, 2007, *Retinotopic organization of Small-Field-Target-Detecting Neurons in the Insect Visual System*, Current Biology (vol. 17, pp. 569-578), https://doi.org/10.1016/j.cub.2007.02.039

22) Brinkworth RS, Mah EL, Gray JP & O'Carroll DC, 2007, Photoreceptor processing improves salience facilitating small target detection in cluttered scenes, Journal of Vision (vol. 8, no. 11, pp. 1-17), https://doi.org/10.1167/8.11.8

23) Busse L, Wade AR & Carandini, 2009, Representation of concurrent stimuli by population activity in visual cortex, Neuron (vol. 64, pp. 931-924)

24) Geurten BRH, Nordström, Sprayberry JDH, Bolzon DM & O'Carroll DC, 2007, *Neural mechanisms underlying target detection in a dragonfly centrifugal neuron*, Journal of Experimental Biology (vol. 210, pp. 3277-3284), https://doi.org/10.1242/jeb.008425

25) Heeger DJ, 1992, Normalization of cell responses in cat striate cortex, Visual Neuroscience (vol. 9, pp. 181-197)

26) Kouh M & Poggio T, 2008, A canonical neural circuit for cortical nonlinear operations, Neural Computation (vol. 20, pp. 1427-1451)

27) Luo SX, Axel R & Abbott LF, 2010, Generating sparse and selective third-order responses in the olfactory system of the fly, Proceedings of the National Academy of Sciences (vol. 107, no. 23, pp. 10713-10718)

28) Mante V, Bonin V & Carandini M, 2008, Functional mechanisms shaping lateral geniculate responses to artificial and natural stimuli, Neuron (vol. 58, pp. 625-638)

29) Mischiati M, Lin HT, Herold P, Imler E, Olberg R & Leonardo A, 2015, *Internal models direct dragonfly interception steering*, Nature (vol. 517, pp. 333-338), https://doi.org10.1038/nature14045

30) Nordström K, Bolzon DM & O'Carroll DC, 2011, *Spatial facilitation by a high-performance target-detecting neuron*, Biology Letters (vol. 7, no. 4, pp. 588-592), https://doi.org/10.1098/rsbl.2010.1152

31) O'Carroll DC & Wiederman SD, 2014, Contrast sensitivity and the detection of moving patterns and features, Philosophical Transactions of The Royal Society B (369: 20130043) https://doi.org/10.1098/rstb.2013.0043

32) Olberg RM, Worthington AH & Venator KR, 2000, Prey pursuit and interception in dragonflies, Journal of Comparative Physiology A (vol. 186, no. 2, pp. 155-162), https://doi.org/10.1007/s003590050015

33) Olsen SR, Bhandawat V & Wilson RI, 2010, Divisive normalization in olfactory population codes. Neuron (vol. 66, pp. 287-299)

34) Sterling P & Laughlin S, 2015, *Principles of Neural Design*, Massachusetts Institute of Technology, Boston US.

35) Troy JB, Enroth-Cugell C, 1993, X and Y ganglion cells inform the cat's brain about contrast in the retinal image, Experimental Brain Research (vol. 93, pp. 383-390)

36) Viollet S & Franceschini N, 2010, *A hyperacute optical position sensor based on biomimetic retinal micro-scanning*, Sensors and Actuators (vol. 160, pp. 60-68)

37) Wiederman SD & O'Carroll DC, 2011, *Discrimination of Features in Natural Scenes by a Dragonfly Neuron*, Journal of Neuroscience (vol. 31, no. 19, pp. 7141-7144)

38) Wiederman SD & O'Carroll DC, 2013, *Selective Attention in an Insect Visual Neuron*, Current Biology (vol. 23, no. 2, pp. 156-161), https://doi.org/10.1016/j.cub.2012.11.048

8 Conclusion

8.1 Wide-field Neurons in Dragonflies

Prior to this work, there was no published research in the wide-field sensitive neurons of the dragonfly lobula (or other neuropils for that matter). Though originally a side-project of limited scope, the apparent complexity of dragonfly Lobula Tangential Cells (LTCs) turned into an entirely new area for researching unique systems (rather than rehashing work from other species).

Building on previous work (Fabian et al., in submission), anatomical staining of LTCs has demonstrated that the dragonfly has two lobula subregions involved in the processing of wide-field motion. Though previously argued that the Lobula Plate and Sublobula were homologous (Strausfeld 2005), it is clear now that with some LTCs arborizing in one, and others in both that the evolutionary picture is far more complex. Separately, each of these areas has been found to perform a similar role in different insect species (Lobula Plate – flies, Sublobula – bees). That the more ancient dragonfly utilizes both areas indicates that perhaps this complexity was redundant (and thus disappeared) in more modern specialists.

Of greatest interest is the separation of motion-adaptation properties between different subclasses of LTC. While the underlying mechanism that enables this difference was not explored here, it remains a fascinating question for future modelling research. Most importantly, it answers a fundamental question about vision processing, namely, how to optimize for two vastly different behavioural repertoires in a constrained neural architecture. Unlike humans, who simply replicate the hardware, the constrained dragonfly instead uses clever encoding strategies to (presumably) reuse existing hardware.

8.2 STMDs in Clutter

Previous research had demonstrated that STMD neurons perform remarkably well in cluttered environments (Nordström et al., 2006, Wiederman et al., 2011). However, even accounting for changes in local contrast, I found the performance of STMDs did deteriorate. Through my research, it has become clear that the target-tracking system is decoupled from the wide-field system and instead treats any natural image as a menagerie of competing targets vying to control ‘attention’.

Interestingly, despite natural images rarely containing perfectly tuned target-like features, there is still a significant reduction in response even with slowly moving backgrounds. Individual trials demonstrate that at least some of these instances are due to responses to

background features. Thus, it appears that background features that only produce modest responses in STMD neurons can still hold their attention 'captive'. This further embeds the idea of attention in these neurons as weakly salient features can be attended to despite 'crafted' high-salience targets being introduced. This can be be partially explained by experiements showing the importance of priming (i.e. temporal precedence).

How these finding relate to behaviour is unclear. In perching dragonflies (not investigated here), the importance of temporal precedence is quite explicable. When stationary, it is unlikely that any background features could strongly excite target-tracking neurons prior to take-off and thus favouring locking-on to old targets would help to maintain track on moving targets during the transition from perching to pursuit. In hawking dragonflies, a similar scenario exists if capturing prey from hovering. However, dragonflies can also capture prey during patrols where significant background motion (and thus false-positive responses) would exist. Thus, one might expect that prior to prey capture locking-on might be detrimental towards success in absence of switching behaviour. Fortunately, even in my limited data set there were numerous examples of switches between foreground and background features, which may help explain how the dragonfly overcomes this preference for old targets over new targets.

Finally, the key finding of my clutter experiments was the primacy of target velocity on selective attention. Fast moving targets appear to have preference over slow moving ones. Thus, it is quite possible that even while patrolling, potential targets may move in a way that gives them a higher salience (i.e. higher speed) enabling preferential switching onto targets over background features (which are fundamentally limited by the dragonfly's own ego-motion).

8.3 Selective Attention

How the finer details of selective attention operate remains a rich area for investigation but numerous aspects of it have been demonstrated here including a preference for targets which have preference for object precedence and objects which better match the underlying velocity tuning of STMD neurons. Interestingly, while strongly target-like features are rare in natural scenes it is apparent that what is weak for one neuron may be strong for another and the selective attention system appears content to track weakly excitatory and strongly inhibitory stimuli as well as perfectly crafted black squares (though it is perhaps less 'inclined' to do so). How this interacts with facilitation also remains unclear but given the moderate sensitisation and inhibition seen through facilitation it is unlikely to be the full picture.

I have also expanded the understanding of selective attention more generally. Firstly, it is now apparent that CSTMD1 is not the only neuron that exhibits the phenomenon with BSTMD2 (another large-field STMD) also exhibiting selective attention. Secondly, I have demonstrated that excitation isn't a prerequisite for selective attention with both CSTMD1 and BSTMD2 capable of selecting or rejecting inhibitory stimuli. This is of particular interest in CSTMD1 where previous research indicated a kind of 'inhibitory side takes all' response (Bolzon et al., 2009). While my experiments confirmed the apparent preference for inhibitory targets, priming experiments have demonstrated that the inhibitory hemifield of CSTMD1 is in fact part of the selective attention network, capable of being selected or rejected.

My experiments have also shown that selective attention can work across hemifields as well as within a single hemifield. They have also shown that targets moving in opposite directions can be selected indicating that selective attention is direction agnostic. Taken together, the presence of selective attention in two different neuron classes and the ability to cross hemispheres and select inhibitory targets indicate that selective attention likely occurs prior to large-field STMDs and that both CSTMD1 and BSTMD2 are simply the outputs of such a network.

8.4 Divisive Normalisation

Divisive Normalisation is a curious mathematical construct. While other examples of Divisive Normalisation have been shown in other species, they have focussed on entirely different purposes. The work here has demonstrated that some of the properties of asymmetric gaussian-like receptive fields could potentially have a purpose beyond being simply haphazard. Using multiple overlapping receptive fields allows signals once thought to be lost to be reconstructed whilst simultaneously removing ambiguities inherent in single encodings systems. While no biological evidence exists for divisive normalisation in insects, it remains a curious avenue of investigation, if not for the bio-mimeticists, certainly for bio-inspired engineers who want to convert ambiguous biological inputs back into something more tractable for engineering analysis.

8.5 Limitations

The electrophysiological recordings herein are subject to one significant limitation. The perils of intracellular electrophysiology demand (at least practically demands) the immobilisation of the dragonfly. This forces the system into a purely open-loop paradigm, which by its nature is limited. Without allowing the dragonfly to suitably react to the stimuli presented on screen it is difficult to guarantee that the responses recorded match those used in a natural context.

