



THE UNIVERSITY *of* EDINBURGH

This thesis has been submitted in fulfilment of the requirements for a postgraduate degree (e. g. PhD, MPhil, DClinPsychol) at the University of Edinburgh. Please note the following terms and conditions of use:

- This work is protected by copyright and other intellectual property rights, which are retained by the thesis author, unless otherwise stated.
- A copy can be downloaded for personal non-commercial research or study, without prior permission or charge.
- This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author.
- The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author.
- When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given.

A biorobotic investigation of multimodal cue integration for insect orientation

Robert Mitchell

Doctor of Philosophy
The University of Edinburgh
2023

Abstract

Insects are capable of great feats of navigation; these range from relatively simple straight-line orientation over a few metres to migration over hundreds of kilometres. A robust compass is therefore critical, and to this end, an animal may use multiple orientation cues simultaneously (a behaviour known as *cue integration*). Recent work has indicated that ball rolling dung beetles perform cue integration during straight-line orientation. This thesis explores dung beetle cue integration using a biorobotic approach, with the goal of understanding how this process could be implemented in the insect brain.

We start by adapting the standard Bayesian approach from human psychophysics to compare mathematical models of cue integration behaviour, assuming that dung beetles weight cues according to their *reliability* (the inverse of their variance). This modelling indicates that beetle cue integration is likely represented by a vector sum which: (1) aligns well with a candidate neural substrate, and (2) can generate a variety of behaviours depending on how the cue weights are set (providing behavioural flexibility despite functionally conserved neuroanatomy).

It has previously been proposed that orientation cues are integrated in a set of neurons known as EPGs which receive input from 'R' neurons via highly plastic connections. Recent evidence has shown that different sets of R neurons encode different cue modalities. Using a computational model, we argue that R neurons could encode vectors as sinusoidal activity across their population and the resultant input to the EPGs represents a vector sum. The model can account for the vector sum, and provides a possible mechanistic explanation for the dung beetle orientation snapshot. Further, the work suggests that a number of factors may contribute to total cue influence; specifically, cue weight, cue reliability, and synaptic strength (which is affected by weight and reliability independently). Moreover, we show that reliability does not seem to be the major factor which determines cue influence in dung beetles (using a combination of modelling and animal behaviour). This result stands in contrast to the vast majority of literature on cue integration, including all previous accounts in insects which discuss concrete models, and prompts a critical re-examination of the explanatory power of the Bayesian approach to cue integration.

Lay summary

Despite possessing tiny brains (around the size of a sesame seed), insects demonstrate complex behaviours and this complexity is particularly apparent in navigation. Famously, bees are able to fly straight back to their nests after foraging and communicate food locations to their nest mates. In order to do this, the bee needs a good sense of direction; that is, it needs an internal compass. In many insects, the compass is difficult to study behaviourally as it is one component in a more complicated behaviour. Of interest then, is the lowly ball-rolling dung beetle.

True to their name, ball-rolling beetles will form dung into a ball and roll it away from a dung pat. Curiously, they roll their balls along stubbornly straight paths in apparently arbitrary directions. If presented with an obstacle they will shuffle sideways to work their way around, then continue undeterred along their chosen bearing. This ability suggests the beetles are using some kind of external compass cue in much the same way a human would use the needle on a magnetic compass. Indeed, the beetles are known to be able to use a variety of different orientation cues to guide their behaviour. It has also been shown that beetles learn the spatial relationship of different cues so that they can continue their path, even if one cue is removed. Most recently, it has been demonstrated that multiple compass cues can affect behaviour simultaneously, a phenomenon known as *cue integration*.

Cue integration has been observed in mammals (particularly humans) for some time. Over the last twenty-five years there has been a concerted effort to describe cue integration behaviour using ideal observer models. These models use probabilistic inference to provide mathematical rules for generating the best possible estimate of the world given some potentially erroneous sensory inputs. While such models can be useful, they mostly remain agnostic as to neural implementation. As such, the process of neural cue integration is very much a black box in the mammalian brain. Demonstrating such behaviour in dung beetles therefore presents a unique opportunity to study neural computation. Insect brains are relatively accessible and the relevant parts are quite highly conserved across species.

This thesis provides an exploration of compass cue integration in the insect brain. Taking advantage of the relative simplicity of beetle behaviour, behavioural modelling is used to show that dung beetle cue integration is likely implemented as a vector sum. This computation supports a wide range of behaviours depending on how it is configured, which means that different insect brains could implement the same core computation but generate different behaviours. Then, taking advantage of the accessibility of the insect brain, a neural model is developed to

demonstrate how the vector sum computation could plausibly be implemented therein. This model has the capacity to learn about the relative positions of different orientation cues and this is put to the test on a robot, operating under similar experimental conditions to the beetles. The neural model indicates that insect cue integration does not obviously correspond to the predictions of ideal observer models and therefore, the thesis also provides a critical analysis of the utility of the ideal observer approach in insects.

Acknowledgements

I consider any development of my own to be a product of the interactions I have had with others and throughout my project, I have had the good fortune of a good community. If you have been part of this community, you have undoubtedly had a positive influence on my experience and I will be forever grateful. Nevertheless, I would like to name some names.

I would first like to thank my principal supervisor, Barbara Webb. Thank you for showing me the wonderful worlds of insect navigation, neural computation, and biorobotics. Thank you for introducing me to Braitenberg's *Vehicles*, which has been foundational to my views on animal cognition. Thank you for teaching me to be critical of the work that I read and the work that I produce, and for encouraging me to defend my position, even where you disagreed. Most of all, thank you for your time. You give a substantial amount in regular (and irregular) meetings and in providing critical feedback. I greatly appreciate the time you were willing to give me as a mentor.

I would also like to thank my assistant supervisor, Marie Dacke. Thank you for adopting me into your lab group. The opportunity for regular discussion with a community of biologists and exposure to animal behavioural work gave me a much-needed dose of perspective, particularly at the beginning of the project. Thank you for having me out to Lund to visit, and in particular for having me in the field. My experience in South Africa was magical, and I will carry it with me for life. Thank you for 'being the rubber duck' and letting me talk at you while I tried to figure out robot problems. Your enthusiasm and positivity helped keep me calm during the more stressful periods.

In addition, I would like to thank the Dacke Lab and the Lund Vision Group. Thank you all for being so welcoming and friendly, and for many interesting and entertaining conversations. In particular I would like to thank Elin Dirlik, Sherry (Shahrazad) Shaverdian, and Lana Khaldy. Elin and Sherry, thank you for asking me to be part of our first paper, for your patience during many frustrating discussions about reliability, and for smiling when asked 'Could you just run this extra experiment?'. Thank you both for being good friends throughout. Lana, thank you for helping me get over my social roadblocks when first visiting the group. Finding a voice as cynical as my own made me feel right at home.

I would also like to thank the InsectRobotics group. Thank you all for helping me with the concepts I struggled to understand and for giving me a space to present my own ideas and get feedback. In particular, I would like to thank James Garforth and Evripidis Gkaniias. James, you were one of the first voices to show me that it was okay to criticise work that was published. You have often made me think twice about things

which I thought obvious, for which I am grateful. Evri, you have improved my graphical communication skills beyond measure. You have always been ready to help and discuss ideas and in many ways, you have acted as a model for how to approach a PhD. In addition, I'd like to thank denizens of G.17 past and present: Jan Stankiewicz, Aruna Raman, Rana El Khouri Maroun, and Le Zhu for many interesting discussions, personal and professional. Particularly the (heated) arguments about the nature of insect cognition.

My thanks also go to my proofreaders: Auguste de Pennart, Anna Hadjitofi, Florent le Moël, Yakir Gagnon, Cora Hummert, Evri, Sherry, Marie, and Barbara. All of you provided valuable feedback which helped to shape this work.

I would like to close by thanking my family. To my Mither and Faither, there is no way I could ever have made it this far without your love and support over the years. Thank you for the opportunities you've provided me, and thank you for letting me choose which opportunities to take. The value in this freedom is something I have come to truly appreciate in recent years. Thank you Lucie for dealing with a rather rubbish brother over the last few years. I hope to be more present both for you and wee Fraser in future. Thank you all for giving me space to work and for keeping up contact in my more distant periods; I appreciate that you didn't give up on me. Finally, thank you Alana for your love and patience throughout. Thank you for being a sounding post when I needed to ramble through some thoughts. Thank you for proofreading my various scribblings throughout. Thank you for being there.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

(Robert Mitchell)

'The world of acronyms is vast and pitiless.' - Alana L. White

Contents

1	Introduction	1
1.1	The insect compass as a target for modelling	1
1.2	From computational neuroethology to biorobotics	2
1.3	The problem: cue integration in the insect compass	3
1.4	Summary of contributions	5
1.5	Thesis structure	5
2	Dung beetle straight-line orientation	7
2.1	Behavioural overview	8
2.2	Available orientation cues	9
2.3	Sensory system	16
2.4	The central complex	18
2.5	The dung beetle dance and the orientation snapshot	22
2.6	Summary	24
3	Cue integration	25
3.1	Insect accounts	25
3.1.1	Descriptive accounts	26
3.1.2	Algorithmic accounts	28
3.1.3	Computational accounts	29
3.1.4	Integrating angular velocity and external cues	31
3.2	The Bayesian perspective	33
3.2.1	Why Bayes?	33
3.2.2	Bayes' theorem in the context of cue integration	33
3.2.3	Removing abstraction	34
3.3	Discussion	35
3.3.1	Linear or Bayesian?	37
3.3.2	Bayesian behaviour	37
3.3.3	Answering 'Why?' questions	40
3.3.4	How does this relate to insects?	40
3.4	Summary	42

4	Paper 1: Weighted cue integration for straight-line orientation	43
5	Paper 2: A model of cue integration as vector summation in the insect brain	64
6	BeetleBot	85
6.1	System description	85
6.1.1	Custom hardware	86
6.1.2	Software	91
6.2	Proposed extensions	93
6.3	Summary	95
7	Demonstrating the orientation snapshot using the BeetleBot	96
7.1	Methods	97
7.1.1	Robot behaviour	97
7.1.2	Multimodal compass model	100
7.1.3	Neural steering circuit	102
7.2	Results	105
7.2.1	Modality transfer	105
7.3	Discussion	111
7.3.1	Modality transfer	111
7.3.2	A possible generalisation of insect steering	112
7.3.3	Repeating experiments <i>in silico</i>	118
7.4	Summary	119
7.5	Acknowledgements	119
8	Conclusions	120
8.1	Cue integration as a vector sum	120
8.2	Reliability does not determine cue influence	121
8.3	Testing models <i>in silico</i> provides benefits beyond simulation	122
8.4	The orientation snapshot requires a firm definition	122
8.5	General rules for insect steering circuits	124
8.6	Bayesian cue integration models are of limited utility in insects	124
	Epilogue	126
A	Polarised light detection	140
B	Derivation of the linear rule for ‘optimal’ cue integration	143
C	Supplementary information for Paper 1	147
D	Supplementary information for Paper 2	157

E	BeetleBot start-up guide	185
E.1	Software availability and usability	185
E.2	Practical considerations for usage	185
E.3	Start-up guide	186
F	Additional paper: Celestial compass design mimics the fan-like polarisation filter array of insect eyes	192

Chapter 1

Introduction

1.1 The insect compass as a target for modelling

Insect navigation provides many examples of compact brains performing complex behaviour, which opens a unique window into neural computation. Monarch butterflies and bogong moths migrate over hundreds of kilometres ([Reppert and de Roode, 2018](#); [Warrant et al., 2016](#)), ants run directly back to their nests after tortuous outbound paths ([Wehner and Wehner, 1990](#)), and dung beetles roll dung balls in straight lines ([Byrne et al., 2003](#)). These behaviours all require a central component: a robust compass ([Cheung et al., 2007](#)).

While it has long been known that insects possess an internal compass, the compass can be difficult to isolate behaviourally. Ball-rolling dung beetles exhibit a repeatable and relatively controllable straight-line orientation behaviour which appears to rely solely on the compass ([Dacke et al., 2021](#)). They are known to be able to orientate to a variety of different cue modalities ([Dacke et al., 2021](#)) and in addition, they have been shown to be able to learn arbitrary arrangements of cues ([el Jundi et al., 2016](#)). Most recently, they have been shown to integrate multiple cues simultaneously to guide their behaviour ([Dacke et al., 2019](#); [Khaldy et al., 2021, 2022](#)). Dung beetle straight-line orientation provides a behavioural window into a robust, flexible orientation system, operating in a brain roughly the size of a sesame seed.