This represents a significant problem as it is already established that motion-sensitive neurons in other animal models are subject to change during locomotion (Longden & Krapp 2009; Maimon et al., 2010; Chiappe et al., 2010). While it may be possible in the future to examine these interactions in dragonflies, technology is currently incapable of such.

Moreover, while the purpose of this thesis was to investigate more realistic stimuli, via the use of natural images. However, all of our experiments rely on the use of simple monitors which can only stimulate subsections of an animal's field of view. Furthermore, while the general statistics of natural images share properties with external environments it is difficult to precisely replicate the intricate interplays of light and shadow, reflections and diffusions which occur in outdoor natural environments. Thus, while this represents a more realistic input, the pursuit of realism has a great deal further it could go. Additionally, the precise contrast variations and trajectories of targets as they move through the field of view represent best-guesses and simplifications from real insects moving through an environment. Light reflections, wing beats and idiosyncrasies of flight are not captured by small black squares moving on grey and white backgrounds. None-the-less, as with all science, even a simplified model can yield a great amount of insight.

8.6 Further Research

With such a broad range of topics touched on in this thesis, the directions of future research are equally broad.

8.6.1 Wide-field Neurons

It is clear from limited examples that the different subtypes of TCs have significantly different anatomy. One obvious area for further exploration is a more detailed account of the anatomy of TCs and in particular, relating the different anatomy back to their differing physiology. In essence, whether or not the additional arborization is responsible for differential motion adaptation and how that mechanism is implemented neuronally.

On a purely physiological front, there are numerous questions that remain unanswered. While the basic properties of TCs have been elucidated, the particular nuances of the adaptation differences are only partially explored. Specifically, the interactions between temporal and spatial frequencies as well as time. Additionally, the interactions of these systems to more complex visual inputs which contain objects moving at different rates as is found in real 3-dimensional environments remains to be explored.

In flies, it has been shown that the tuning optima of the ON and OFF pathways are asymmetric. As FATC and SFATC neurons exhibit different responses to ON and OFF flicker (FATC responding to ON and OFF while SFATC only responding to OFF) finding differences in the tuning for ON and OFF edges may yield further insight into the fundamental differences between these subtypes.

Finally, while numerous simple models for LPTCs and other WFMS neurons exist, the particular implementation of dragonfly TCs and their differential motion adaptation remains an area for further investigation.

8.7 STMDs in Clutter

Having largely answered the fundamental question (how do STMDs interact with clutter) further research in this area will likely be skewed far more to an investigation into the mechanisms of selective attention. Some obvious follow-ons include, further investigation of selective attention across hemifields especially with regard to how facilitation operates and whether each hemifield uses its own form of facilitation. Additionally, the observation that STMDs appear happy to select sub-optimal background features over carefully selected target features indicates that the ‘selection criteria’ of selective attention requires significant further investigation. This may help explain how STMDs are intended to function in complex scenes, especially those with many distracting features.

8.8 Divisive Normalisation

While the basic mechanism has been well established, how one might integrate divisive normalisation into a larger modelling context remains unclear. Arrays of divisive normalisation pairs do produce interesting outputs (unpublished observation). How these integrate into a two-dimensional context as well as control systems for feedback remain unclear and an area ripe for research.

9 Appendix I: Quantifying Asynchrony of Multiple Cameras using Aliased Optical Devices

9.1 Preamble

Early in this project, one aim was to record behaviour of dragonflies using multiple off-the-shelf consumer-quality cameras. One important challenge in this approach is fusing data from multiple sources, specifically the synchronization of the cameras. While a hardware approach to this problem is the most reliable, it was unavailable early in the project. Thus, we instead developed a visually-based system designed to allow a good approximation of the camera asynchrony. It is important to understand that this asynchrony is not as simple as working out how many frames out-of-sync each camera was but relied on the sub-frame asynchrony. This was due to the rapid motion of our test subject where even a fraction of a frame could represent large displacements of the subject.

Statement of Authorship

Title of Paper	Quantifying Asynchrony of Multiple Cameras using Aliased Optical Devices
Publication Status	Published
Publication Details	Evans BJE, Parslow BA, O'Carroll DC & Wiederman SD, 2015, *Quantifying Asynchrony of Multiple Cameras using Aliased Optical Devices*, International Conference on Image Processing Theory, Tools and Applications

Principle Author

Name of Principal Author	Bernard Evans		
Contribution to the Paper	Experiment conceptualisation, data collection, model development, model analysis, data interpretation, figure generation, manuscript authorship.		
Overall percentage (%)	70%		
Certification	This paper reports on original research I conducted during the period of my Higher Degree by Research candidature and is not subject to any obligations or contractual agreements with a third party that would constrain its inclusion in this thesis. I am the primary author of this paper.		
Signature		Date	14/12/18

Co-Author Contributions

By signing the Statement of Authorship, each author certifies that:

i. the candidate's stated contribution to the publication is accurate (as detailed above);
ii. permission is granted for the candidate in include the publication in the thesis; and
iii. the sum of all co-author contributions is equal to 100% less the candidate's stated contribution.

Name of Co-Author	Ben A Parslow		
Contribution to the Paper	Experiment conceptualisation, data interpretation, manuscript evaluation		
Signature		Date	26.11.18

Name of Co-Author	David C O'Carroll		
Contribution to the Paper	Experiment conceptualisation, data interpretation, manuscript evaluation		
Signature		Date	4/12/2018

Name of Co-Author	Steven D Wiederman		
Contribution to the Paper	Experiment conceptualisation, data interpretation, manuscript evaluation		
Signature		Date	14/12/18

9.2 Abstract

We present a new method for estimating asynchrony between off-the-shelf consumer-targeted camera systems. The system utilizes local aliasing induced by oscillating LEDs to perform post-hoc calculations. The system is capable of determining the asynchrony between cameras down to fidelity of less than a tenth of a frame. This system has numerous applications in the field of behavioural analysis of animals using arrays of inexpensive cameras.

9.3 Introduction

Video cameras and image processing techniques assist in the study of animal behaviour, but the spatial and temporal resolution of the captured video is a limiting factor for behavioural analysis. For example, an aim of our laboratory is to reconstruct predator-prey interactions in the field between flying insects. This is a significant challenge given their small size and dynamic behaviour over a large volume of their surrounding world. While it is possible to recreate pseudo-natural laboratory environments for some insects (Mischiati & Lin 2015), for others (e.g. hawking dragonflies) this remains practically impossible. Quite simply, the animals under investigation will not exhibit their full behavioural repertoire in anything other than real-world conditions.

For our purposes, we require data that captures target detection, tracking and selection (amidst multiple targets). We can then place these results in context of our electrophysiological recordings from target-detecting neurons likely to underlie such behaviours (Wiederman et al., 2013; O'Carroll & Wiederman 2014; Wiederman & O'Carroll 2013).

The expense, size, immobility and limited field of view of traditional high-speed cameras has constrained practical recording of insect behaviour for such tasks. However, with the advent of consumer-targeted video cameras with high spatial resolution and reasonable frame-rates, an opportunity has arisen to utilize them in ethological experiments.

Consumer-targeted video cameras have several advantages for capturing insect behaviour over their more expensive counterparts. They are, typically compact, robust, and unobtrusive. They carry low cost and highly portable power and data-storage options on board. Moreover, their low cost permits the use of several (perhaps dozens) cameras simultaneously, capturing large volumes of space from multiple points of view.

One significant challenge of using multiple cameras is ensuring that either the cameras are synchronized or that their asynchrony is known. Assuming arbitrary starting times, two

cameras with identical frame-rates can be up to half a frame out-of-sync. In many applications this asynchrony is not relevant. However, in highly dynamic scenes objects can move several dozen pixels between individual frames. Therefore, small timing errors correspond to significant errors in localizing features of interest when using interpolation techniques.

There are two broad strategies for ensuring synchrony between multiple cameras, namely (i), to control the timing behavior of the cameras externally using a common time-keeping system or (ii) to use information extracted from the captured data to back-calculate the asynchrony. The former hardware approach has a number of practical limitations, especially since most consumer electronics do not provide native support for hardware synchronization (Bradley et al., 2009). Those cameras that do provide this capability (e.g. Marshall GenLock), do not support the high frame rates (up to 240 fps at 720p resolution) possible with cameras such as the Hero3/4 series (GoPro, Inc.) and iPhone 6 Plus (Apple, Inc). Additionally, hardware solutions impose physical constraints upon the potential designs of any multi-camera gantries whether on the physical arrangement (Wilburn et al., 2005) or by using external lighting constraints (Bradley et al., 2009).

Numerous approaches utilizing a post-processing resynchronization approach have been employed, including matching visual and audio features (Shrestha et al., 2010), using space-time interest points (Yan & Pollefeys 2004) and matching sequences using trajectories (Caspi et al., 2006). However, most do not achieve better than a nearest-frame asynchrony estimate and often rely on specific scene events (such as collisions), rigid objects (Wolf & Zoemt 2002) or camera motion to achieve their estimates. Moreover, the capture of insect behavior is predominantly in natural scenes with very few rigid objects. Motion of the insects themselves, wind-swept plants and water movements all make many geometry-based approaches impractical. Finally, these techniques rely heavily on shared information which places immediate limitations on the locations of cameras.

Here we describe a novel technique to calculate asynchrony by using a visual stimulus temporarily within the field of view of the multiple cameras. This approach employs a pair of LEDs (one flashing periodically, one with a luminance step) and properties of aliased signals to provide an estimate of sub-frame synchronization. This paper will describe the LED circuit and the theoretical framework underlying the technique. Furthermore, we validate our approach with a ‘falling ball’ experiment demonstrating that the prediction aligns the outputs of the two cameras with a precision and accuracy that exceeds the requirements for the

behavioral capture. We achieved estimates of asynchrony with fidelity of less than one tenth of a frame, confirming the potential for exploiting large arrays of such cameras to monitor a large volume of space at high resolution in space and time.

9.4 Methods

Our approach mimics that of the clapper board used to synchronize imagery and sound recording in the film industry. That technique results in a maximum precision limited by the camera's frame-rate (of up to half a frame). In order to accomplish an accurate sub-frame estimate of asynchrony a high-resolution clock could be placed in the field of view. However, this will interfere with the behaviour under investigation as the clock cannot be removed during image acquisition. Another obtrusive approach would be to sample a continuous visual stimulus in the field of view (e.g. a LED waveform) and garner temporal information via interpolation. In addition to the obtrusiveness, any interpolation techniques are limited by the spatial resolution and intensity sensitivity of the camera. The event being interpolated will either consist of a moving target and thus be limited by the speed of the target and the field of view of the camera or if intensity-based will be limited by the quantized encoding of light intensity (8-bit, 0-255).

Our alternative approach overcomes these limitations by extracting information over a series of frames, exploiting properties of the aliased signal sampled over an extended duration (10 seconds). With these multiple aliased signals (from each camera) we can deduce the sub-frame time offsets between the cameras.

9.4.1 Cameras and Video Processing

All experiments detailed in this project utilized Hero 3+ Black Edition cameras (GoPro, Inc.). The stock (fish-eye) lenses were replaced with longer focal length (5.4mm), 1/2.3" format MP-10 lenses (RageCams). Experiments were done using 720p resolution with narrow field of view filmed at a 120 frames per second. The camera's option Protune image processing settings were used to enable capture of footage in RAW mode. ISO was set to 400 and all other settings were disabled or switched to zero. An ISO of 400 was chosen in order to reduce noise and also to ensure the half-wavelength shutter speed was achieved. Files were un-encoded using GoPro Studio into raw .avi format before being read into MATLAB for post-processing.

9.4.2 Aliasing Model

Temporal aliasing is a phenomenon where periodic signals of different frequencies can appear identical when sampled. For example, either a 1 Hz or 121 Hz signal will produce an identical output when sampled at 120 Hz. Importantly, this aliasing process preserves the phase of the original signal. Thus, the phase difference between two aliased signals is equivalent to the phase difference between the original signals. This is illustrated in Figure 65A, depicting two 121 Hz signals sampled at 120 Hz, which both appear as 1 Hz signals. One signal has a 10% phase delay (λ/10 s). As can be seen in Figure 65A, this delay is preserved in the apparent 1 Hz signal.

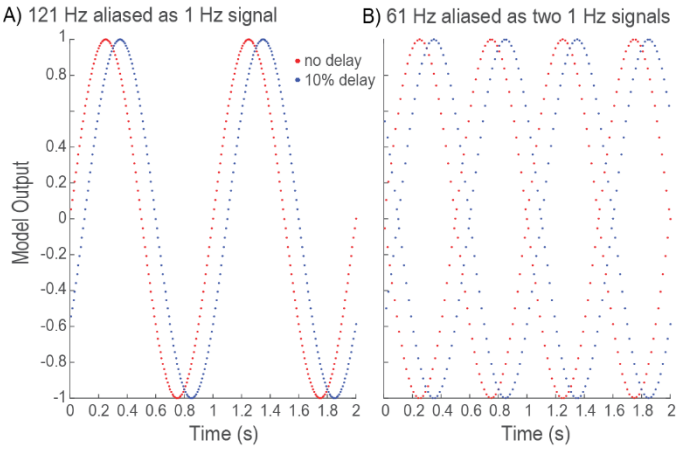

Figure 65: *Model example of the aliasing effect of* ***A)*** *a 121 Hz signal sampled at 120 frames per second with a resultant 1 Hz aliased signal. The underlying 121 Hz signals (not shown) have a 10% (λ/10 s) phase delay, which is also apparent in the aliased 1 Hz signals (red versus blue points).* ***B)*** *a 61 Hz signal sampled at 120 frames per second, generating two alternating 1 Hz signals. The 'double' aliased 1 Hz signals also preserve the phase difference of the original source.*

By comparing the phase delays of the aliased signals (1 Hz), the phase delay of the actual signal (121 Hz) can be deduced. The real-world change in LED intensity is fixed, which means that variation in the observed phase delays in each camera is the result of the sub-frame asynchrony between the cameras. A low-frequency aliased signal is preferred to accurately estimate the phase delay. It would seem intuitive to pick a frequency such as 121Hz (1Hz alias); however, this is subject to shutter speed limitations that are discussed in section 9.4.3. Instead we chose a frequency of 61 Hz, i.e. 1Hz above half the camera sample rate. This results in the appearance of two 1 Hz signals out of phase by 180 degrees (Figure 65B).

9.4.3 Integration Timing

Unlike idealized MATLAB models, cameras do not integrate instantaneously, instead capturing photons over a specified period. The cameras used in this study (GoPro) have a variable integration time dependent on scene luminance ranging from the full frame length (1/120s = 8.3ms) in low intensity scenes to a short 1/8192s in bright scenes.

In low light conditions, using an LED signal oscillating at 121 Hz results in an integration being performed over almost a full cycle (Figure 66A). Contrasting this with a 180 degree (half frame) shift of the same signal (Figure 66B), the value of the signal integrated over the full frame period (8.3ms) is only marginally different (red lines). Consequently, the resultant amplitude of the aliased oscillation would be too small to be detectable above the background noise.

Fortunately, it is possible to circumvent this problem by choosing a LED frequency slightly offset from half the camera sampling frequency (i.e. 61 Hz). When sampled, two components will be produced with a 180 degree phase shift. This is due to the sampling frequency being approximately twice the oscillating frequency producing two samples per LED cycle instead of one. As the signal is only integrated over half a cycle, the difference between the maximum (Figure 66C) and the minimum (Figure 66D) is more than 60% of the original LED dynamic range.

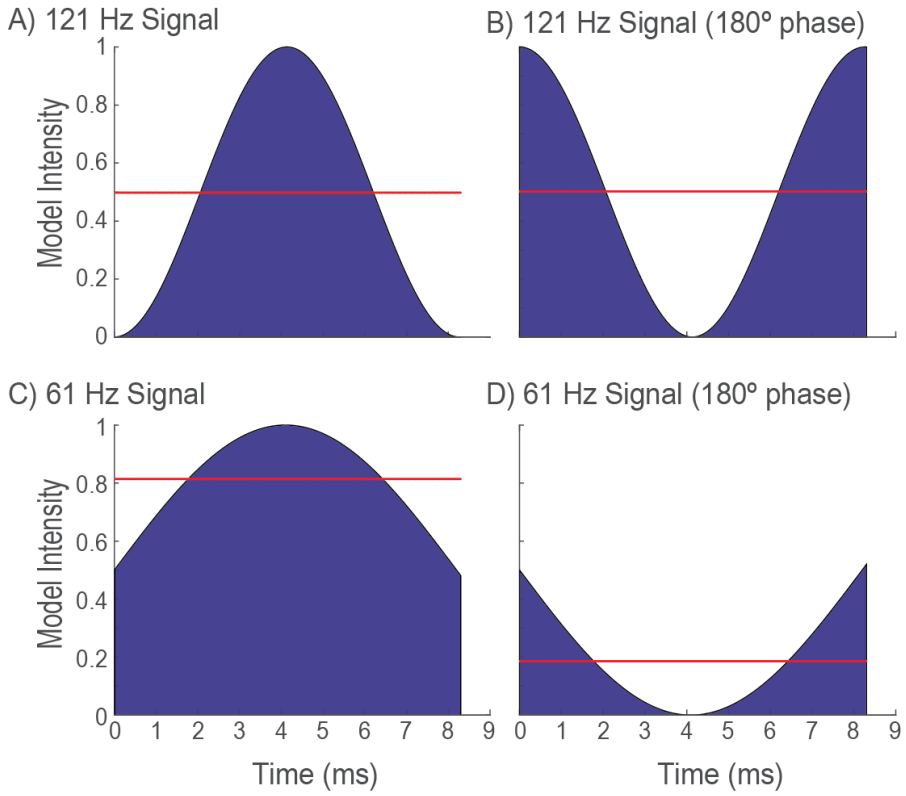

Figure 66: *Model output illustrating the effect of integration time on an oscillating signal (**A** & **B** : 121Hz, **C** &**D**: 61Hz, red line indicates the mean of the blue area). In **A** & **B**, the difference of the mean integration comparing 180 degrees phase shift is negligible. The 61Hz signal (**C** & **D**) displays a much higher dynamic range.*

While this approach solves the issue of detection, the presence of the second aliased frequency introduces an additional ambiguity to the asynchrony estimate. Fortunately, these two components are 180 degrees out of phase, so this ambiguity can be simply resolved with a second LED generating an exponential step over several frames.

9.4.4 LED Circuit

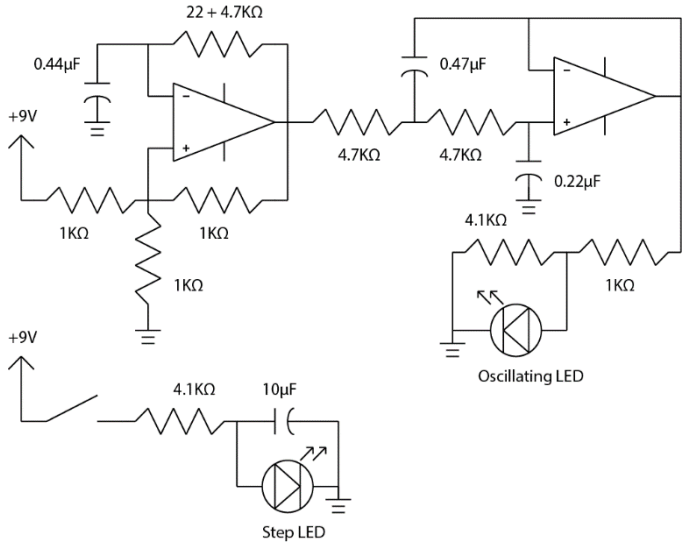

Figure 67: *LED Circuit Design: Fixed frequency LED oscillator and a simple slow-rise-time step LED.*

In order to provide a near-sinusoidal stimulus in the visual field of the cameras (given the nonlinearity of a LED), we therefore designed an oscillating circuit to drive a single LED (Figure 67 upper circuit). In addition to this 'blinking' LED a simple step LED (a switch activated once) permits synchronizing to the nearest frame (Figure 67 lower circuit). This 'step' LED incorporated a capacitor to ensure a rise time of more than a single frame. This allowed interpolation of the LED intensity resulting in an approximate estimate of the sub-frame synchronization (i.e. to the nearest half frame). This second LED allows us to remove any multiple-frame asynchrony (i.e. triggering latency from the GoPro WiFi remote control) and to resolve the ambiguity inherent in the 61 Hz aliased frequency. The LEDs were placed adjacent to one another to avoid any rolling shutter effects.