It is therefore of interest to understand how such an apparently powerful system can be implemented in such limited hardware. Insects are a rare example of complex behaviour and in a brain we can practically study. Useful quantities neuroanatomical data can be collected by manual tracing of individual neurons ([el Jundi et al., 2018](#); [Sayre et al., 2021](#)) (even if this is very laborious). Recent breakthroughs in automatic neuron tracing have resulted in a functional ‘connectome’ of a significant chunk of the fruit fly brain ([Scheffer et al., 2020](#)), including the region

which houses the compass, known as the central complex (Hulse et al., 2021). This connectome gives us a reasonable map of what connects where, and with roughly what strength. In addition, neurophysiological data is also relatively accessible. Classic intra- and extracellular recording techniques are available (Nguyen et al., 2021; el Jundi et al., 2015b; Stone et al., 2017; Beetz et al., 2023) and, in fruit flies specifically, calcium imaging permits researchers to examine how specific neuron populations respond to stimuli (Seelig et al., 2010). Further, in fruit flies, optogenetic manipulation provides a way to causally link neuron activity to behaviour (Deisseroth, 2011). Notably, the insect central complex (the putative navigational centre) is well conserved across different insects (Honkanen et al., 2019). We can (with care) use fruit fly neurobiology to provide models for dung beetle behaviour.

Altogether, insect navigation presents a remarkable opportunity to place behaviour directly into a neural context. To take it a step further, we can provide a concrete computational model of the brain and test it against behavioural data to see if it truly does capture the target behaviour. This approach is generally referred to as *computational neuroethology*: loosely, using computer modelling to tie brain to behaviour.

1.2 From computational neuroethology to biorobotics

Computational neuroethology is interesting as an approach from two perspectives.

1. Brains tend to be naturally adept at tasks for which it is difficult to program a computer (Moravec's paradox).
2. Building concrete computational models of (bits of) brains usually exposes gaps in our understanding of the target neural circuit or behaviour.

To expand further on the second point, models of animal behaviour have a tendency of being 'verbal' which can make them overly accommodating. Producing a *concrete* computational model highlights which assumptions need to be made in order to make the model work as proposed. Thus, we can make sure our models are well-defined and make solid behavioural predictions.

That said, how do we determine what behaviour a model predicts (or generates)? A common answer (at least as pertains to models of insect navigation) is to place the model in a simulation environment and see if the model captures the target behaviour. This is a valid approach but

it is worth noting that simulations inevitably end up being artificially sterile; we ‘include only what *we think* is important’ (Webb, 2000) (emphasis added). One can argue that a simulation can always be improved, more powerful computing hardware should lead to more accurate simulations. One then glances at current robotics literature and sees that the ‘sim-to-real’ gap persists. Alternatively, we can take our model out of simulation and test it in reality by using a robot, termed a *biorobot* (Webb and Consi, 2001). This has added appeal in that we can use the same behavioural assay that was used to test the animal in the first place.

Even with an imperfect simulation, one may argue that a biorobot does not provide additional insight. After all, a robot approximates reality via its sensors and the chosen hardware may not mimic the target animal well. We can counter by observing that hardware choices make assumptions (and related limitations) about the sensory system explicit. The requirement to specify each step from input through perception to action to generate behaviour forces us to critically assess that specification (Webb, 2000). Furthermore, the possibility of testing the hardware in the same behavioural setup used for the animal eliminates any gaps between the simulated approximation and the real behavioural experiment. The biorobot will experience the same problems that the animal did. Thus, if we want to understand insect navigation and the insect compass, a biorobot provides a useful tool with which to do so.

To be clear, simulations are by no means useless and provide a valuable tool for model construction and basic functionality testing. The biorobotic implementation simply provides an additional, powerful test for biological plausibility. In much the same way that a concrete computational model highlights gaps in our understanding of the biological system, the biorobot can highlight problems with the model which may have been hidden by the simulation. Thus, a complete computational approach to understanding the insect compass might first perform a behavioural analysis, then provide a simulated computational model, and finally trial this model on a robot.

1.3 The problem: cue integration in the insect compass

Having justified the approach, we return to the robust, multimodal insect compass. Compass cue integration behaviour in dung beetles has received seemingly conflicting reports. It has been described as ‘winner-take-all’ (Dacke et al., 2019) (the most reliable cue controls behaviour) and has also been described as ‘Bayesian’ (Khaldy et al., 2022) (in this context, cues are smoothly weighted according to their reliability). In

cue integration literature, these are generally considered as separate outcomes (cue capture vs. cue integration), even if the line is fuzzy (Ernst and Bühlhoff, 2004).

Behaviourally disambiguating these behavioural models requires quantifying single cue reliability and determining whether or not relative cue influence corresponds with its relative reliability. However, this form of analysis had not been performed.

Behaviour is one avenue of attack, but remember, we also have access to the brain. A neural compass has been identified in the insect brain. A set of neurons in the ellipsoid body of the central complex, known as EPGs, represent animal direction as a sinusoidal bump of activity (Seelig and Jayaraman, 2015). This peak will shift around the ellipsoid body in lock-step with head direction; in flies the ellipsoid body is toroidal and the bump behaves like a literal compass needle tethered to the environment. In other insects, the correspondence is not so neat but the functionality appears to be identical (e.g. in dung beetles (el Jundi et al., 2015b)). EPG neurons receive visual cue input via a class of neurons known as R neurons (or ER neurons in more recent descriptions), and the connections between R and EPG neurons are plastic (Fisher et al., 2019; Kim et al., 2019). The compass can also update using idiothetic angular velocity input from PEN neurons, though it is not currently known if this input comes from proprioception or motor efference (Turner-Evans et al., 2017; Fisher, 2022). It is now known that angular velocity regulates plasticity between R and EPG neurons (Fisher et al., 2022).

Behaviourally, dung beetles are known to learn relative cue positions and it has been proposed that relative cue position could be stored in plastic synapses between compass inputs and compass neurons (el Jundi et al., 2016). This idea is supported by recent evidence that wind and skylight polarisation inputs can update the compass via separate sets of R neurons (Okubo et al., 2020; Hardcastle et al., 2021) and has been carried forward in the literature (Hulse et al., 2021; Fisher, 2022; Wilson, 2023). However, no concrete computational model has been provided to demonstrate the principle. In addition, it has been proposed that this learning mechanism could facilitate reliability-based (i.e. 'Bayesian') weighting of cues, but again, this has not been tested.

This thesis addresses these literary gaps through the use of computational models. Bayesian modelling of dung beetle cue integration behaviour indicates that they integrate cues according to vector summation and this computation could be implemented by the putative multimodal compass circuit. I then provide a neural model to demonstrate that the circuitry can account for a vector sum computation and can also account for the dung beetle orientation snapshot (the ability to learn the offsets between different cues). The model's ability to implement the snapshot

is further tested on a custom robot platform operating under comparable conditions to the beetles.

1.4 Summary of contributions

1. Dung beetle compass cue integration behaviour is best explained by a vector sum calculation. Cue influence does not seem to be well predicted by measured single cue reliability suggesting that the beetles are not 'Bayesian' actors as has previously been suggested ([Paper I](#)).
2. A biologically plausible model of the insect head direction circuit can encode a vector sum. The neuroarchitecture does not couple cue influence and cue reliability as has been suggested. Instead, cue influence is largely determined by cue weight which, in our model, is a property of the input. This model observation is bolstered by behavioural evidence cue influence is not determined by reliability in dung beetles ([Paper II](#)).
3. The model also accounts for the dung beetle orientation snapshot which is demonstrated on a robot platform operating under similar behavioural conditions to the beetles (Chapter 7).
4. The conceptual separation of cue influence and cue reliability provides an opportunity to critically examine the utility of the Bayesian approach to cue integration and determine whether or not it is useful in insects. The Bayesian approach appears to provide little insight in insects due to our ready access to relevant neuroanatomy (Chapter 3).
5. The currently proposed insect steering mechanism ([Mussells Pires et al., 2024](#)) is used to determine a set of general rules which should govern steering circuit construction. These rules permit a wide variety of possible circuitry and future work could explore why specific examples arise in nature.

1.5 Thesis structure

Chapter 2 provides the necessary dung beetle behavioural background. Chapter 3 provides a summary of previous cue integration experiments in insects, the Bayesian framework which is commonly used to study this problem, and a critical examination of the utility of this framework

in insects. Chapter 4 is a collaborative study in which I provided modelling for a dung beetle behavioural experiment to place the results into a computational context ([Paper I](#)). Chapter 5 presents the neural model with simulated results and an accompanying dung beetle behavioural experiment ([Paper II](#)). Chapter 6 describes the robot platform developed for this project as a testing platform for insect navigation models. Chapter 7 then gives the example biorobotic experiment. Chapter 8 places the contributions of the thesis into the context of the literature and suggests future directions.

Chapter 2

Dung beetle straight-line orientation

Dung beetles are commonly split into three loose categories: dwellers, tunnelers, and rollers. Dwellers will live on or in a dung source, tunnelers construct a tunnel beneath the source to sequester their food away from competition, and rollers will form a piece of dung into a ball and transport it away from the dung pile before burying it somewhere safe ([Hanski and Cambefort, 1991](#)). It is the rollers that are of interest in this thesis and we will generally ignore the others.

The earliest known account of dung beetle ball rolling behaviour is given by Plutarch ([Babbitt, 1936](#), pg. 29), however, the later (substantively similar) description given by Horapollo is arguably more poetic (at least as translated into English):

“Rolling it with its hind legs from east to west, it faces the east, so as to give it the shape of the world, for the world is borne from the east to the west.”

[Boas \(1993, pg. 48\)](#)

As we will see, poeticism does not guarantee correctness. While there are later accounts of ball rolling behaviour in the scientific literature (e.g. [Matthews, 1963](#)), the body of literature of interest to us (with a specific focus on straight-line orientation) begins with [Byrne et al. \(2003\)](#) and continues up to the present. Herein, we will be focussing on a collection of species native to South Africa (see Table 2.1). It is worth noting that not all beetles now classed as ‘rollers’ physically roll their balls, some will drag ([Dacke et al., 2020](#); [Khaldy et al., 2019b](#)); that said almost all beetles discussed here do roll balls. The important distinction is that they convey dung away from the dung pile before consuming it.




Figure 2.1: *Kheper lamarcki* rolling its dung ball. Ball rolling beetles are able to use a variety of orientation cues to roll their balls along straight paths. Beetle image provided by the Lund Vision Group.

2.1 Behavioural overview

Guided by olfaction, dung beetles fly in search of sustenance. Once fresh dung has been found, the beetles will land nearby and walk the remaining distance. At the dung pat, a ball roller will cut off a piece of dung, shape it into a ball, perform a brief orientation dance (a rotation atop the dung ball) (Baird et al., 2012; el Jundi et al., 2016), and roll the ball away in a (roughly) straight line (Dacke et al., 2021). The beetles do not roll from east to west, instead rolling their balls in a seemingly random direction with respect to their surroundings (Baird et al., 2010; Khaldy et al., 2019b). This behaviour, known as menotaxis (travelling in an arbitrary direction with respect to some external cue), is thought to be a strategy for avoiding competition at the dung pile (Byrne et al., 2003); after all, why make your own ball when you can steal one, ready-made? Dung is a valuable resource for feeding and breeding and the competition is ruthless (Figure 2.2).

It is therefore unsurprising that straight-line orientation is remarkably robust in dung beetles. Once on their bearing, beetles can recover after falling from a ramp, being passively rotated, or being forced around an obstacle (Byrne et al., 2003; Dacke et al., 2004, 2003b; Baird et al., 2012). If the beetles were using idiothetic information exclusively, it seems un-



Figure 2.2: A dung thief in action. Made using video provided by the Lund Vision Group.

likely that they could recover their bearing after such physical disruption. Instead, this robust orientation performance points to the use of allothetic orientation cues (Cheung et al., 2007). Indeed, dung beetles appear to have a range of orientation cues available; work to date has found beetles able to orientate to: the sun, the moon, polarised skylight (at night, or during the day), the bright band of the Milky Way, intensity gradients, spectral information, and even the wind (Dacke et al., 2021) (see Table 2.1 for a breakdown by species). Moreover, it has recently been discovered that the beetles integrate multiple orientation cues simultaneously (Dacke et al., 2019; Khaldy et al., 2021, 2022)(Paper I; Paper II) (see Chapter 3).

2.2 Available orientation cues

Point sources

Point sources are single salient sources of light, the sun and moon being the most prominent natural examples. Artificial point sources can also be created by using LEDs or lamps. *Kheper lamarcki* and *Scarabaeus satyrus* will both treat a green LED as if it were the sun or moon (el Jundi et al., 2015b). All ball rollers studied to date have been shown to be able to use point sources as an orientation cue (see Table 2.1). The sun appears to be a primary orientation cue among diurnal beetles, however, this clearly depends on the species' natural environment (Khaldy et al., 2021, 2019b). For example, the woodland-living *Si. fasciculatus* will ignore point sources in favour of polarised skylight (see below) unless point sources are presented in isolation (Khaldy et al., 2019b). Dial group also plays a clear role. *S. satyrus* (nocturnal) and *S. zambesianus* (crepuscular), will preferentially use polarised light under natural circumstances (Dacke et al., 2003a,b, 2004; el Jundi et al., 2015b) though both will respond to an isolated artificial point source (Dacke et al., 2004; el Jundi et al., 2015b). Interestingly, *S. satyrus* will also respond to the position of the sun if made to roll during the day. Conversely, if forced to

Table 2.1: Dung beetle species studied for straight-line orientation. Orientation cues are given as well as a list of additional research performed using the same species. *These species are described as using the sun in (Byrne et al., 2003), then using polarisation in (Dacke et al., 2011) but Dacke et al. (2011) do not state a reason for this change in view. ***K. lamarcki* is described as using polarised light in (Dacke et al., 2011), but evidence for orientation to polarised light comes later from (el Jundi et al., 2014). ***These beetles will respond to a mirrored sun only if polarised light is simultaneously manipulated.

Species	Dial group	Cues confirmed	Also studied for
<i>Scarabaeus rugosus</i>	Diurnal	point source (artificial) (Byrne et al., 2003), polarised light* (Dacke et al., 2011)	
<i>S. rusticus</i>	Diurnal	point source (artificial) (Byrne et al., 2003)	
<i>Kheper nigroaeneus</i>	Diurnal	point source (artificial, sun) (Byrne et al., 2003; Khaldy et al., 2021)	bearing selection (Baird et al., 2010), dance behaviour (Baird et al., 2012), visual ecology of compass cue integration (Khaldy et al., 2021), landmark based orientation (Dacke et al., 2012)
<i>Pachylomerus femoralis</i>	Diurnal	point source (sun) (Byrne et al., 2003), polarised light* (Dacke et al., 2011)	
<i>Gareta unicolor</i>	Diurnal	point source (artificial, sun***) (Byrne et al., 2003; Khaldy et al., 2021), polarised light* (Dacke et al., 2011; Khaldy et al., 2021)	visual ecology of compass cue integration (Khaldy et al., 2021)
<i>S. zambesianus</i>	Crepuscular	point source (artificial) (Dacke et al., 2004), polarised light (Dacke et al., 2003a,b, 2004, 2011)	morphological comparison of dial groups (Tocco et al., 2019)

<i>S. satyrus</i>	Nocturnal	point source (sun, artificial) (el Jundi et al., 2015b; Smolka et al., 2016), polarised light (Dacke et al., 2011; el Jundi et al., 2015b), Milky Way (Dacke et al., 2013; Foster et al., 2017), intensity gradient (Foster et al., 2017, 2021)	gross neuroanatomical analysis (Immonen et al., 2017), fine neuroanatomical analysis (el Jundi et al., 2018), electrophysiology (el Jundi et al., 2015b), orientation precision under varying degrees of polarisation (Foster et al., 2019), the effects of light pollution on orientation (Foster et al., 2021), activity window (Smolka et al., 2016; Tocco et al., 2021), morphological comparison of dial groups (Tocco et al., 2019)
<i>K. lamarcki</i>	Diurnal	point source (sun, moon, artificial) (Dacke et al., 2014; el Jundi et al., 2015b,a; el Jundi et al., 2016; Smolka et al., 2016; Khaldy et al., 2019b,a; Dacke et al., 2019; Khaldy et al., 2022)(Paper I; Paper II), polarised light** (Dacke et al., 2011; el Jundi et al., 2014; el Jundi et al., 2016; Khaldy et al., 2022), intensity gradient el Jundi et al. (2014, 2015a), colour (el Jundi et al., 2015a; Yilmaz et al., 2022a), wind (Dacke et al., 2019)(Paper I; Paper II)	gross neuroanatomy (Immonen et al., 2017), fine neuroanatomy (el Jundi et al., 2018), electrophysiology (el Jundi et al., 2015b), compass cue integration (Dacke et al., 2019; Khaldy et al., 2022)(Paper I; Paper II), effect of step-size on orientation performance (Khaldy et al., 2019a), bearing selection (Khaldy et al., 2019a), activity window (Smolka et al., 2016; Tocco et al., 2021), habitat (Khaldy et al., 2019b), morphological comparison of dial groups (Tocco et al., 2019)
<i>S. ambiguus</i>	Diurnal	point source (sun) (Khaldy et al., 2019a)	effect of step-size on orientation performance (Khaldy et al., 2019a)
<i>Sisyphus fasciculatus</i> ***	Diurnal	point source (artificial) (Khaldy et al., 2019b), polarised light (Khaldy et al., 2019b, 2021)	habitat (Khaldy et al., 2019b), bearing selection (Khaldy et al., 2019b), visual ecology of compass cue integration (Khaldy et al., 2021)
<i>G. nitens</i>	Diurnal	point source (sun***) (Khaldy et al., 2021), polarised light (Khaldy et al., 2021)	visual ecology of compass cue integration (Khaldy et al., 2021)

roll at night, *K. lamacki* will continue to use the available point source (the

moon) for orientation (el Jundi et al., 2015b; Smolka et al., 2016) which makes this diurnal beetle the only one on our list (Table 2.1) to perform lunar orientation under a natural sky.

Orientation to point sources under natural conditions is usually tested by shading the beetle's view of the sun or moon, then using a mirror to reflect the source (resulting in an apparent 180° rotation of the cue). Artificial sources are usually tested inside a blackout tent so that the stimulus can be completely controlled (e.g. see Appendix C), however there are some examples of artificial sources being used outdoors to create a point source cue (Dacke et al., 2004; Smolka et al., 2016).

Polarised light

If we consider a light ray emanating from the sun, the ray can be described as a wave oscillating in the plane perpendicular to the direction of travel. *Unpolarised* light oscillates equally along all axes which lie in the plane, *linearly polarised* light oscillates along exactly one axis. Light can also be (and is usually) partially polarised and the degree to which it is polarised is called *degree of linear polarisation*, the angle of the (major) axis along which the light is polarised is referred to as the *angle of polarisation*. Our ray of solar light is completely unpolarised until it collides with particles in the atmosphere where it is scattered, creating a regular, wide-field pattern of polarisation across the sky (see Figure 2.3 for an illustration). Only a few samples of this pattern are needed to be able to recover the position of the sun (Gkaniyas et al., 2019) (see Section 2.3). The effect of polarisation is invisible to humans (but see (O'Shea et al., 2021)), but many other animals are capable of sensing polarised light. (Note that light can also be 'circularly polarised'; circular polarisation is not known to be of behavioural relevance to ball rolling beetles and will not be discussed.)

Orientation to polarised light in dung beetles was first demonstrated in the crepuscular *Scarabaeus zambesianus* (Dacke et al., 2003a,b), followed by *Kheper lamarcki* (el Jundi et al., 2014), *S. satyrus* (el Jundi et al., 2015b), *Sisyphus fasciculatus* (Khaldy et al., 2019b), and finally *Gareta nitens* (Khaldy et al., 2021). While they are able to use polarised light for orientation, *K. lamarcki* have only been shown to do so under a natural sky when the sun is unavailable. In a light-controlled setup, *K. lamarcki*'s use of polarised light can be modulated by changing the intensity or degree of polarised light in combination with a point source cue (Khaldy et al., 2022). *S. satyrus*, *S. zambesianus*, and *Si. fasciculatus* all appear to use polarised light as their primary orientation cue. It remains to be seen if they also respond to polarised light in a weighted fashion. The beetles do not appear to expect any set relationship between the position of the sun

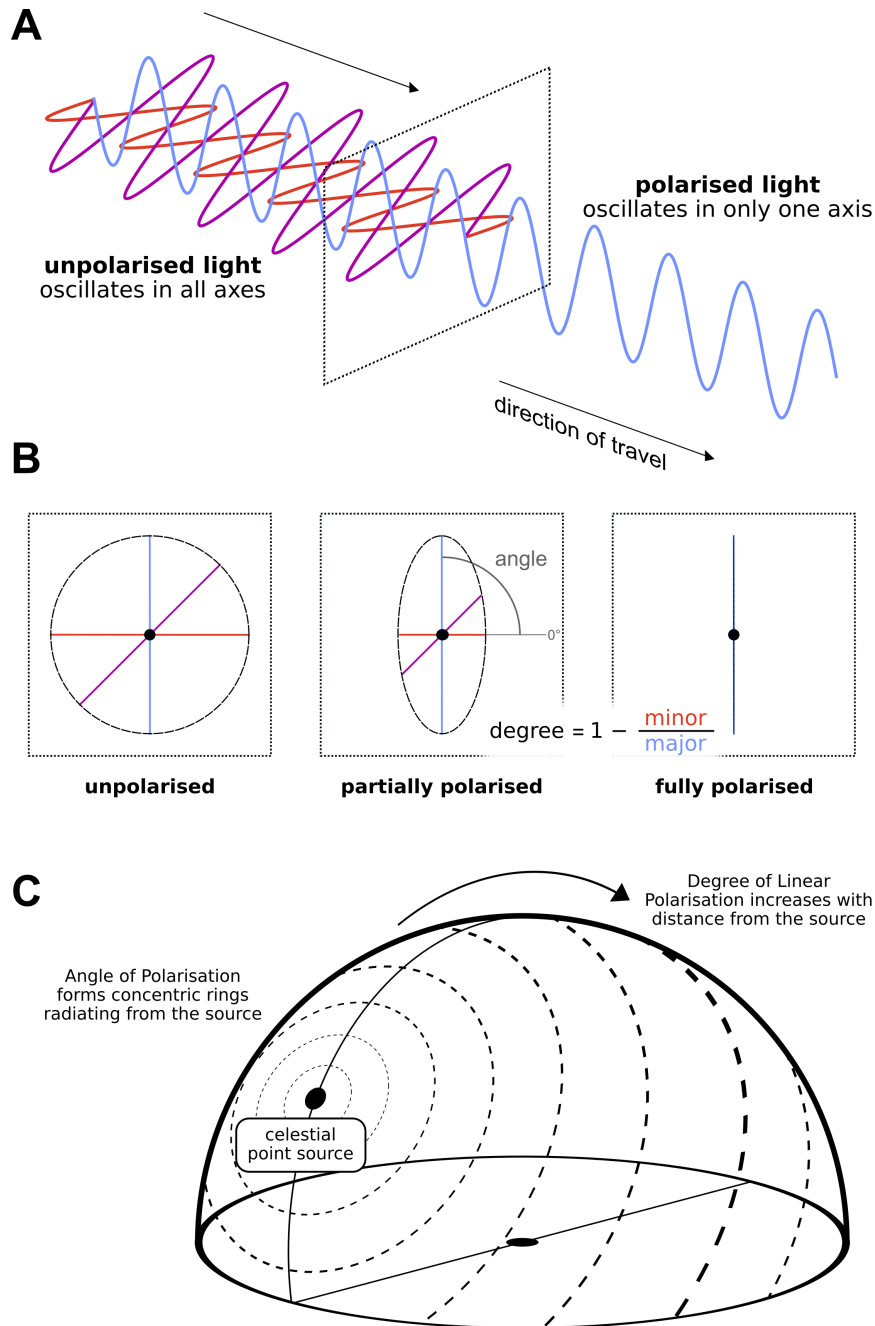


Figure 2.3: Conceptual illustration of polarised skylight. **A** Light which is *unpolarised* oscillates in all axes in a plane perpendicular to the direction of travel (dotted plane, only three shown for clarity), light which is *polarised* oscillates in one axis. **B** Projection of oscillatory axes into the perpendicular plane. Properties of polarisation can be visualised by drawing an ellipse around the tips of the waves. This results in different degrees of ellipse depending on how polarised the light is. The degree of polarisation is given by the ratio of major to minor axes, and the angle of polarisation is given by the angle of the major axis (relative to some reference, here given as horizontal). **C** An illustration of the polarisation pattern over the natural sky.

and pattern of polarised light, they will happily orientate to unnatural patterns (el Jundi et al., 2016).

Polarised light orientation is tested by using a linear polariser (a material which only transmits light which is polarised at a specific angle) placed over the animal (Dacke et al., 2003a; el Jundi et al., 2016; Khaldy et al., 2022, 2019a). This creates a stimulus which is of uniform degree of linear polarisation across the field of view with the only cue to orientation being in the angle. Thus, polarised light experiments tend to generate axial orientation patterns, demonstrating that, under the filter, beetles cannot determine the difference between an orientation of say, θ and $\theta + 180^\circ$. However, it remains unclear if this is a limitation of the beetles or of the filter stimulus used to test the beetles (Gkanias et al., 2019). The degree of polarisation can be regulated by adding diffusing filters below the polarising filter which will partially depolarise the light (Foster et al., 2019; Khaldy et al., 2022).