9.4.5 Data Processing

The video files were read into MATLAB using VideoReader. The location of the LEDs in each camera image were manually segmented to a 5x5 pixel region. The cropped regions were then averaged to create a simple intensity signal for i) the oscillating LED and ii) the step LED. This region of the image approximates the physical size of the LED's when viewed from 1 meter away. This averaged estimate of LED intensity reduces 8-bit quantization errors of the intensity values. For an intensity measure we also took the mean of the 3 RGB channels.

We took account of the 'half-frequency' aliasing by separating two candidate signals derived from alternating frames. Figure 68A shows the original sampled data from the two cameras (red and blue points) and Figure 68B and Figure 68C reveals that there are two possible signal pairings between two cameras. We analyzed both of these pairings to calculate the phase offset. We did this by applying a discrete Fast Fourier Transform (FFT) to the candidate, aliased signals (Figure 69). Despite nonlinearities in the recorded signal (due to a variety of cause), the fundamental frequency is easily distinguished.

In the phase domain of the FFT, we then determined the phase of this peak (fundamental) frequency. This represents our desired phase offset for a particular camera. At this stage we still have two candidates (for a single camera), and for each we can calculate the respective phase offsets between the two cameras (difference between FFT phase values).

After determining sub-frame asynchrony estimates, the step-LED signals were used to provide the gross frame synchronization. Each signal was normalized and interpolated to allow direct comparison. Candidate asynchrony estimates were then applied to the processed step-LED signals and the best match (via regression) was selected as the true estimate.

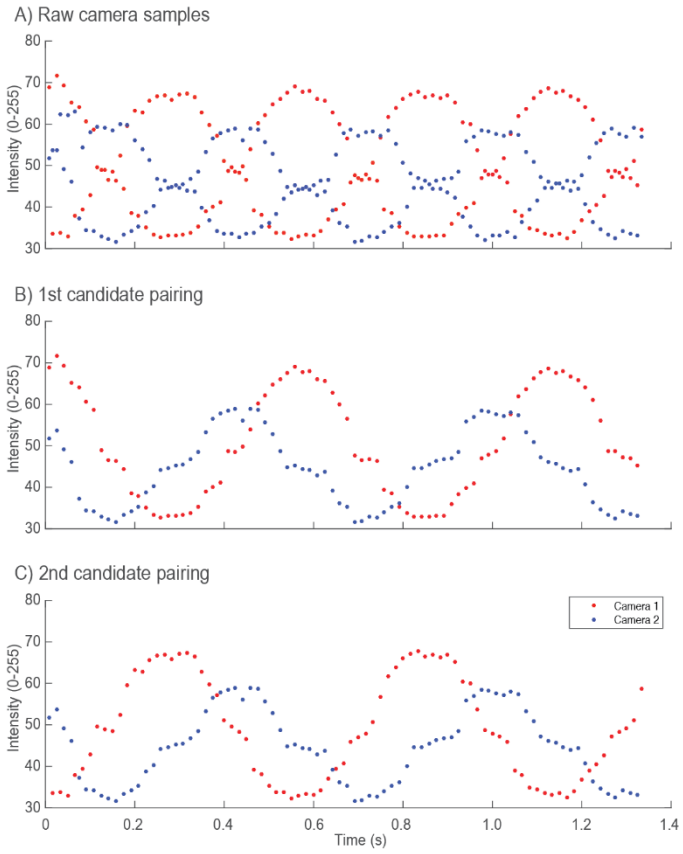

Figure 68: *Measured intensity samples from two cameras (red versus blue points) and the two candidate pairings of the aliased signals after separation (alternating points from A).*

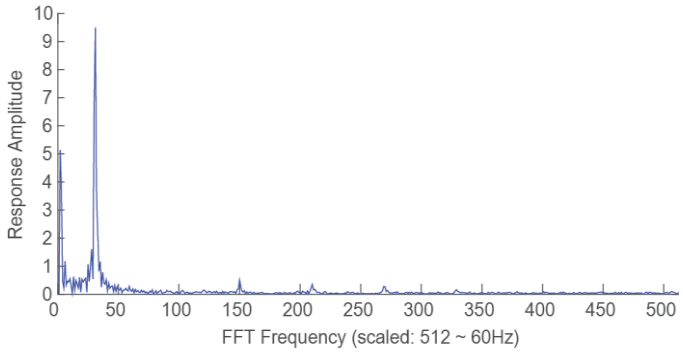

Figure 69: *Fast Fourier Transform depicting the fundamental frequency of the oscillating LED.*

9.4.6 Falling Ball Validation

We validated the asynchrony estimate by performing a simple falling ball experiment (Figure 70). This experiment involved two cameras, triggered using a standard GoPro WIFI remote control. The two cameras viewed the scene via a box-prism to ensure a similar point of view and removing the need for parallax correction. With both cameras recording, we turned on the LED device. We dropped a small white Styrofoam ball down a chute visible from both

cameras. Markers at either end of the chute allowed for image registration that accounted for any spatial offset in each cameras field of view.

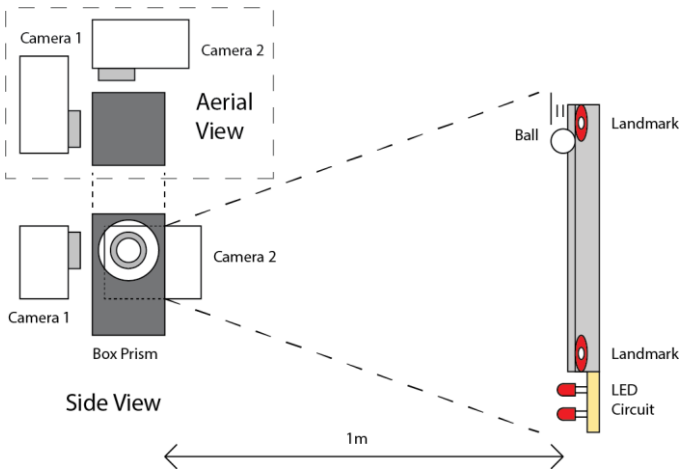

Figure 70: A Falling Ball experiment allowed for validation of our approach and an objective measure of the asynchrony between cameras. This value was compared to our estimated value to provide a measure of accuracy.

For each experiment, one video was flipped (to account for the box-prism reflection). The markers at each end of the chute were cross correlated to determine their relative offsets between images (these offset values differed due to minor rotations between camera points of view). We tracked the trajectory of the ball with a template-correlation match algorithm. The resulting trajectories were then corrected spatially by interpolating the calculated marker offsets to produce a spatial offset for each ball position. The trajectories were then temporally offset using the asynchrony estimate calculated using the LED circuit.

These offset trajectories (both in space and time) were then compared to establish validation for the technique.

9.5 Results

9.5.1 Falling Ball Results

We conducted eight falling ball experiments in low-light conditions. The camera and chute positions were not modified between experiments but the cameras were restarted to induce a range of asynchrony values. Results prior to and after the ball falling were discarded, as the algorithm tracked arbitrary points in the scene during this time (see Figure 71A artefacts in the blue line up to 100 ms). The difference between the predicted ball location from Camera 1 were aligned with Camera 2 were calculated for each point along the interpolated trajectory.

The mean difference between the predicted location (LED technique) and interpolated location (ball validation) was 0.65 pixels with a standard deviation of 0.31 pixels. With the ball moving on average 11.33 pixels/frame this equates to a temporal error of 0.5 ± 0.35ms (approximately a 16th of a frame). As this value is less than the potential error introduced during the alignment of the two images (which had single pixel fidelity) the limitation of this validation technique has been reached. Given that this technique will be applied to reconstructing spatial locations of insects, a sub-pixel level of accuracy is sufficient.