Intensity gradients

An intensity gradient is a wide-field monochromatic gradient from dark to bright over the sky. This could be a smooth gradient or could be considered any light pattern where there is a relatively bright and dim side of the field of view.

Both *Kheper lamarcki* and *Scarabaeus satyrus* have been demonstrated to orientate to intensity patterns though the methods used have been mixed. *K. lamarcki* was tested by using an intensity filter which was constructed from strips of material which transmitted different amounts of light across the diameter of the filter (el Jundi et al., 2014); in addition there have been experiments with two LED point sources of different intensities situated 180° apart (el Jundi et al., 2015a). *S. satyrus* has been tested using a similar dual LED stimulus, but has also been tested with a full LED band (1.6m radius, arching over the animal) set to varying degrees of contrast between each end (Foster et al., 2017).

The Milky Way

On moonless nights, the Milky Way galaxy creates a bright band of light across the South African sky (Foster et al., 2017). *Scarabaeus satyrus* has been found to use this band for orientation on moonless nights (where lunar polarisation is unavailable) and in a planetarium (Dacke et al., 2013). Additional analysis has been performed by using patterns of white LEDs arranged to mimic the Milky Way from the perspective of the beetle (Foster et al., 2017). The combined research on this phenomenon indicates that the beetles do not use the star patterns themselves or a sin-

gle bright reference, but appear to rely on the pattern of contrast available along the major axis of the band of the Milky Way (Dacke et al., 2013; Foster et al., 2017). Of particular note is that *S. satyrus* could not orientate under light patterns which provided an orientation cue in their spatial arrangement but lacked any intensity difference between the opposing ends of the pattern (Foster et al., 2017). It therefore seems likely that the sensory hardware which enables orientation to intensity gradient cues more generally also permits *S. satyrus* to orientate to the Milky Way.

Wind

Wind is the most recently discovered orientation cue in dung beetles and so far, *Kheper lamarcki* is the only ball roller which has been found to use the wind for orientation. In the field, orientation to wind cues can be tested by placing wind barriers and using large domestic fans (Dacke et al., 2019). In the laboratory, custom wind generators (made up of three smaller fans) are used to try and provide the beetles with uniform laminar flow across the radius of the experimental arena (Paper I) (see Chapter 7 and Appendix C).

Spectral gradients

The visual system of *Kheper lamarcki* is sensitive to green and UV light (el Jundi et al., 2015a; Yilmaz et al., 2022a). As sunlight is scattered by the atmosphere, more green light is transmitted in the solar region of the sky and more UV is transmitted in the anti-solar region. This creates a spectral gradient across the sky which can serve as an orientation cue. The most comprehensive exploration of spectral orientation in dung beetles is provided by Yilmaz et al. (2022a). Generally they report that spectral orientation can be explained by perceived differences of intensity of the different wavelengths presented (i.e. an intensity gradient cue). However, in the particular case of UV (365nm) and green (530nm) light, the beetles do appear to orientate according to the colour gradient (Yilmaz et al., 2022a). It is noted that beetles do not interpret green light as coming from the sun and UV as being anti-sun, as has been described in bees (Brines and Gould, 1979). The beetles do not appear to have any built-in expectation about the relationship of green and UV signals across the sky.

The effect of landmarks

Terrestrial landmarks do not appear to affect dung beetle straight-line orientation behaviour. In an experiment using *Kheper nigroaeneus*, Dacke

et al. (2012) demonstrate that the beetles ignore landmarks in favour of skylight cues. In particular, these beetles are disorientated under overcast skies (where the terrestrial panorama could be of use in guiding behaviour), showing that terrestrial landmarks do not have any effective input to the beetle's compass.

The definition of a terrestrial landmark is a little fuzzy. It is obvious that a relatively large proximal landmark, such as a tree or a human holding a shading board, may affect a beetle's view of the sky and therefore its behaviour. The distinction is that, visually, it is the skylight information that guides the beetle.

2.3 Sensory system

The eyes [point sources, intensity gradients, spectral gradients]

Ball rolling beetles have two superposition compound eyes, though in some species a cuticular ridge (the canthus) partially or completely separates each eye into a dorsal and ventral section (for simplicity, these are treated as eyes in their own right). For the purposes of orientation, beetles appear to use only their dorsal eyes; beetles whose dorsal eyes are obscured by a cardboard cap are unable to orientate (Dacke et al., 2012, 2013).

It is not currently known how beetles compute their orientation from non-polarised visual cues. For monochromatic cues, one obvious strategy would be to average intensity across the dorsal field of view; this would align nicely with a proposed mechanism used to compute polarisation orientation (Gkaniias et al., 2019) (Gkaniias et al., in press) (see below). This would also suggest that any stimulus which generates an intensity gradient would use the same underlying computational mechanism. They do not appear to match views (Foster et al., 2017).

Eye size varies between species and generally correlates with body size (Khalidy et al., 2019b). Eye size is also related to dial group (Tocco et al., 2019). For example, *S. satyrus*, and *K. lamarcki* are of similar body size, but *S. satyrus* has much larger eyes to better collect light under the dim night sky (Smolka et al., 2016; Tocco et al., 2019). The crepuscular *S. zambesianus* has eyes which are of intermediate size (Tocco et al., 2019). The dorsal eyes can be further divided into a dorsal and ventral region.

The dorsal rim area [polarised light]

Insects detect polarised light using a specialised set of photoreceptors which run along the dorsal rim of the eye, a region therefore called the




Figure 2.4: Figure 2 (with caption) from [Dacke et al. \(2002\)](#); comparison of DRA vs. non-DRA rhabdoms in *Pachysoma striatum*. (Original caption) Electron micrographs showing cross-sections through rhabdoms of the dorsal eye. **A** Specialised ommatidia of the dorsal rim area. **B** Unspecialised ommatidia of the rest of the dorsal eye formed by seven receptor cells (1-7). The rhabdoms of the dorsal rim area are dumb-bell-shaped with orthogonal microvilli, in contrast to the slightly lobed rhabdoms that are found in the rest of the eye. Slightly lobed ommatidia can also be found in the entire ventral eye. Scale bar: 500nm

dorsal rim area (DRA). The DRA is defined by the morphology of the photoreceptors. In the insect eye, the photoreceptive structure is known as the *rhabdom*; the rhabdom is made up of different compartments called *rhabdomeres*, each of which contains a cluster brush-like photoreceptors known as *microvilli*. In the DRA the microvilli are arranged along two perpendicular axes; in the unspecialised ommatidia in the rest of the eye they are spread at intermediate angles. As an example, Figure 2.4 shows the different rhabdom morphologies in the diurnal dung beetle *Pachysoma striatum*. A clear distinction can be made between those in the DRA and those in the rest of the eye. Similar distinctions can be found in a variety of insects ([Labhart and Meyer, 1999](#)). The microvilli are naturally sensitive to light polarised along their length ([Land and Nilsson, 2012](#)) and therefore the perpendicular arrangement permits downstream cells to act as polarisation opponent detectors ([Labhart, 1988](#)) (also see Appendix A).

Where DRA rhabdoms have been identified in dung beetles, their size and frequency seem to reflect the visual ecology and behaviour of the animal. For example, *S. satyrus* and *S. zambesianus* have relatively large DRA rhabdoms ([Dacke et al., 2003b](#)) and the DRA of *S. zambesianus* takes up around half of the dorsal eye ([Dacke et al., 2003b](#)). These beetles operate under low-light conditions and both are known to use skylight polarisation as their primary orientation cue ([Dacke et al., 2004](#); [el Jundi et al., 2015b](#)). In contrast, the diurnal, savannah living *K. lamarcki* ap-

pears to have only a single row of polarisation sensitive ommatidia along the top edge of the dorsal eye (Khaldy et al., 2021, citing Dacke, unpublished data). Further differences between nocturnal and diurnal beetles can be found in the early stages of visual processing where *S. satyrus* has an additional region of the lamina (the first optic lobe) dedicated to information coming from the dorsal rim area (Immonen et al., 2017).

The fan-like arrangement of the rhabdoms in the DRA has been argued to be a consequence of the construction of the insect eye (Dacke et al., 2003b). However, this arrangement could serve to mimic the distribution of different angles of polarisation across the sky, allowing the animal to unambiguously determine its orientation (Gkanias et al., 2019) (Gkanias et al., in press) (see Appendices A and F).

The antennae

It is known that the antennae are used for wind detection in beetles (Dacke et al., 2019). It is not known how the beetles extract wind direction from their antennae or how this information directs straight-line orientation. Work in the fruit fly *Drosophila melanogaster* has shown that wind direction can be extracted by taking the difference between antennal displacements (Suver et al., 2019). Furthermore, a complete pathway from the antennae to the central complex (see below) has been described (Suver et al., 2019; Okubo et al., 2020). Given the high degree of conservation of insect neuroanatomy it seems plausible that beetles use more or less the same decoding technique and pathway, however this is an area of active research.

2.4 The central complex

The central complex (CX) (Figure 2.5) is widely believed to be the navigational centre of the insect brain (Honkanen et al., 2019) and is therefore of relevance to straight-line orientation. The CX is a set of four midline neuropils which has been identified in every insect brain studied to date (Honkanen et al., 2019). The canonical CX consists of the ellipsoid body (EB), protocerebral bridge (PB), fan-shaped body (FB), and the paired noduli (NO). A fifth neuropil, the asymmetric body, has also been included in fruit flies (Wolff and Rubin, 2018). It is worth noting that there is a difference in nomenclature between fruit flies and 'other insects'. Specifically, the EB and FB are often referred to as the lower division of the central body (CBL) and upper division of the central body (CBU) respectively. I generally use the fruit fly nomenclature for neuropils and neuron classes as the bulk of current insect neuroanatomical data comes from fruit flies.

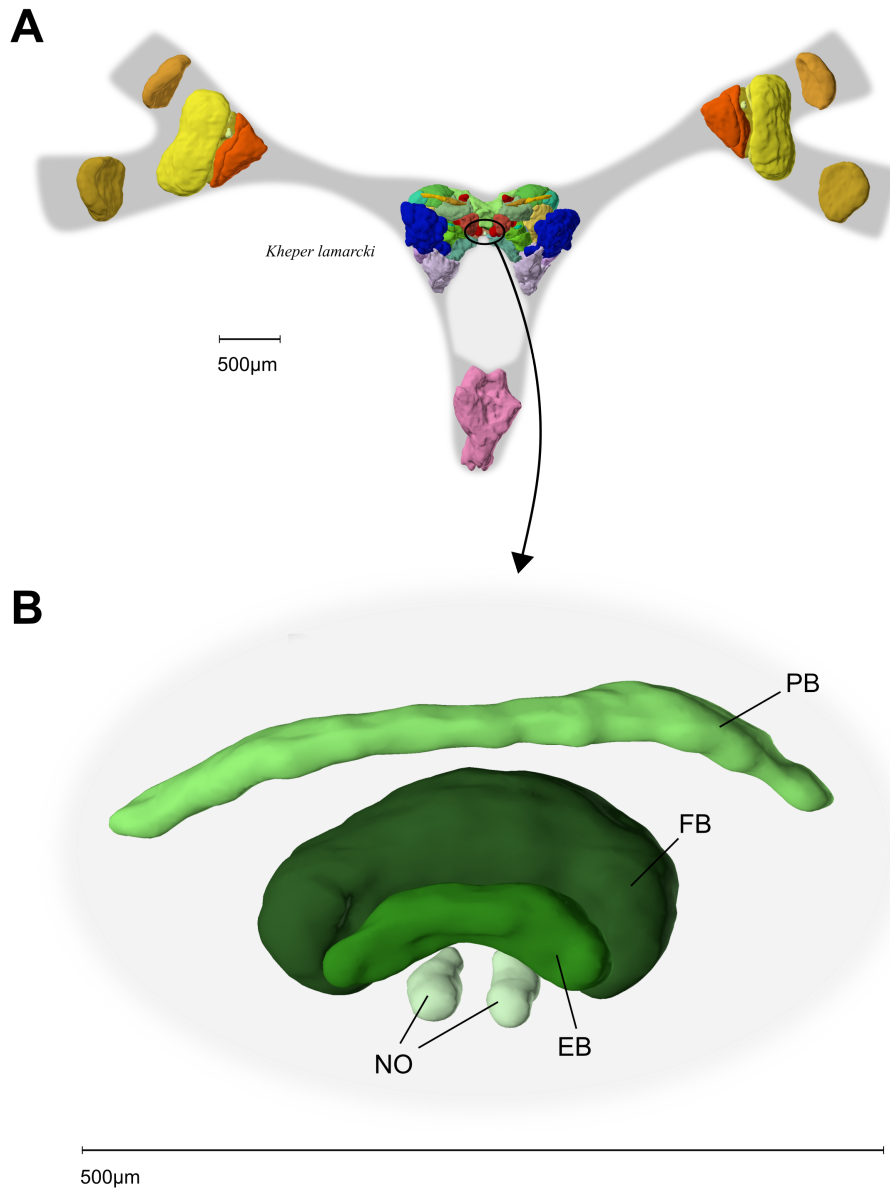


Figure 2.5: **A** High-level (anterior) view of the brain of *Kheper lamarcki*, different neuropils (identified by Immonen et al. (2017)) are shown as coloured volumes. **B** The central complex (CX) neuropils, the protocerebral bridge (PB), fan-shaped body (FB), ellipsoid body (EB), and noduli (NO). Neuropil images from InsectBrainDB, scale-bars are based on those provided by InsectBrainDB (Insect Brain Database) (Heinze et al., 2021) Full brain outline based on images provided by Immonen et al. (2017). The CX has been rotated (pitch) slightly for better neuropil visibility.

While the broad structure of the CX is consistently identifiable, there are clear differences in gross morphology between species (Honkanen et al., 2019), as well as connectomic differences (Pisokas et al., 2020). The CX has been identified and studied in *Kheper lamarcki* and *Scarabaeus satyrus* with no differences noted between the species (Immonen et al., 2017; el Jundi et al., 2015b, 2018).

The central importance and complexity of the CX appears to grow by the day, so we will focus on two functional roles which are of immediate interest in dung beetles, the storage of head direction, and goal direction (i.e. the compass needle and the bearing taken).

Head direction

In fruit flies, head direction is represented by a sinusoidal pattern of activity distributed across neurons in the ellipsoid body known as EPG neurons (Seelig and Jayaraman, 2015). Information about external cues is thought to reach the EPGs via a class of neurons known as ring neurons. Ring neurons innervate EPGs in the ellipsoid body and there are groups of ring neurons which are directionally tuned and respond separately to visual cues (Seelig and Jayaraman, 2013), polarised light cues (Hardcastle et al., 2021), and wind cues (Okubo et al., 2020). Once a head direction signal is established in the ellipsoid body, it persists in the absence of external cues and will update according to the animal's movements (Seelig and Jayaraman, 2015; Turner-Evans et al., 2017). These updates are known to be idiothetic in nature but it is not yet known whether they are driven by proprioception or motor efference (Turner-Evans et al., 2017; Fisher, 2022).

Neuron activity is more difficult to study in beetles (and other insects in general) when compared to the fruit fly; that said, intracellular recordings from *K. lamarcki* and *S. satyrus* indicate similar directional tuning (for point source and polarised light stimuli) in ring neuron and EPG homologues (el Jundi et al., 2015b). Further, an anatomical study conducted by el Jundi et al. (2018) showed that the fine anatomy of the beetle head direction circuit (what connects where) is generally consistent with descriptions of the locust CX (and other non-dipteran CXs). A later comparative modelling study by Pisokas et al. (2020) compared the head direction circuits of the locust and fruit fly and found them to be functionally consistent despite their differing connectivity (though they did note differences in response dynamics). Thus, while the anatomy of the beetle does not exactly match that of the fly, available evidence suggests that the functional role and behaviour of this circuit is conserved.

Goal direction

Beetles are largely unique amongst insects commonly studied for navigation in that their straight-line orientation does not have a specific end point (but see [Giraldo et al. \(2018\)](#)). Most of the literature on maintaining a bearing concerns path integration (integrating distance and direction over time to determine displacement). As a path integration process always maintains a current displacement vector (relative to the start point), an animal can always follow a straight path home regardless of its outbound path. There are many insects which are known to path integrate, including the South African dung beetle *Scarabaeus galenus* ([Dacke et al., 2020](#)). Interestingly, despite the extensive research on path integration in insects, whether the CX is actually involved is still unknown ([Hulse et al., 2021](#)).

Modelling work by [Stone et al. \(2017\)](#) proposed that a goal vector (home vector in this case) is stored as a sinusoid distributed across the population of PFN (CPU4) neurons. As an animal forages for food, PFN neurons ‘charge up’ by integrating speed input from optic flow and orientation information from the head direction circuit. On the homeward path, structured connections draw a comparison between the current head direction signal and PFN signal in order to steer the animal or agent towards its goal. While this model was based on the anatomy of the sweat bee *Megaopta genalis*, PFN neurons are known to exist in the ball-rollers *K. lamarcki* and *S. satyrus* ([el Jundi et al., 2018](#)).

While PFN neurons are not confirmed to be the substrate for goal direction in ball-rolling beetles, directional memories exhibit similar decay dynamics across species. In the jack jumper ant *Myrmecia crosslandi*, path integration memories have been shown to decay over time, and in response to cold-induced anaesthesia ([Pisokas et al., 2022](#)). The same has been shown in a path integrating dung beetle *Scarabaeus galenus* ([Yilmaz et al., 2022b](#)). While the effects of cold have not yet been investigated, preliminary data collected in *Kheper lamarcki* indicates that goal direction memory decays over time in a similar fashion to *S. galenus* (pilot data, not yet published). It is possible that the substrate for vector memory in path integrating insects serves as the substrate for goal direction memory in ball rolling beetles ([el Jundi et al., 2018](#)). The same neurons can be identified, and the memory dynamics appear to match insects which are more closely related to the bees used for the original modelling effort.

Recent work in fruit flies has proposed an alternative set of neurons which could act as goal direction neurons, namely FC2s ([Mussells Pires et al., 2024](#)). These neurons maintain a sinusoidal bump of activity (similar to the EPG bump in the EB) which appears to encode the animal’s goal direction. (Note that this means we might expect to see similar

memory decay dynamics to those described for path integration as the goal direction memory is encoded in the same way as is suggested for the path integration vector.) It was possible to direct flies by stimulating specific FC2 neurons. The authors go on to describe a full steering circuit. This circuit and its operation are given a more detailed treatment in Chapter 7, in which a bio-inspired implementation is provided to steer a robot (Chapter 6). In brief, PFL3 neuron pairs (putative steering neurons) sample from different EPGs (compass) which produces an imbalance in activity across the PFL3 pair. By using FC2 activity to select specific PFL3 pairs, the brain can use this activity imbalance to steer the fly in a specific direction (see Chapter 7). PFL3 neurons try to rotate the agent such that the EPG bump is in a specific location in the EB and the placement is selected by FC2 activity. Neurons which exhibit the characteristics of goal direction neurons have also been identified in the monarch butterfly *Danaus plexippus* (Beetz et al., 2023). Whether these correspond with FC2 neurons is not yet known.

2.5 The dung beetle dance and the orientation snapshot




Figure 2.6: An example dung beetle dance. The beetle climbs atop its ball, extends its head, rotates, and then rolls the ball away. The duration of rotation is known to vary (e.g. (Baird et al., 2012) define a dance as any ball-mounted rotation $> 90^\circ$). Consistency of dance direction is not known. In this instance, the beetle danced 180° .

Before setting out from the dung pile, dung beetles will first climb atop their balls, extend their heads, and pirouette. This rotation behaviour is known as the dung beetle *dance* (Matthews, 1963; Baird et al.,

2012). The dance is consistent enough to be of behavioural importance but inconsistent enough to make targeted research difficult. Baird et al. (2012) (working with *Kheper nigroaeneus*) suggest that, at the dung pile the initial dance stores the positions of the available orientation cues on the desired heading; subsequent dances are then interpreted as scans to efficiently re-align the beetle with its chosen direction. The first component of this interpretation has been validated by subsequent work. el Jundi et al. (2016) performed a related experiment whereby beetles (*Kheper lamarcki*) were made to dance in the centre of a circular arena and then permitted to roll (their ball) to the edge. The authors varied which cues were available during the dance and which were available during the rolling phase and show that beetles were aware of all cues present during the dance and that beetles will ignore cues which were not present during the dance. Thus, we know the dance appears to encode the spatial relationship between the available cues in what we will call an *orientation snapshot*. Note that in the literature, this is usually referred to as a ‘celestial snapshot’ (el Jundi et al., 2016). Historically, dung beetles were thought to respond only to celestial cues for orientation (Dacke et al., 2014); it is now well-established that wind can also be used (Dacke et al., 2019) (Paper I; Paper II).

The literature describes the orientation snapshot as playing two roles: (1) it encodes the relative positions of different orientation cues allowing beetles to orientate to arbitrary cue arrangements (el Jundi et al., 2016), and (2) the snapshot is also described as being the reference against which the beetle orientates (i.e. a compass bearing) (Baird et al., 2012; el Jundi et al., 2018). These conceptual roles are not necessarily in conflict but evidence suggests these functions are encoded in different parts of the brain (see Section 2.4).

The first conceptual role is the easiest to relate to the dance behaviour and the snapshot experiments performed by el Jundi et al. (2016). The spatial relationship of available cues appears to be encoded by plastic inputs to the insect central complex, where the dung beetle compass likely resides (el Jundi et al., 2015a, 2018) (and see (Paper II)). This plasticity is known to be regulated by idiothetic angular velocity via ExR2 neurons in *Drosophila* (Fisher et al., 2022); further, anti-Hebbian plasticity in combination with angular velocity inputs can encode cue spatial relationship and also results in cues being ignored if they are not present during learning (Paper II). The dance behaviour seems an efficient way to form these connections but it is not essential that individuals ‘dance’, only that they have experienced sufficient idiothetic angular velocity in combination with external cue rotation. If these connections are already well-structured when a beetle goes to leave the dung pile then there is no need for it to dance. Given that plasticity could be driven dynami-

cally by angular velocity (rather than ‘enabled’ during a dance), it seems plausible that dances would not always be necessary to encode relative cue positions (which may explain why the dance is so consistently inconsistent).

The second conceptual role (setting the goal direction) cannot be directly coupled to the dance; there is no obvious need to rotate to set the ‘goal direction’. However, changes in apparent goal direction could be indirectly coupled to the dance behaviour. Remaining agnostic as to whether FC2 neurons or PFN neurons represent the beetle’s goal, the goal direction is given its directional context by the compass. As already mentioned, the mappings between the compass inputs and compass neurons are highly plastic (Fisher et al., 2019) so it is not unreasonable to suggest that this relationship could drift. Such a drift would mean that the relationship between the compass and world had shifted and thus, if the neural representation of goal direction remained constant, the animal’s actual trajectory in the world would change. An example of this effect can be seen in the results presented in Chapter 7 where a drift in compass representation caused a drift in the orientation of a robot.

2.6 Summary

Ball rolling dung beetles perform a robust straight-line orientation behaviour for which they have multiple orientation cues available. Where other navigation behaviours can be difficult to study due to the interaction of multiple high-level navigation systems (e.g. the interaction of landmark guidance and path integration - see Chapter 3), the relative simplicity of beetle behaviour makes them a highly accessible model in which to study the insect compass. In particular, the recent discovery of cue integration in the dung beetle compass presents a unique opportunity to study a relatively controlled cue integration behaviour in a non-human animal.

Chapter 3

Cue integration

In the wider literature (i.e. that which focuses on humans), cue integration has largely been approached as a Bayesian inference problem (Trommershauser et al., 2011; Rahnev and Denison, 2018). In insects, cue integration studies split into three broad categories. *Descriptive* accounts provide textual descriptions and possibly statistical analyses of behaviour but do not provide a computational model which describes the integration process. *Algorithmic* accounts describe potential computational models of behaviour but do not provide concrete implementations of these models. *Computational* accounts will provide a concrete computational model and compare simulated results to real behavioural results. Reflecting the wider literature, algorithmic and computational accounts also tend to approach the observed behaviour from a Bayesian perspective (Collett, 2012; Legge et al., 2014; Wystrach et al., 2015; Hoinville and Wehner, 2018; Khaldy et al., 2022).

Here I will give a brief overview of the literature on insect cue integration. I will then discuss the Bayesian perspective on cue integration which forms the basis for the work presented in (Paper I). In addition, I want to take advantage of the experience gained since writing (Paper I) to provide a more critical examination of the utility of the Bayesian approach.

3.1 Insect accounts

The most prominent examples from the insect literature concern the integration of navigation system outputs to steer an animal; for example, combining path integration and visual navigation systems to navigate home (Narendra, 2007; Collett, 2012; Wystrach et al., 2015; Hoinville and Wehner, 2018). While this is arguably different to orientation cue integration, the underlying principles are identical. Each navigational sub-

system tells the animal how to steer, and the steering system integrates these directional cues to produce a single behaviour; analogously, each sensory system gives a different directional cue which is integrated into a compass to give a single direction estimate. The literature surveyed here is not meant to be exhaustive.