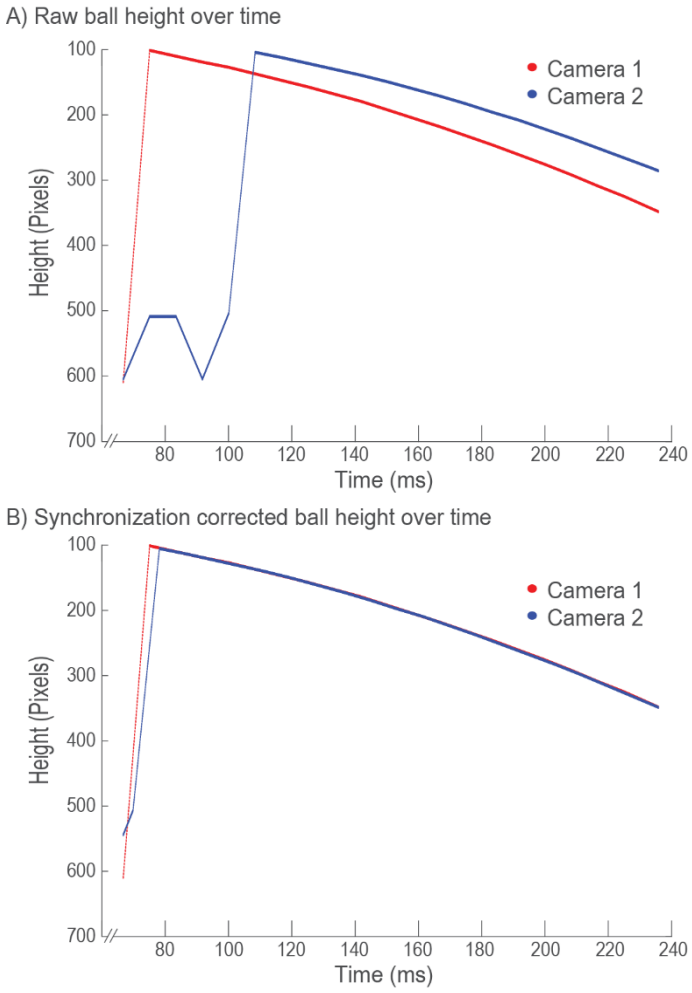

Figure 71: ***A)*** *The interpolated trajectory of a falling ball is reconstructed from two cameras (red and blue points).* ***B)*** *After correcting for the asynchrony using the LED aliased technique, the trajectories of the falling balls align, thus validating the method.*

9.5.2 Frame-rate Timing

While off-the-shelf cameras boast impressive frame-rates, there will inevitably be small differences between individual cameras of the same model. While these differences may be small and indeed negligible in typical consumer purposes, in the highly dynamic scenes

within insect pursuits, these small timing errors translate into significant spatial errors for reconstruction.

Small differences in frame-rates will compound over time resulting in the cameras slowly drifting from any initial synchronization. We quantified this effect using our LED synchronization method. A rolling sampling period of 10 seconds was taken over a one minute period of footage using two cameras. As shown in Figure 72, the frame asynchrony between the cameras slowly drifts over time in a linear fashion. Over a 50 second period, the asynchrony changes by approximately 1/5 of a frame, which is equivalent to having an extra frame in one of the cameras every five minutes.

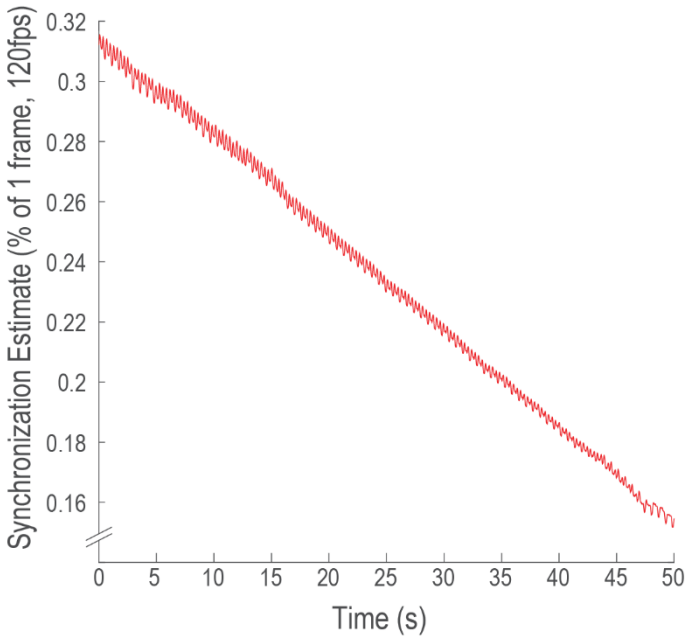

Figure 72: *Asynchrony estimate over extended period. The above figure depicts the result of minor differences in camera frame rates which results in a slow phase drift in synchronization over time.*

9.6 Discussion

9.6.1 Applicability for Insect Behaviour Capture

Results from the falling ball experiment demonstrate that the asynchrony offsets can be overcome to achieve subpixel accuracy for fast-moving objects. This is sufficiently accurate to allow reconstruction of insect trajectories using multiple cameras whose accuracy is fundamentally limited by the resolution of the camera involved. Additionally, as the LED stimuli only needs to be in the field of view for a limited period of time to establish the asynchrony estimate, it means that the synchronizing device can be used at the beginning of a recording and then removed subsequently to prevent the LED's interfering with insect behaviour during filming.

Furthermore, the technique does not rely on any specific scene geometry, scene events or hardware timing systems. This enables any camera gantry to be designed to accommodate the insect behaviour (rather than the algorithm) with cameras boasting different resolutions, fields of view and zoom levels. Being a light-weight device, this allows the system to be easily transported to the often remote and logistically difficult locations where insects live. Moreover, as the LED circuit is mobile it can even be used to calculate asynchrony between cameras with no or minimal overlap in their field of view or cameras which never have a suitable geometry even in their common field of view. This can be achieved by simply moving the circuit through the field of view of all of the cameras (perhaps mimicking the presumed trajectory of the subject). This will be especially useful for insects whose behaviour may necessitate cameras in multiple locations to follow long or convoluted trajectories and traverse large volumes of space. The processing techniques involved are not computationally expensive and are eminently scalable to any number of cameras.

This technique also helps overcome the slow phase drift identified in the long-duration experiment performed in this article. Correcting for these frequency misalignments is quite practical given that the rate of phase shift is quite low (less than 0.2 frames per minute) and follows a predictably linear trend. Thus, cameras can either be calibrated to pre-calculate the expected drift during experiments relative to one another (perhaps only selecting cameras with highly matched frame-rates) or multiple synchronization events could be used (i.e. before and after filming) to facilitate calculation of the interim phase difference. As most insect engagements are very brief (perhaps a few seconds) any phase difference during the actual engagement will likely be negligible.

9.6.2 Transferability and Lighting Conditions

The falling ball experiment was performed inside a laboratory under flat lighting; however this represents a fraction of the intensity compared to outdoor environments. In full daylight, normal LEDs do not provide sufficient dynamic contrast to be registered by Hero3+ sensors. However, this can be overcome through use of higher-power lighting (i.e. high power LEDs) or the use of a temporary shade to provide local contrast while the synchronization step is performed. Alternatively it is also possible to use other lighting technologies or incorporating the lighting system to directly stimulate the sensors internally via minor hardware adjustments.

Additionally, in daylight conditions, reduced shutters speeds will enable frequency choices which do not produce secondary aliases (e.g. 121Hz). It is also possible to extend this alias

theory to synchronize cameras with different frame rates assuming appropriate frequencies can be chosen to correctly produce the aliased signal.

9.6.3 Future Work

While the basic principle has been established, there are numerous improvements that can be yet made to any asynchrony estimating system. Firstly, many processes (such as detecting LEDs) can be automated to minimize the amount of post-processing setup time. Secondly, the dynamic range of the LED's intensity may need to be altered depending on the ambient lighting conditions. Frequency choice is also dependent on ambient lighting and thus further work is required to better characterize the effects of shutter speed on optimal frequency choice.

9.7 References

1) Bradley D, Atcheson B, Ihrke I & Heidrich W, 2009, *Synchronization and Rolling Shutter Compensation for Consumer Video Camera Arrays*, IEEE CVPR Workshops (1-8).
2) Caspi Y, Simakov D & Irani M, 2006, *Feature-Based Sequence-to-Sequence Matching*, International Journal of Computer Vision (vol. 68, no. 1, pp. 53-64)
3) Mischiati M, Lin HT, Herold P, Imler E, Olberg R & Leonardo A, 2015, *Internal models direct dragonfly interception steering*, Nature (vol. 517, pp. 333-338), https://doi.org10.1038/nature14045
4) O'Carroll DC & Wiederman SD, 2014, *Contrast sensitivity and the detection of moving patterns and features*, Philosophical Transactions of The Royal Society B (369: 20130043) https://doi.org/10.1098/rstb.2013.0043
5) Shrestha P, Barbieri M, Weda H & Sekulovski D, 2010, *Synchronization of Multiple Camera Videos using Audio-Visual Features*, IEEE Transactions on Multimedia (vol. 12, no. 1, pp. 79-92)
6) Wiederman SD & O'Carroll DC, 2013, *Selective Attention in an Insect Visual Neuron*, Current Biology (vol. 23, pp. 156-161), http://dx.doi.org/10.1016/j.cub.2012.11.048
7) Wiederman SD, Shoemaker PA & O'Carroll DC, 2013, *Correlation between OFF and ON Channels Underlies Dark Target Selectivity in an Insect Visual Systems*, The Journal of Neuroscience (vol. 33, no. 32, pp. 13225-13232), https://10.1523/JNEUROSCI.1277-13.2013

8) Wilburn B, Joshi N, Vaish V, Talvala E, Antunez E, Barth A, Adams A, Horowitz M & Levoy M, 2005, *High Performance Imaging Using Large Camera Arrays*, Proceedings of SIGGRAPH (765-776).

9) Wolf L & Zomet A, 2002, *Correspondence-free Synchronization and Reconstruction in a Non-Rigid Scene*, Proceedings Workshop on vision and Modelling of Dynamic Scenes

10) Yan J & Pollefeys M, 2004, *Video Synchronization via Space-Time Interest Point Distribution*, Proceedings of Advanced Concepts for Intelligent Vision Systems.

10 Appendix II: Multi-Focal Video Fusion with a Beam Splitter Prism

10.1 Preamble

Early in this project, one aim was to record behaviour of dragonflies using multiple off-the-shelf consumer-quality cameras. One important challenge in this approach is fusing data from multiple sources, specifically cameras from different view-points. To minimize the need for parallax compensation we sought to investigate the viability of using beam-splitters for aligning the field of view of two cameras. The technique proved entirely effective which created opportunities for data collection. Among these were increasing the time resolution (having two cameras taking alternate images) or especially the working distance (improving the depth available).

Statement of Authorship

Title of Paper	Quantifying Asynchrony of Multiple Cameras using Aliased Optical Devices
Publication Status	Published
Publication Details	Evans BJE, Parslow BA, O'Carroll DC & Wiederman SD, 2015, *Quantifying Asynchrony of Multiple Cameras using Aliased Optical Devices*, International Conference on Image Processing Theory, Tools and Applications

Co-Author

Name of Co-Author	Bernard Evans		
Contribution to the Paper	Experiment conceptualisation, data analysis, data interpretation, figure generation, manuscript evaluation.		
Overall percentage (%)	20%		
Certification	This paper reports on original research I conducted during the period of my Higher Degree by Research candidature and is not subject to any obligations or contractual agreements with a third party that would constrain its inclusion in this thesis. I am the second author of this paper.		
Signature		Date	14/12/2018

Co-Author Contributions

By signing the Statement of Authorship, each author certifies that:

i. the candidate's stated contribution to the publication is accurate (as detailed above);
ii. permission is granted for the candidate in include the publication in the thesis; and
iii. the sum of all co-author contributions is equal to 100% less the candidate's stated contribution.

Name of Primary Author	Ben A Parslow		
Contribution to the Paper	Experiment conceptualisation, data collection, data analysis, data interpretation, figure generation, data interpretation, manuscript evaluation		
Signature		Date	26.11.18

Name of Co-Author	David C O'Carroll		
Contribution to the Paper	Experiment conceptualisation, data interpretation, manuscript evaluation		
Signature		Date	4/12/2018

Name of Co-Author	Steven D Wiederman		
Contribution to the Paper	Experiment conceptualisation, data interpretation, manuscript evaluation		
Signature		Date	14/12/18

10.2 Abstract

This paper addresses the feasibility of using inexpensive, compact consumer-oriented cameras to resolve fundamental issues in capturing video footage of biological interactions in the field for quantitative analysis. The strengths and weaknesses of using a multi-camera beam splitter system to capture closely align video footage and its applications for studying biological interactions are discussed. The strength of the approach is that the closely aligned video streams are readily focal stacked to increase spatial information in a frame. Weaknesses include the limitation of the system in low light conditions where the beam splitter reduces the light captured by the cameras. Our novel system lends itself to incorporation into a camera array to capture large amounts of biological information.

10.3 Introduction

Many ethological studies examine animal behaviour within natural scenes, often quantified with video footage captured on location. The quality of such observations is limited by the cost and capabilities of the camera system. Traditionally, capturing fast moving biological interactions (e.g. insect flight) is done using expensive high frame rate cameras. These systems allow for variable frame rates, up to very high speeds at moderate resolution and with the ability to synchronize electronically with external devices (including triggers and other cameras). Beside their expense, the major down side is their bulk and power requirements, which impedes versatile field deployment. Each camera must trade-off lens focal length for improved resolution and working distance against limited depth of field and field of view (FOV). Therefore, most camera systems limit the volume over which animal behaviour can be studied in fine detail. Previously employed high-speed systems have therefore limited their detailed analysis to large artificial environments and constrained types of natural animal behaviour, where the precise location of the interactions is known in advance. For example, such systems have been used to study take-off and prey capture by perching dragonflies where only a single possible perching location was provided in the artificial arena (Mischiati & Lin et al., 2015).

Recently, less expensive cameras have become commercially available, with reasonable spatial and temporal resolution. We propose that it is feasible to capture relevant behaviours, e.g. complex predator–prey interactions, in a more natural setting by fusing recorded data from multiple cameras. By using large arrays of compact, consumer level cameras with a lower frame rate and quality, the images can be combined to increase spatial and temporal

resolution (Shectman et al., 2002). In our laboratory, we record from 'target-detecting' neurons within the insect brain (Wiederman & O'Carroll 2013, Wiederman et al., 2013, O'Carroll & Wiederman 2014), which are likely neuronal correlates for such predator-prey interactions. Therefore, obtaining behavioural data with high spatiotemporal resolution places our electrophysiological results into a relevant context.

One readily available and comparatively affordable camera system is the Hero 3+ Black (GoPro, Inc.). These cameras were designed for capturing extreme sports from a first-person perspective. However, their usage has rapidly expanded into many other applications. They have the advantages of a high quality sensor capable of 720p resolution at 120 frames per second (fps). They are highly compact (59 x 41 x 21mm, with a mass of 28g), portable, waterproof and robust, yet with a built in battery and storage allowing continuous capture for many minutes at a time. We propose that a large array of dozens of such cameras could be harnessed on a lightweight support structure, deployed in the field. This system would capture complex animal behaviour within a large volume of space. Many of these cameras could be combined for less than the cost of a single high-speed unit.

Hero 3+ cameras are limited by a fixed aperture, fixed focus lens providing a large field of view and simple operation. Although there is no way to adjust focus and depth of field (DOF) in the stock configuration, the fixed aperture (f/2.8) and very short focal length (2.77mm) yield a large (170°) field of view and a DOF that covers ~150mm to infinity. For its original application as an 'action camera' this setup is perfect as it allows capture of the complete scene from a point of view style, but for capture of behavioural interactions the "fish eye" distortion and limited resolution provided by the wide-angle lens is undesirable. Aftermarket companies have identified these limitations, and have developed longer focal length aftermarket M12 thread lenses to replace the stock lens. This allows altering the FOV and effective removal of 'fish-eye' distortion, while also allowing adjustment to shift the plane of optimal focus and thus the DOF in the scene. DOF can be described as the distance between the nearest and farthest in-focus object of a scene. This area of focus can change within a scene and increase or decrease in size depending on the lens focal length, sensor pixel size, available light and aperture ratio of the lens. Longer focal length aftermarket lenses can improve the Hero 3+ working distance at which fine detail can be resolved by a given sensor, but at the expense of reduced DOF. This is a particular limitation when capturing behavioural data, particularly of unconstrained animals moving unpredictably in natural environments.

Features of interest can rapidly move either in front or behind the DOF, resulting in loss of spatial information in the out of focus imagery.

A solution widely used for still-frame systems is the technique of focal stacking images, an approach that combines multiple images at different focal depths to create an all-in-focus image (Antunes et al., 2005). The resultant enhanced spatial information can be achieved by several different stacking methods and the recent algorithm development emphasizes reduced computational complexity and increased accuracy of reconstruction (Guo et al., 2015; Zhong & Blum 2001; Liu et al., 2015; Pertuz 2013; Zhang & Ge 2009). These algorithms assume that the images are correctly registered (spatially aligned). This is less of a concern in still photography when a single camera from the same viewing angle takes multiple images at varying focal depths. Importantly, this imagery for focal stacking is inevitably captured at different points in time requiring a static (or near static) subject. To permit the simultaneous capture of behavioural data of moving subjects at varying DOFs, we propose a multiple-camera system with each individual camera set at a different focal distance.

Multiple camera systems permit simultaneous capturing of the scene with the individual cameras set to identical or different properties (e.g. FOV, frame rate, spatial resolution). However, camera systems with adjacently positioned units (minimizing intraocular distance between lenses) result in parallax and geometry errors that need to be corrected for in post-processing. In such setups, some errors cannot be accounted for, since the projection geometry can result in relevant information (e.g. feature markings) being obscured from other cameras. Although these adjacent cameras could provide additional spatial content, the focal stacking issues become increasingly complex.

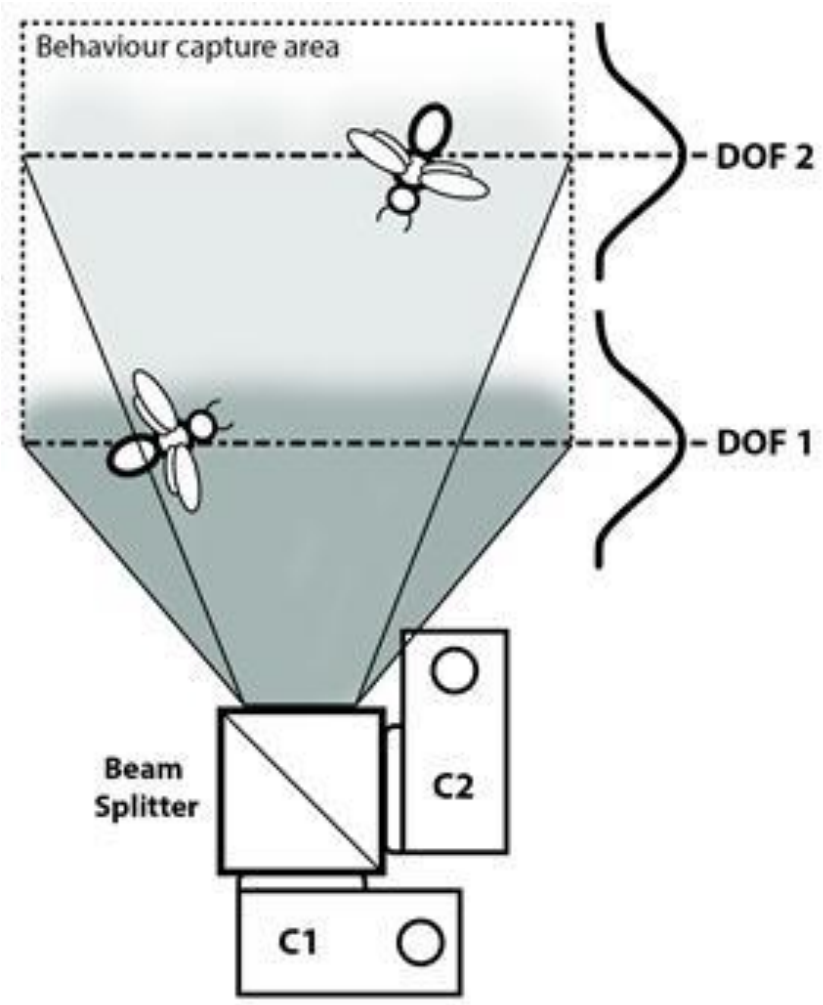

Figure 73: *Schematic of experimental setup. Cameras (C1 and C2) view the same image via a beam splitter capturing a behavioural volume of space represented by the dotted lines. The illustrative Gaussian curves are centred (dot-dashed lines) at DOF 1 (C1) and DOF 2 (C2). DOF 1 is 300 mm from camera system, with the FOV indicated by dark grey shading. DOF 2 is 1000 mm away from camera system with approximate FOV shown by light grey shading.*

To address these limitations, we present a novel technique for using Hero 3+ cameras and a beam splitter prism to obtain video footage that is captured coaxially for two different DOFs. We test this approach by fusing these data with previously developed image registration and focal stacking algorithms. For our testing, we used two co-axial cameras, however, due to the compact camera size additional beam splitters could be used to capture a larger depth range. In bright, natural, environments the reduction in light intensity due to the beam splitter does not impede the capture of behaviour over our desired volume of space.