3.1.1 Descriptive accounts

It seems likely that the first cue integration experiment performed in insects is by von Frisch (1967, pg. 338-341), in which bee dances are observed with respect to the position of the sun and the angle of a linear polariser when these cues are placed in conflict. von Frisch (1967) reports that bees dance at intermediate angles to those predicted by each cue in isolation, with the relative influence given to each depending on the time of day. As solar elevation increased, the relative influence of the sun appeared to increase (though von Frisch (1967) argues that this is not actually to do with the solar elevation). Ugolini (2001) reports that the equatorial sandhopper *Talorchestia martensii* is capable of using both the position of the sun and the Earth's magnetic field to orientate along shorelines. The sandhoppers appeared to rely on the magnetic field only during the highest solar elevations. Müller and Wehner (2007) describe the desert ant *Cataglyphis fortis* orientating at angles intermediate to those predicted by solar azimuth and wind direction, and the weight given to each cue again depended on solar elevation. In these experiments, the ant dorsal rim area was covered with light-proof paint to remove polarised light as an orientation cue. Where polarisation is available, these ants will exclusively use polarisation for orientation during navigation (Wehner and Müller, 2006; Müller and Wehner, 2007). Narendra (2007) demonstrates that, in familiar surroundings, the Australian desert ant *Melophorus bagoti* will travel both towards a learned route and towards home if displaced, indicating that both landmark guidance and path integration processes are playing a role in steering the animal.

For dung beetles specifically, Dacke et al. (2019) present the first evidence of multimodal cue integration in dung beetles, wherein the beetles (*Kheper lamarcki*) are shown to integrate wind and sun cues to guide their straight-line orientation. The beetles are described as following a 'winner-take-all' strategy and whether they use the sun or the wind depended on solar elevation. This raises an interesting question as to when an animal can be considered to be integrating different cues.

Other cue integration literature draws a distinction between *cue capture* and cue integration, but on review, this distinction seems fuzzy. Cue capture describes the case where a single cue dominates the percept (Knight et al., 2014; Ernst and Bühlhoff, 2004). However, cue capture is

not a strict separation; the non-dominant cue can still have some influence. In (Paper I) we drew a strict distinction between winner-take-all (where one cue exclusively drives beetle behaviour) and cue integration (where both cues have an influence, even if one dominates). This distinction seemed fair at the time as we wanted to know if wind and sun cues were both responsible for guiding behaviour or if the animals switched between them (i.e. trying to distinguish between cue capture and strict cue switching). However, this strict separation ignores the fact that the beetles learn about all cues, even if they are not dominant (el Jundi et al., 2015b; el Jundi et al., 2016).

If an animal follows one cue exclusively but still learns about the relative position of another, does this count as cue integration? An animal could maintain two separate orientation estimates and switch between them in such a way that neither one influences the other. (Interestingly, Dacke et al. (2019) actually describe the beetles as having separate sun and wind compasses and the ability to transfer information between them.) Whether or not this counts as ‘cue integration’ seems to depend on implementation. (i.e. Does this information get combined in the brain?) Subsequent investigations have shown that the beetles do use multiple cues simultaneously (Khaldy et al., 2022) (Paper I), and this can also be explained as an anatomical integration (Paper II). Dacke et al. (2019)’s use of ‘winner-take-all’ could instead be replaced by the weaker ‘cue capture’.

A more convincing behavioural demonstration is given by Khaldy et al. (2021). The authors investigate the interaction of the sun and polarised light in beetle species from regions with different visual surroundings. *Sisyphus fasciculatus*, a woodland-living dung beetle species, preferentially orientated to polarised light. Conversely, *Kheper nigroanaeus*, a savannah-living beetle, preferentially orientated to the position of the sun. Both species were capable of using the other cue and their behavioural differences were explained as adaptations to their visual surroundings. Skylight polarisation is a wide-field cue which can be fully reconstructed from only a few sparse samples. Intensity cues are not so robust to occlusion (and so is less suited to environments with lots of canopy cover).

A third species, *Garetta unicolor* (also savannah living) was shown to ignore manipulations of both the sun and polarised light (Khaldy et al., 2021). However, when both cues were changed together, changes in bearing became uniformly distributed. Moreover, when cues were returned to their original position, the beetles resumed their original bearing. In all cases the beetles continued to follow straight paths (so they were not disorientated), and the return to their original bearing indicates that they were still trying to follow their original bearing when the

cues were manipulated (behavioural goals had not changed). Thus, for *G. unicolor* both cues were apparently working together to drive the behaviour of the animal.

While this is a better behavioural demonstration of cue integration than [Dacke et al. \(2019\)](#), it is not completely unambiguous. For example, probabilistic cue switching could perhaps explain this behaviour (e.g. if half of the population followed the sun and half followed polarised light, then a similar result could be obtained). This form of ambiguity is not unique to [Khaldy et al. \(2021\)](#) and reflects a more general problem with cue integration literature. Observing the average behaviour of a population or an individual does not always allow us to distinguish between different computational models which could generate that behaviour ([Scarfe, 2022](#)).

3.1.2 Algorithmic accounts

These examples are still technically ‘descriptive’ as they do not provide any computational modelling. However, examples included here explicitly suggest that behaviour could be explained by Bayesian principles (discussed in Section 3.2) which starts to carry implications about the behaviour and experimental manipulations. For the following discussion, note that *reliability* is defined as the inverse variance of a noisy estimate, see Section 3.2 for a detailed discussion.

[Collett \(2012\)](#) noted that desert ants (*Cataglyphis fortis*) will take intermediate paths to those predicted by path integration and landmark guidance and proposed Bayesian inference as a potential model to explain this behaviour. While the term ‘Bayesian inference’ is arguably quite vague, it is likely that the [Collett \(2012\)](#) meant that behaviour was described by a specific linear model (Equation 3.13); a weighted average of navigational systems, weighted by the relative reliability of those systems (see Section 3.2).

[Leggett et al. \(2014\)](#) perform two cue integration experiments in the desert ant *M. bagoti*. The first is a straightforward conflict experiment in which terrestrial and celestial navigation cues are set in 0° , 90° , 180° , and 270° of conflict. They observe ants taking intermediate courses between the conflicting cues and conclude that the ants are integrating the two sources of information. In the second experiment, the authors attempt to manipulate the weights associated with each cue by displacing ants to unfamiliar environments. As environmental familiarity dropped, the weight associated with terrestrial cues also dropped. The authors interpret this as reliability-based weighted average of celestial and terrestrial cues (making reference to the same linear cue integration model as ([Collett, 2012](#))). While their displacement did increase spread in the popula-

tion, it is not clear how an unfamiliar environment increases the variance (decreases the reliability) of a visual guidance system. In addition, reliability is meant to be quantified per individual which is not the case here (see discussion below).

[Khaldy et al. \(2022\)](#) give the most recent (and most convincing) evidence of multimodal cue integration in ball-rolling dung beetles (*Kheper lamarcki*). The authors show that the beetles will respond flexibly to polarised light and point source cues. In particular they show that, in cases of cue conflict, the beetles will take intermediate courses which are dictated by the degree of polarisation presented. The manipulation is assumed to affect the reliability of the polarisation cue but no evidence is given that this is the case. The authors conclude that *K. lamarcki* integrates orientation cues in a 'Bayesian manner', presumably referring to the same linear model as [Collett \(2012\)](#).

3.1.3 Computational accounts

There are two prior computational accounts ([Wystrach et al., 2015](#); [Hoinville and Wehner, 2018](#)). These accounts differ in their modelling approaches which results in some key differences in outcome. They are therefore discussed in more detail.

[Wystrach et al. \(2015\)](#) provide the first computational account of insect cue integration (using the desert ant *Cataglyphis velox*). They perform a similar cue integration experiment to ([Legge et al., 2014](#)), but in this instance the authors provide different candidate models which are fit to, and subsequently compared against, the behavioural data. Their behavioural data is concerned with the integration of path integration (PI) and landmark guidance (LG) navigation and their core model is the reliability weighted average suggested by [Collett \(2012\)](#).

Following [Cheung et al. \(2007\)](#), [Wystrach et al. \(2015\)](#) provide a quantitative prediction for the change in *directional variance* of a PI estimate over time. The authors argue that the directional variance is proportional to $\frac{d_T}{d_H^2}$ (distance travelled divided by squared displacement). That is, if an ant is travelling on a relatively straight path ($d_T \approx d_H$), then directional variance decreases inverse proportionally to displacement. However, if the path is tortuous, this can increase the variance of the directional estimate. The example given by [Wystrach et al. \(2015\)](#) is $d_T = 13$, $d_H = 3$, $\frac{d_T}{d_H^2} = 1.4$ as opposed to 0.33 for a straight path. These variance estimates are used to determine what the relative weight of PI should be for a given home vector. This model is then fit to the behavioural data by choosing the LG variance which best matches the apparent relative weighting of LG and PI.

The authors perform two behavioural experiments. In experiment 1,

ants are trained to run along a channel to a feeder. Experienced ants are captured at different distances along the channel ($0m$, $1m$, $3m$, or $7m$), given food, and released in a familiar location, such that the direction indicated by PI and LG will conflict (by 110°). They then examine the angles taken by the ants to see where they lie with respect to the PI and LG predictions. They report that the ants integrate PI and LG strategies in accordance with a weighted average, where PI weight is proportional to its directional reliability (as predicted by the model).

In experiment 2, the authors capture ants at $3m$ and allow them to run for five minutes in a fluon-lined pot such that their distance travelled increases compared to their home vector length; with the expectation that this will change the weight given PI. This is then compared to a quantitative prediction of the weight based on distance travelled and home vector length. On average, the ants travelled for $10m$ in the pot and the authors predicted that PI weight should decrease as distance travelled grew while displacement did not. They instead found that PI weight did not change as expected with the distance travelled. They argue that the ants could be using home vector length as a proxy for directional variance, rather than maintaining a true estimate of the variance which should have been affected by the additional distance travelled.

It is critical to highlight that, while the authors comment on the average distance travelled in the pot, they do not discuss the average displacement. They seem to assume that the ants did not increase their home vector length while running in the pot (i.e. that they all ran in circles). That said, if we assume that ants travelled in straight lines in the pot, then PI weight should increase to the point that PI alone predicts behaviour which is not reflected in their behavioural data. Without also knowing the average displacement of ants in the pot, it is difficult to interpret this result.

[Hoinville and Wehner \(2018\)](#) change tack by providing an exclusive modelling study without original behavioural data. They aim to provide an explanation for previous cue integration experiments (including [Legge et al., 2014](#) and [Wystrach et al., 2015](#)), again in terms of Bayesian principles, but accounting for the angular nature of the problem. [Hoinville and Wehner \(2018\)](#) correctly note that, for noisy directional cues, the classic linear model (given below and used by [Wystrach et al. \(2015\)](#)) is not appropriate. Instead, they use a model derived from Bayes' theorem for cues with von Mises noise ([Murray and Morgenstern, 2010](#)) which can be expressed as a weighted vector sum. The authors then fit this model to a selection of previous behavioural experiments in ants and bees and show that the model is capable of generating data which cannot be statistically distinguished from the true experimental data. In contrast to the modelling given by [Wystrach et al. \(2015\)](#),

Hoinville and Wehner (2018) argue that the directional variance of a PI estimate does not depend on distance travelled, and instead depends only on displacement.

Their argument is as follows. A path integration vector results from the summation of discrete step vectors which tell us the distance and direction travelled on each step. If we assume step lengths are constant then an integrated home vector can be expressed as

$$\mathbf{v}_H = L \sum_{t=0}^T e^{i\theta_t} \quad (3.1)$$

where L is a step length and θ_t is the noisy compass angle at time t , with von Mises error $\epsilon \sim M(0, \kappa_\epsilon)$. If we multiply both sides by κ_ϵ we get

$$\kappa_\epsilon \mathbf{v}_H = L \sum_{t=0}^T \kappa_\epsilon e^{i\theta_t} \quad (3.2)$$

The position of κ_ϵ on the right-hand side is deliberately chosen for illustration. This scaling means that each step vector is ‘weighted’ according to the reliability of the compass. The magnitude of the sum of such weighted vectors (termed ‘belief vectors’ by Hoinville and Wehner (2018)), encodes the angular variance of the sum (Murray and Morgens-tern, 2010). In context, this means that the directional variance of the home vector shrinks with displacement but does not depend on distance travelled at all. Hoinville and Wehner (2018) use this reasoning to argue that ants do indeed weight PI and LG according to their relative variance, as opposed to the proxy mechanism proposed by Wystrach et al. (2015).

It is worth noting that Hoinville and Wehner (2018) do not include step-size error in their model. Further, if the compass reliability is itself variable over the course of PI vector construction, then PI vector length will no longer be proportional to the angular variance. An agent would need to retain a separate accumulator which weighted every step according to the instantaneous reliability of the compass. In reality, scaling home vector length by average compass reliability may be sufficient.