10.4 Methodology

10.4.1 Camera characterization

We purchased off-the-shelf Hero 3+ Black cameras (GoPro Inc.). We removed the original wide-angle lens (170° angle) and replaced them with aftermarket lenses from RageCams (ragecams.com). The two versions of lenses used were the 5.4mm f/2.3 IR cut MP-10 flat lens and the Mega IR 12mm f/2.5Both lenses can have the focal distance set manually. We set the camera's resolution to 720p, 120 frame per second (FPS), Narrow FOV and low light adjustment off. We turned Protune on, with RAW white balance, flat colour space and a stock ISO value of 1600.

10.4.2 Image capture

We paired the cameras using the GoPro WiFi remote control, with synchronization to the nearest frame performed in post-processing. The cameras were attached to a gantry that incorporated a 30mm beam splitter (VIS coated, non-polarizing, 50/50) where they were carefully positioned to mitigate geometrical errors. Figure 73 is a schematic of the experimental set up, which is attached to our gantry for field deployment. To evaluate the beam splitter setup, we also applied the same techniques to another arrangement, with two cameras placed adjacently to one another (smallest intraocular distance). For testing stationary objects, we created foam cubes and filmed cardboard boxes. For dynamic scenes, we tested in the field, recording moving bees *(Apis mellifera*) in full sunlight. We imported each camera's footage into the proprietary GoPro studio software and manually aligned them to within a single frame, before saving short clips in the least compressed avi format (Cineform).

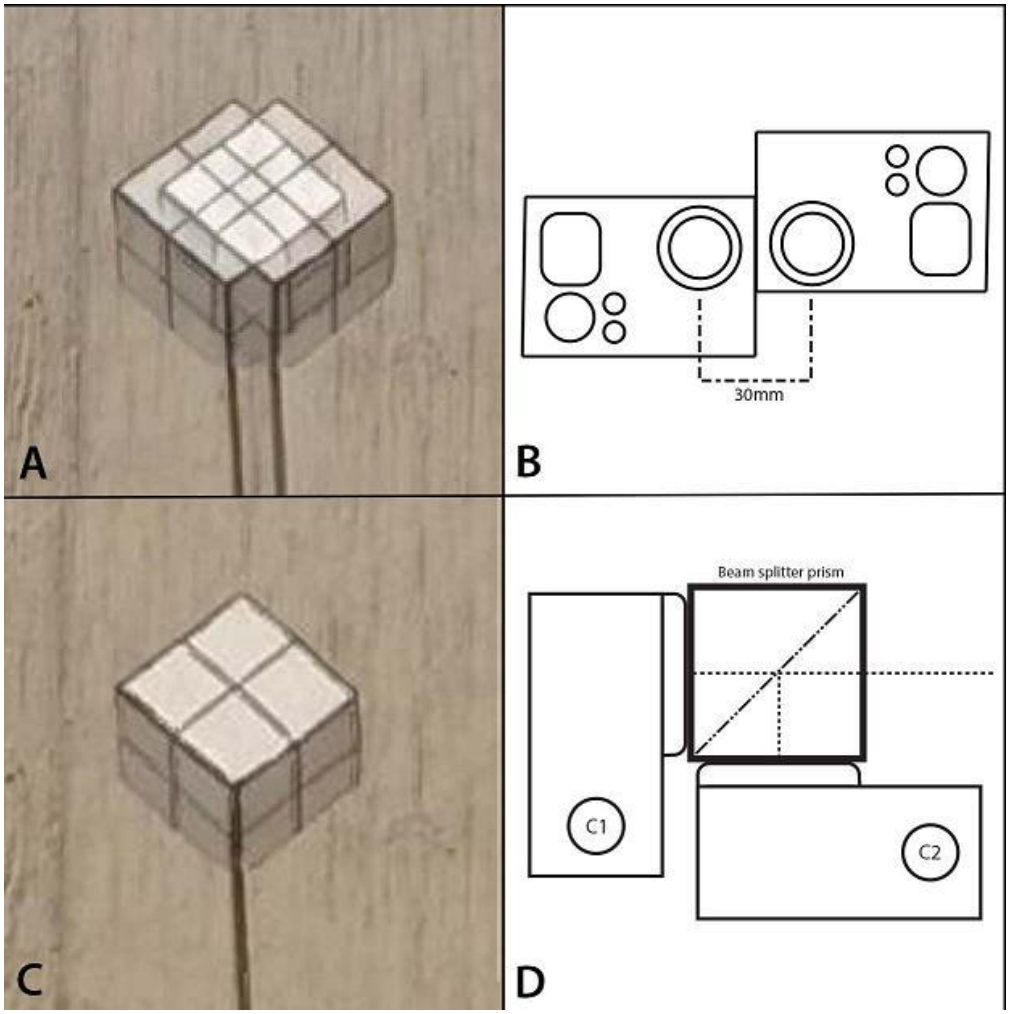

Figure 74: *Comparison between adjacent or beam splitter camera systems.* ***A)*** *the focal-stacked image from* ***B)*** *adjacent (frontal view) cameras.* ***C)*** *focal-stacked images from the* ***D)*** *beam splitter prism cameras (top view).*

10.4.3 Image processing

We read the avi files into MATLAB r2015a (Mathworks Inc.) using the Image Processing Toolbox to create individual image frames for each camera stream. We matched the image histograms for each camera at each time sample, to account for slight variations in image contrast and brightness. The images were then focal-stacked in Zerene Stacker software (Zerene Systems LLC) using the Pmax algorithm, a Pyramid method designed to work well with overlapping structures. This technique preserves detail but results in increased contrast and noise. The resulting stacked images were combined back into a single output video stream using MATLAB.

10.5 Results

In order to examine the effectiveness of our technique, we examined individual frames from our focal-stacked video sequence (Figure 74). We clearly see the improvement over the

adjacent camera system (Figure 73A, B) by the beam splitter prism setup (Figure 73C, D). The adjacent camera stack results in severe alignment errors following stacking due to the attempted fusion between the different projections. Furthermore, the beam splitter technique produces images that require less image registration (if any) than the adjacent camera system.

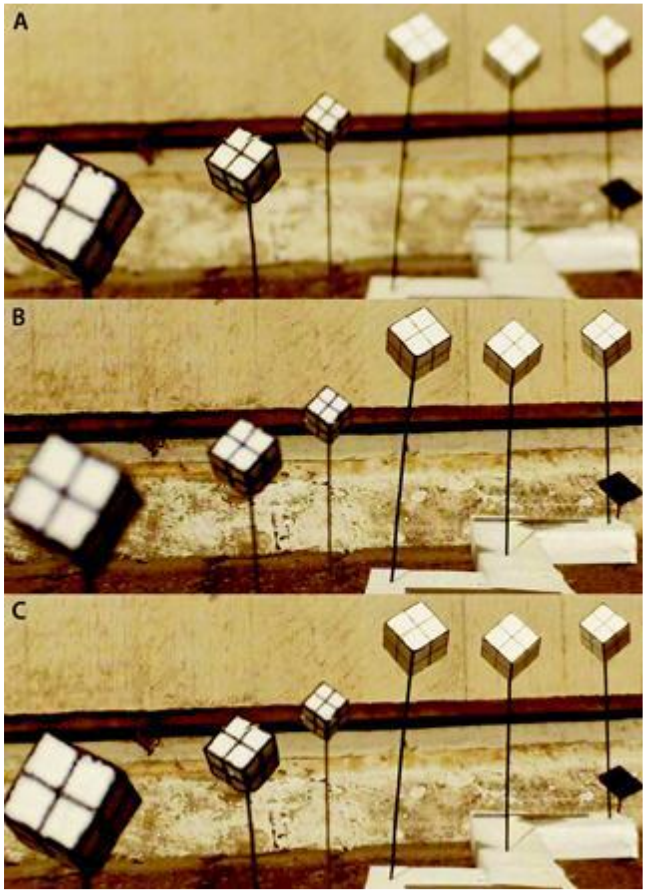

Figure 75: ***A)*** *Image from camera 1 (C1) with foreground in focus (Depth of field set at 300mm from camera).* ***B)*** *Image from camera 2 (C2) with background in focus (DOF set at 1000mm from camera).* ***C)*** *focal stacked image using Zerene stacker using Pmax method.*

The practicality of the beam splitter capture system when partnered with the focal stacking technique is evident in Figure 75. We positioned static foam targets at 200mm, 300mm, 400mm, 900mm, 1000mm and 1100mm distances. These distances are typical of the working range over which this camera/lens combination is capable of providing adequate spatial resolution for filming behaviour of larger flying insects such as bees or dragonflies. Optimal focal distances were set at 300mm for C1 and 1000mm for C2. The images are near DOF (Figure 75A), far DOF (Figure 75B) and the focal-stacked merger of images from C1 and C2 (Figure 75C). The technique successfully brings into focus our 6 features of interest, however, intriguingly the stacking algorithm has lost some of the background detail (Figure

75B cf. Figure 75C). This is a result of using the propriety focal stacking algorithm and we are currently developing alternative versions.

An alternative application of the beam splitter prism is to combine imagery taken from cameras with different lenses. In this way, we can capture information from the same scene and viewing axis for two alternative trade-offs between the DOF / FOV and the resolution within the scene. For example, one camera (5.4mm lens) can capture overall predator-prey interactions contained within the wider FOV. This video (after accounting for camera asynchrony) could be used to track feature position using interpolation techniques. Simultaneously, a second camera (12 mm lens) captures head and body movements of the predator at a finer spatial resolution, albeit within a more limited volume of the scene. Figure 76 demonstrates the practicality of using the different lens combination with the beam splitter prism system. The first image (Figure 76A) is taken with the wider angle 5.4 mm lens (focused at 300mm) with the background target out of focus. The second image (Figure 76B) is from a camera with a 12mm lens, which has a smaller FOV but captures more spatial resolution at a distance. The fine detail of the box at 1000 mm from the camera captured by the 12mm lens, however, the 5.4mm image has a larger FOV, capturing a larger extent of the scene.

Figure 76: *Comparison of images captured simultaneously through the beam splitter prism with* ***A)*** *5.4mm lens and* ***B)*** *12mm lens.*

We are currently testing this system in real-world environments with moving features in the scene. Due to the temporal asynchrony of the cameras (up to 4.16 ms, 0.5 frames at 120fps), objects within the scene may still have moved significantly, irrespective of the use of the beam splitter prism. These effects are evident in Figure 77. This shows the result of stacking a moving target in a natural scene (high clutter) and an artificial scene (low clutter). The images to be stacked in the natural scene (Figure 77A, C) show finer resolution in the foreground and background respectively. The resulting stacked image is displayed in Figure 77E, revealing minimal distortion of the target (bee in centre frame). In the artificial scenes, the tethered insect is in focus in Figure 77B, however, is out of focus in Figure 77D. The tethered insect in the stacked image (Figure 77F) is therefore duplicated. The 'ghosting' effect is due to the asynchrony between each camera; here the stacking algorithm deciding on a frame-by-frame basis which part of the image represents the best spatial frequency content. This is beneficial for our current work as high spatial information is required to reconstruct feature information in biological interactions (e.g. head movements, body orientation, identification and categorization of target). Even without focal stacking, our combination of two DOFs increases the probability of capturing an in-focus version of the target.

It should be noted that the reconstructed videos are not designed to be perceptually appealing, with such 'artefacts' (extra data) in the reconstructed video. However because this focal stacking technique merges the two points in time for the captured frames into a single output image, this also makes analysis of temporal positioning and trajectories difficult. Instead this can be calculated using the original camera streams.

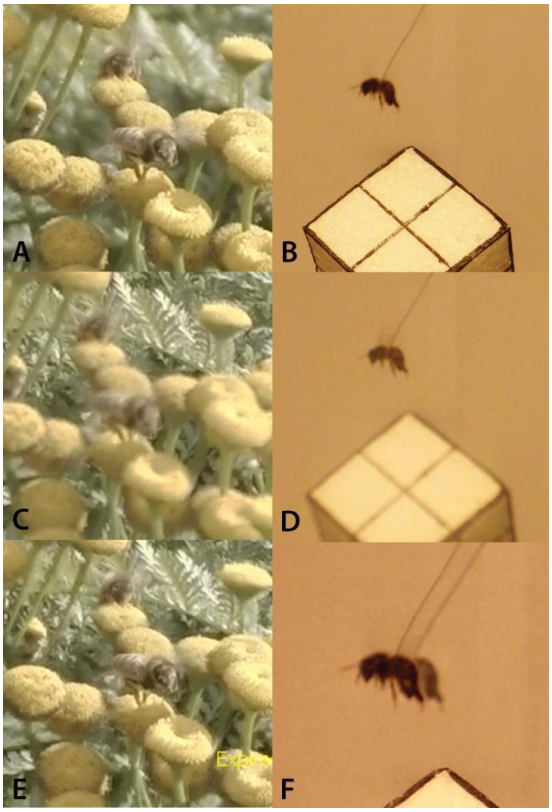

Figure 77: *Comparison of focal stacked moving insect targets in a natural scene and artificial scene.* ***A*** *and* ***B****, camera 1 (C1) image of natural and artificial scene.* ***C*** *and* ***D****, Camera 2 (C2) image for natural and artificial scene.* ***E*** *and* ***F****, produced stack showing alignment errors produced by lack of data in background.*

In principle, these problems can be mitigated using the original individual camera streams. Moreover, we have developed a post hoc estimate of asynchrony (Evans et al., 2015) between the multiple cameras that allows us to use this additional information for reconstructing feature trajectories with better temporal resolution than the default frame rate (e.g. knowing the precise time points of the in-focus and out of focus bee before image fusion).

Another limitation of our beam splitter system is the availability of light for the cameras. Each beam splitter prism approximately halves the light received by each sensor. In many consumer level camera systems, lens aperture can be increased to mitigate effects of low light, which results in a brighter image but a shallower DOF. Because of the fixed aperture of the lenses we used (either f/2.3 or f/2.5), the camera compensates for lower levels of light by using a slower shutter speed to allow more light to reach the sensor. This can result in additional temporal blurring in the image and, as a result, a loss of detail for fast moving targets (e.g. insects in flight). In ideal conditions (full sunlight) this effect is minimal. But the reduced light capture could be a more serious problem in restricted lighting conditions or for studying animals active at low luminance.

Despite these limitations, our beam splitter system provides a larger volume over which in-focus information can be captured when compared to a single camera. In focus-stack mode this is still limited, however, to a single FOV. Methods in the literature have used combination of single sensor cameras to create an array to capture increased spatial and temporal information (Wilburn et al., 2004; Wilburn et al., 2005). The incorporation of the beam splitter system into an array configuration would be an ideal way to increase the total captured volume. In ideal conditions a 'multiple' beam splitter system could be implemented. This would divide each output of the original beam splitter into two paths into which additional beam splitters could be used to gather data form several cameras. The Hero 3+ cameras are sufficiently compact for this to be practical. This piggybacking of beam splitters could greatly increase the amount of data that can be collected from the same FOV or even multiple FOVs if different focal lengths were used, although this would reduce the light gathered by each camera sensor. However, for field work on day-active insects, this limitation is frequently mitigated by the abundance of available light. This application is thus practical when considering the inexpensive setup of these cameras and will be assessed in our future work with this system.

10.6 Conclusion

Our novel method of combining cost effective commercial level cameras with a beam splitter prism to capture footage suitable for focal stacking has useful applications in capturing behavioural information for field research. Modification to the currently proposed camera system with the addition of lighting to mitigate any low light information loss, and a robust synchronization technique to control camera synchronization will increase the usability of the system. With the advancement in available camera technology, it won't be long before better camera systems are available commercially. Combining such advances with our approach should make it practical to capture unprecedented amounts of information about the natural behaviour of animals in the field, an endeavour that would not have been possible with earlier generation equipment.

10.7 References

1) Antunes M, Trachtenberg M, Thomas G & Shoa T, 2005, *All-in-focus imaging using a series of images on different focal planes*, in Image Analysis and Recognition, M. Kamel and A. Campilho, Editors., Springer-Verlag Berlin: Berlin. *p. 174-181*.

2) Evans B, Parslow B, O'Carroll DC and Wiederman SD, 2015, *Quantifying Asynchrony of Multiple Cameras using Aliased Optical Devices*, IPTA 2015

3) Guo D, Yan J & X Qu, 2015, *High quality multi-focus image fusion using self-similarity and depth information*. Optics Communications (vol. 338, pp. 138-144)

4) Liu Y, Liu S & Wang Z, 2015, *A general framework for image fusion based on multi-scale transform and sparse representation*. Information Fusion (vol. 24, no. 147-164)

5) Mischiati M, Lin HT, Herold P, Imler E, Olberg R & Leonardo A, 2015, *Internal models direct dragonfly interception steering*, Nature (vol. 517, pp. 333-338), https://doi.org10.1038/nature14045

6) O'Carroll DC & Wiederman SD, 2014, *Contrast sensitivity and the detection of moving patterns and features*, Philosophical Transactions of The Royal Society B (369: 20130043) https://doi.org/10.1098/rstb.2013.0043

7) Pertuz S, 2013, *Generation of All-in-Focus Images by Noise-Robust Selective Fusion of Limited Depth-of-Field Images*. IEEE Transactions on Image Processing (vol. 22, no. 3, pp. 1242-1251)

8) Shechtman E, Caspi Y, and Irani M, 2002, *Increasing Space-Time Resolution in Video*, in Computer Vision — ECCV 2002, A. Heyden, et al., Editors., Springer Berlin Heidelberg. p. 753-768

9) Wiederman SD, Shoemaker PA & O'Carroll DC, 2013, *Correlation between OFF and ON Channels Underlies Dark Target Selectivity in an Insect Visual Systems*, The Journal of Neuroscience (vol. 33, no. 32, pp. 13225-13232) https://10.1523/JNEUROSCI.1277-13.2013

10) Wiederman SD & O'Carroll DC, 2013, *Selective Attention in an Insect Visual Neuron*, Current Biology (vol. 23, pp. 156-161), http://dx.doi.org/10.1016/j.cub.2012.11.048

11) Wilburn B, Joshi N, Vaish V, Talvala E, Antunez E, Barth A, Adams A, Horowitz M & Levoy M, 2005, *High Performance Imaging Using Large Camera Arrays*, Proceedings of SIGGRAPH (765-776).

12) Wilburn B, Joshi N, Vaish V, Levoy M & Horowitz M, 2004, *High-speed videography using a dense camera array*. In *Computer Vision and Pattern Recognition* Proceedings of the 2004 IEEE Computer Society Conference

13) Zhang YJ & Ge LL, 2009, *Efficient fusion scheme for multi-focus images by using blurring measure*. Digital Signal Processing (vol. 19, no. 2, pp. 186-193)

14) Zhong Z & Blum R, 2001, *A hybrid image registration technique for a digital camera image fusion application*. Information Fusion (vol. 2, no. 2, pp. 135-149)