3.1.4 Integrating angular velocity and external cues

A related, though conceptually separate cue integration problem is the integration of idiothetic angular velocity with allothetic orientation cues in the insect head direction circuit. While this is a relevant puzzle piece, it is not explored in detail in this thesis (beyond some discussion in (Paper II) and Chapter 7).

Neuroethological accounts for fruit flies have been given by [Seelig and Jayaraman \(2015\)](#) and [Turner-Evans et al. \(2017\)](#), which demonstrate firstly that angular velocity inputs can update the head direction circuit in darkness (using angular velocity inputs), and secondly how these inputs arrive in and update the compass. The first model implementation of this integration process is given by [Cope et al. \(2017\)](#) and the combination of angular velocity and external cue information is now usually included as a standard feature of head direction models (e.g. ([Goulard et al., 2021](#))). Anatomically derived circuits from fruit flies and locusts bear a striking resemblance to a theoretical model of the rodent head direction circuit proposed by [Skaggs et al. \(1994\)](#) ([Pisokas et al., 2020](#)). This integration process has also been modelled as a Bayesian phenomenon where idiothetic and allothetic compass inputs are weighted according to their relative reliability ([Kutschireiter et al., 2023](#)).

From the literature it seems that this is largely considered a solved problem but the interaction of the different neurons involved seems to be quite complex. Idiothetic angular velocity regulates plastic connections between allothetic compass inputs (ER neurons) and the compass neurons (EPGs) ([Fisher et al., 2022](#)). Further, angular velocity should be the main influence on the compass during learning in order to form meaningful connections; this requirement is discussed explicitly in ([Paper II](#)) but has been quietly assumed in other models ([Cope et al., 2017](#); [Dan et al., 2021](#)).

In order for this learning mechanism to work, angular velocity inputs need to be upregulated relative to allothetic cue inputs during learning (i.e. during rotation), and the angular change they impart on the compass must be tightly coupled to the apparent rotation of the world. The changes in synaptic strength during learning will also change the relative inputs of allothetic and idiothetic information as learning occurs. Changing relative inputs to the compass seems to change the response dynamics throughout the circuit due to the high number of recurrent connections (personal observation). The proposed angular velocity update and learning mechanism is elegantly simple, but the fact that this physiological dance works is quite remarkable (see Chapter 7). This interaction deserves more targeted exploration. For example, it would be useful to know what feedback mechanisms are required for activity regulation and how these could be implemented.

3.2 The Bayesian perspective

3.2.1 Why Bayes?

Perception is inherently probabilistic. Sensed values are never perfect and even where sensing is relatively reliable, the stimulus itself can be noisy. It therefore makes sense to discuss issues of perception in terms of distributions of possible events, that is, in terms of probabilities. Perception then becomes an estimation problem: given a set of sensor signals, what is the stimulus which is most likely to have caused these signals. This idea was first proposed by [von Helmholtz \(1925\)](#), translated into English as unconscious conclusion or unconscious inference. Animals are thought to infer a likely state of the world based on a combination of past experience and noisy sensor inputs. Bayes' theorem provides a theoretical framework to build models of perception from this perspective (i.e. Bayesian Encoding models which treat the brain as a Bayesian inference engine ([Lange et al., 2023](#)) - see Section 3.3).

3.2.2 Bayes' theorem in the context of cue integration

Bayes' theorem (or Bayes' rule) is a straightforward result from probability theory:

$$P(X|Y) = \frac{P(Y|X)P(X)}{P(Y)} \quad (3.3)$$

where $P(X)$ is the probability of some event X occurring, $P(X|Y)$ is the probability of X occurring *given* Y has occurred. $P(Y)$ acts as a normalising constant and it is common to ignore it:

$$P(X|Y) \propto P(Y|X)P(X). \quad (3.4)$$

We can interpret X as property we are trying to estimate and Y as a sensor reading which allows us to estimate that property (i.e. a cue). $P(X|Y)$ is referred to as the *posterior*, $P(Y|X)$ the *likelihood*, and $P(X)$ is the *prior*. Given a sensor reading $Y = y$, the goal is then to find the value for X which maximises the posterior ("What is the most probable state of the world X given sensory input $Y = y$?").

If we want to maximise the posterior, the first task is to estimate the prior distribution over possible values of X . This is immediately problematic. In most cases, it is practically impossible to estimate the prior probability of a certain stimulus $X = x$ (and therefore the distribution of X). A popular (but controversial) technique for dealing with the problem of priors is to assume that they are uniformly distributed ([Rahnev and Denison, 2018](#)). In certain contexts an experimenter may also try to

engineer a uniform prior in their test subjects by controlling the presentation of different stimuli; however, this does not guarantee that a test subject actually forms said prior (Rahnev and Denison, 2018).

If we carry the assumption of uniform priors, $P(X = x)$ becomes constant and can conveniently be ignored. Bayes' rule has been reduced to:

$$P(X|Y) \propto P(Y|X) \quad (3.5)$$

This can be applied to cue integration simply by considering multiple sensory inputs,

$$P(X|Y_1, Y_2) \propto P(Y_1, Y_2|X) \quad (3.6)$$

If we assume that Y_1 and Y_2 are independent ($P(Y_1, Y_2) = P(Y_1)P(Y_2)$) then we get

$$P(X|Y_1, Y_2) \propto P(Y_1|X)P(Y_2|X) \quad (3.7)$$

Equation 3.7 allows us to determine what value of X is most likely given sensory inputs Y_1 and Y_2 . It is worth noting that the assumption of independence is usually made without any verification (Scarfe, 2022) (though not always (Oruç et al., 2003)).

3.2.3 Removing abstraction

To know which sensor values are most likely given stimulus X we have to know how our sensors respond with respect to the true value of the stimulus. For example, we could show a human participant a bar and ask them how tall they perceive it to be. By recording their answers we can determine (1) how variable the participant's responses are and, (2) whether the participant's perception is biased. Response distributions can usually be fit by Gaussian probability density functions:

$$f(x; \mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2} \quad (3.8)$$

The mean μ describes the bias in the participant's responses and the variance σ^2 describes how variable the participant's responses were. For a given X , we can therefore describe $P(Y_1|X)$ and $P(Y_2|X)$ by Gaussian distributions $f(y_1; \mu_1, \sigma_1)$ and $f(y_2; \mu_2, \sigma_2)$, and Equation 3.7 becomes:

$$P(X|Y_1, Y_2) \propto f(y_1; x, \sigma_1)f(y_2; x, \sigma_2) \quad (3.9)$$

$$\propto f(x; y_1, \sigma_1)f(x; y_2, \sigma_2) \quad (3.10)$$

$$\propto f(x; \mu_1, \sigma_1)f(x; \mu_2, \sigma_2) \quad (3.11)$$

(The change from y to μ is simply to standardise notation for the mean.) Fortunately, the product of two Gaussian distributions has a known (proportionally Gaussian) form with:

$$\mu_c = \frac{\mu_1\sigma_2^2 + \mu_2\sigma_1^2}{\sigma_1^2 + \sigma_2^2} \quad \sigma_c^2 = \frac{\sigma_1^2\sigma_2^2}{\sigma_1^2 + \sigma_2^2} \quad (3.12)$$

(See Appendix B.) In the context of cue integration, this can be more usefully (if less intuitively) written as:

$$\mu_c = \frac{\frac{1}{\sigma_1^2}\mu_1 + \frac{1}{\sigma_2^2}\mu_2}{\frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2}} \quad \frac{1}{\sigma_c^2} = \frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2} \quad (3.13)$$

The inverse of the variance of the noise distribution is known as the *reliability* of the cue. The less noisy a cue is, the more reliable it is and vice versa. The mean of the combined distribution as written in Equation 3.13, can be seen to be sum of the means of the constituent distributions, weighted by their relative reliabilities. As the product of two Gaussians is proportionally Gaussian, the mean is therefore the combination which yields the most likely value of the stimulus. It can also be seen that the reliability of the combined distribution is greater than reliability of either constituent distribution.

Equation 3.13 makes a concrete prediction about how combined cue estimates should be distributed. Where the prior is uniform and the cue noise distributions are independently Gaussian, Bayes' theorem predicts that the combined estimate will be proportionally Gaussian with mean (μ_c) and reliability ($1/\sigma_c^2$). The combined reliability will be greater than that of either of the constituent cues and, moreover, it can be proved that reliability-based weighting provides the best possible estimate of the stimulus (Oruç et al., 2003; Shahar, 2017).

This is one method (see Section 3.3.1) of arriving at the classic linear model for cue integration (Ernst and Banks, 2002). Using the same reasoning Murray and Morgenstern (2010) provide expressions for the integration of angular cues with von Mises noise distributions. This work was used as the basis of the modelling for (Paper I) which then led to the neural model implementation in (Paper II).

3.3 Discussion

The body of literature which uses Bayes' theorem to analyse animal behaviour one form or another is vast. It would therefore be wise to set some scope for the following discussion.

There is an ongoing debate as to the general utility of using Bayesian models to describe behaviour (Bowers and Davis, 2012a; Griffiths et al.,

2012; Bowers and Davis, 2012b; Rahnev and Denison, 2018). This debate exists because of the way these models are typically constructed, applied, and communicated rather than the merit in the actual approach. In cue integration specifically, there is a critical issue in that ‘optimal’, ‘Bayes-optimal’, or ‘Bayesian’ behaviour has generally come to refer to Equation 3.13. That this model has become essentially the default analysis tool for cue integration can be seen most clearly in large recent surveys of the literature (Rahnev and Denison, 2018; Scarfe, 2022). Thus, despite the fact that Bayes’ theorem is presented above as a framework for understanding perception, which we then use to derive a specific rule, the literature has fixated on the rule itself. This is an issue in cue integration literature more generally, but has become a problem in insect literature too. Algorithmic accounts of insect cue integration describe the behaviour as Bayesian and the chain of references leads back to Equation 3.13. Computational accounts have referred both to Equation 3.13 (Wystrach et al., 2015) and to its circular counterpart (Hoinville and Wehner, 2018). I therefore think it is worth discussing the shortcomings of Equation 3.13 in general, but also as is specifically relevant to insects (note that the points raised apply equally to the circular variant used by Hoinville and Wehner (2018)).

In addition, there is a (recently highlighted) distinction between models of Bayesian Encoding and Bayesian Decoding (Lange et al., 2023). Bayesian Encoding models propose that Bayesian inference (or an approximation) is implemented in the brain. Bayesian Decoding models use Bayesian inference to construct models which decode neural activity to recover the original stimulus. This can be used to place restrictions on how well it is theoretically possible for the brain to perform (i.e. how much information is actually available to downstream neurons) or to provide an accurate decoding tool for analysis (for a clear example, see Salinas and Abbott (1994)). While many of the psychophysical examples of Bayesian cue integration simply use Equation 3.13 to describe behaviour, this is commonly accompanied by the strong claim that the brain is performing Bayesian inference (Knill and Pouget, 2004). There is a resultant conflict between those who treat Bayesian inference as a theoretical framework to describe behaviour and those who treat it as a theory of mind to explain behaviour (Colombo and Seriès, 2012; Rahnev and Denison, 2018). This appears to have caused a great deal of confusion as to the goal of the approach and the meaningfulness of a Bayesian result, with some arguing that we should be able to find neural correlates of priors, likelihoods, and posteriors (Knill and Pouget, 2004). I would therefore classify Equation 3.13 as a Bayesian Encoding hypothesis but I want to clarify that this may not be true for all authors who employ the model. I will not discuss Bayesian Decoding models as I want to focus

specifically on the issues with Equation 3.13 (and its circular counterpart) as commonly applied and its relevance to insect studies.

3.3.1 Linear or Bayesian?

The history of Equation 3.13 is a little unclear from the literature. Some works discuss it explicitly as a special case of Bayes' theorem (Ernst and Banks, 2002; Murray and Morgenstern, 2010; Wystrach et al., 2015). Other works describe Equation 3.13 as a 'linear model' and describe 'Bayesian' models as being conceptually separate (Oruç et al., 2003; Landy et al., 2011) (even if it is also noted that the 'linear' model is a special case of Bayes' theorem). Scarfe (2022) describes Equation 3.13 as a minimum variance unbiased estimator model which is 'related to Bayesian models of sensory perception' implying it is not a Bayesian model in itself.

Oruç et al. (2003) attempt to justify the separation by proving that the linear rule with weights given by reliability is optimal (gives a minimum variance estimate) for any underlying noise distribution (assuming only finite mean and variance). But their proof assumes that a weighted average is the correct combination rule (Oruç et al., 2003, Appendix A), which in Bayesian terms, assumes Gaussian noise. It is therefore unclear how general this proof really is as different distributions give different combination rules. This is explored in a navigational context below, but a separate example is given by Rosas and Wichmann (2011) who note that the Cauchy distribution has undefined mean and variance, which precludes the use of Equation 3.13. It is not clear why the linear/Bayesian distinction is made, beyond perhaps attributing the 'linear model' to Cochran (1937) who discussed it in terms of a maximum-likelihood estimate but without the Bayesian start point.

3.3.2 Bayesian behaviour

The result in Equation 3.13 has seen a great deal of success in human psychophysics. The idea that we would behave according to probability theory has obvious appeal. Nevertheless there have been plenty of examples of sub-optimality in human decision making, including cue integration tasks (Rahnev and Denison, 2018). What it means for behaviour to be optimal, sub-optimal, or Bayesian ends up being quite unclear. Equation 3.13 seems clearly 'Bayesian' but it only tells us what Bayes' theorem predicts under quite specific circumstances and assumes that we are able to validate those circumstances in our model system (Scarfe, 2022).

Despite its popularity, the behavioural model given by Equation 3.13 carries significant baggage which is often not addressed.

Uniform prior Equation 3.13 carries a built-in assumption of a uniform prior distribution over world states. It would be straightforward to include and test different priors to see how this changes the interpretation of an experiment. In cue integration contexts, it is not clear how often this assumption is considered or tested (Rahnev and Denison, 2018).

Independent noise Equation 3.13 also assumes that noise distributions are independent and this assumption is usually not verified (Scarfe, 2022) (but see (Oruç et al., 2003) for a counter-example).

Gaussian noise Continuing the theme, Equation 3.13 carries a built-in assumption of Gaussian noise distributions. This makes intuitive sense and I would consider this a fair assumption for the most part (although, like independence, this assumption is rarely given any scrutiny (Scarfe, 2022)). However, this can cause problems when we switch or combine different domains.

Some studies conflate distance (linear) and direction (angular) cues into ‘spatial’ cues which can make it difficult to interpret results. As an example Sjolund et al. (2018) use Equation 3.13 to model a ‘spatial’ cue conflict task. For small conflicts, subjects integrated cues as expected, however, for large cue conflicts, subjects tended to follow one cue or the other which is not explained by Equation 3.13. The authors verbally justify this by appealing to causal inference. To paraphrase Cheng et al. (2007), if I see a tiger’s stripes straight ahead and hear a growl to my right, I should infer that these cues came from two different tigers (what I am then meant to do with this inference is unclear). However, the appeal to causal inference may not be necessary. If we consider angular cues with independent von Mises noise (and again assume a uniform prior), then cue averaging at low conflicts and cue capture at large conflicts emerges as a consequence of Bayes’ theorem (Murray and Morgenstern, 2010) (see (Paper II)). (It is worth noting that this assumed von Mises noise distribution should also be justified.)

Assuming Gaussian noise is reasonable but the domain does require some consideration, especially in studies concerning navigation.

Estimating noise distributions The framework laid out above requires measuring noise distributions and doing this properly is only really practical in human test subjects. Noise distributions should be characterised on the individual level. This is because cue reliability will vary from person to person, and different people will have different perceptual biases. To determine if someone is behaving in a Bayesian fashion, it is critical to consider Bayesian behaviour *with respect to their sensor system*. This

typically requires performing a large number of trials with a single individual which is rarely practical in other animals. In humans, we can also set the context of the experiment. The participant knows that they are to give the height of the bar to the best of their ability, for example. We cannot be sure that a dung beetle is so singularly focused.

As an aside, work in cue integration (particularly in insects) can be very fuzzy about where error is actually coming from. For example, if we measure a behavioural response to a stimulus, such as dung beetle precision with respect to wind speed or solar elevation, it is not immediately clear how this adds noise to the system response. An orientation cue could have azimuthal variability, it could be difficult to detect, the sensor may not be very good, and the pathway from sensation to motor command could add error. Motor control could also be a confound; for example, dung beetles rock side-to-side when they roll balls and this almost certainly adds variability to the image of the sun on the retina (dung beetle eyes are fixed in the head and they do not appear to stabilise their heads). Potential sources for noise and the effects of these specific manipulations are discussed in ([Paper II](#)) (and the short Conclusion included in the thesis).

Implied goal It is then also important to consider the goal of the system. What is it trying to do? How do we define failure? The derivation laid out in the previous section implicitly assumes that the goal of the system is to produce the best possible estimate the stimulus X given two noisy sensor readings. However, precision can be sacrificed depending on the goal. For example, in humans, there is trade-off between response time and response accuracy; the faster the response, the less accurate the estimate will be.

That the goal affects behavioural outcome seems almost asinine to point out, but for non-human animals, we do not know what the relevant goal is. Take the dung beetle rolling its ball. It could be argued that precision is of the utmost importance if they want to avoid competition, but then speed could also be important. Perhaps speed is important during the initial stages of rolling behaviour and precision becomes more important later on. We do not know if there is a trade-off between rolling speed and rolling precision, but the idea that the system may have different goals during different segments of the same observable behaviour seems plausible. This has significant implications for what behaviour can be considered Bayesian or optimal; optimal with respect to what goal?

Meaningfulness The problem of the implied goal can be addressed by including concepts of action and cost and looking for the action which

generates the least cost with respect to the most likely state of the world (see (Rahnev and Denison, 2018)). However, this now means we define the goal of the system and describing behaviour as ‘according to Bayesian reasoning’ or ‘Bayes-optimal’ has become meaningless. If we can choose the goal then any behaviour can be regarded as Bayesian (or optimal).

3.3.3 Answering ‘Why?’ questions

The objections laid out above are nothing new. In particular, the implied goal and meaningfulness of ‘optimal’ and ‘Bayesian’ as descriptors of behaviour have seen a great deal of critical discussion (Bowers and Davis, 2012a,b; Colombo and Seriès, 2012; Rahnev and Denison, 2018) (also see (Scarfe, 2022)).

In responding to criticism, proponents argue that the idea that one can choose the cost function is a core feature of the approach (Griffiths et al., 2012; Rahnev and Denison, 2018). By choosing different cost functions, we can see which leads to the best description of behaviour which allows us to explain *why* an animal (typically a human) behaves in the way that it does (Griffiths et al., 2012; Kording, 2007). If we assume that an agent is behaving optimally with respect to some goal, then we can determine what the goal is (or is most likely to be).

This is reasonable but does not generally align with the reporting of results (at least in those papers which receive the most attention, e.g. Ernst and Banks (2002)). That a system performs Bayesian reasoning or performs in a Bayes-optimal fashion is generally treated as a key insight, not a starting assumption (Knill and Pouget, 2004) (also see (Rahnev and Denison, 2018)). The unfortunate side-effect of this method of reporting is that the reader completely misses the underlying point of using Bayes-optimality as a benchmark for behaviour: to explain *why* an animal behaves in the way it does.

3.3.4 How does this relate to insects?

The general impression one will glean from a brief survey of this literature is as follows. Animals integrate cues in order to minimise the variance of their perception of the world. They do so by averaging different cues, weighting each according to its relative reliability. This averaging procedure is in fact optimal, and suggests that animals behave according to Bayesian principles.

The amount of relevant detail omitted here should be obvious. This creates a paradoxical situation where ‘Bayesian behaviour’ both means something very specific (and requires certain conditions to be met (Scarfe,

2022)), yet could also describe any behaviour so long as we make the correct assumptions (Bowers and Davis, 2012a,b; Rahnev and Denison, 2018). The idea that there is a single Bayesian view on cue integration has appeared in previous publications in insects (Collett, 2012; Legge et al., 2014; Wystrach et al., 2015; Khaldy et al., 2022) and in those presented in this thesis (Paper I; Paper II). Frustratingly, this is a communication problem (Rahnev and Denison, 2018), rather than a fundamental issue with the actual analysis technique. This singular view of Bayesian behaviour also prevents us from considering whether or not this framework is useful in insects, particularly if we are interested in *how* tiny brains perform complex behaviours such as cue integration.

If we start with the 'general impression', then it seems difficult to properly test for so-called Bayesian or Bayes-optimal behaviour in insects. Estimation of cue noise distributions is not feasible on the individual level in insects. It can be done on the population level (e.g. (Paper I)) but it then becomes unclear what we can say with regards to the idea of Bayes-optimality. Even if we push this issue aside, the underlying assumption that insects are always trying to maximise the precision of their estimate may also be flawed. An ant trying to find its nest may well be trying to maximise the precision of its estimate (Wystrach et al., 2015; Hoinville and Wehner, 2018). In the case of a dung beetle one could argue that orientation cue redundancy improves robustness independently of any change in precision. This precision-focused idea of Bayes-optimality therefore seems an unhelpful benchmark. We could fall back on the more general 'why' question. But then we go from "How do insects integrate orientation cues?" to "Why do insects integrate orientation cues?" and the former question is still left without an answer.

Using Bayes' theorem to analyse behaviour at a population level is not useless. However, we need to be careful with established language. *Reliability* has a specific meaning and is assumed to be quantified on the individual level. Where we cannot do this and cannot make reasonable assumptions about goals, the typical notion of *optimality* is not a meaningful bar against which to compare behaviour. It is also worth highlighting that the narrow definition of Bayesian/optimal behaviour popularised by classic human studies such as (Ernst and Banks, 2002) and propagated by the subsequent psychophysical literature (Rahnev and Denison, 2018) does not need to be the bar we use to measure performance insects.

An animal may well weight cues according to their reliability in order to generate a minimum variance estimate (and this would certainly be interesting), however it is not the only lens through which to examine cue integration. Limiting ourselves to discussing cue integration as optimal or sub-optimal according to Equation 3.13 (or a domain-appropriate

equivalent) gives us a myopic view of the subject. We can of course still quantify population noise to see how this changes for different cue states (e.g. light elevations or wind speeds ([Paper I](#))). We can also see if this corresponds to cue weighting in any way and try to determine if this says anything about how the animal weighs cues. That said, where we are no longer bound by preconceived notions of what an animal ‘should’ do, we can also look for alternative explanations for cue weighting behaviour (e.g. ([Paper II](#))).

We cannot instruct a dung beetle in the nature of its task, we cannot be sure as to its motivation, and we cannot work with one beetle for hundreds of trials. But, we can look into its brain in ways which we cannot in humans. In insects more generally, we have an unparalleled ability to place behaviour into a neuroanatomical and physiological context. This can allow disambiguation of computational models which generate the same behavioural predictions (see earlier discussion on ([Khaldy et al., 2021](#))). Modern analyses are giving us a wealth of anatomical data which can allow us to propose and test models which are plausible given a putative substrate ([Scheffer et al., 2020](#); [Hulse et al., 2021](#)). Using this information, we can propose models based on behaviour which are constrained by known anatomy. Constructing a full anatomical model is not immediately necessary to prove the principle, an abstracted version could do (e.g. a vector sum versus the neural model given in ([Paper II](#))). Trying to find a simple, anatomically feasible computation which explains behaviour may not generate a nice theoretical result but it would be more meaningful than trying to determine whether we can express this as an example of Bayesian inference in the brain.

3.4 Summary

Bayes theorem’ has become the default lens through which to view cue integration problems. As a result, existing literature on insect cue integration has discussed behaviour in Bayesian terms. In spite of its popularity (and its merits), the approach has several problems, primarily the communication and verification of the assumptions underlying the most popular Bayesian model (Equation 3.13). The classic Bayesian model is hard to apply in insects as it is difficult to estimate the required quantities reliably (though see [Wystrach et al. \(2015\)](#)), and also difficult to verify assumptions with regards to goals. Nevertheless, applying Bayesian models to insect behaviour can still lead to valuable insight where we can draw parallels with neuroanatomy.

Chapter 4

Paper 1: Weighted cue integration for straight-line orientation

Introduction

This paper represents the logical first step in investigating orientation cue integration strategies for dung beetle straight-line orientation. Previous works have suggested both that dung beetles use a ‘winner-take-all’ (Dacke et al., 2019) strategy and that dung beetles weight cues in a ‘Bayesian’ fashion (Khaldy et al., 2022) (see Chapter 3). In both cases, it was argued that cue influence was determined by reliability, but neither work actually quantified single cue reliability in order to verify this claim ((Dacke et al., 2019) do reference earlier results from (Dacke et al., 2014) as justification but see Discussion in (Paper II)).

In this paper, we investigated the integration of wind and sun cues for orientation. We began by measuring orientation precision under single cue conditions which provides the best available proxy measurement for cue reliability. Then, we presented a controlled conflict experiment in which the wind is shifted with respect to the sun by either 60° or 120°. By using the single cue reliability measurements, we were able to characterise noise distributions (assumed von Mises) for each cue and simulate the conflict experiment. This allowed us to determine what a noisy population would do under different weighting/integration models for different degrees of conflict. We then compared five such candidate models to see which was most likely to generate our behavioural data. In summary, models which could be encoded as a vector sum were all more likely than the classic linear model (Equation 3.13). Moreover, cue weights were not determined by reliability; reliability measurements (as estimated) required significant adjustment to correspond to

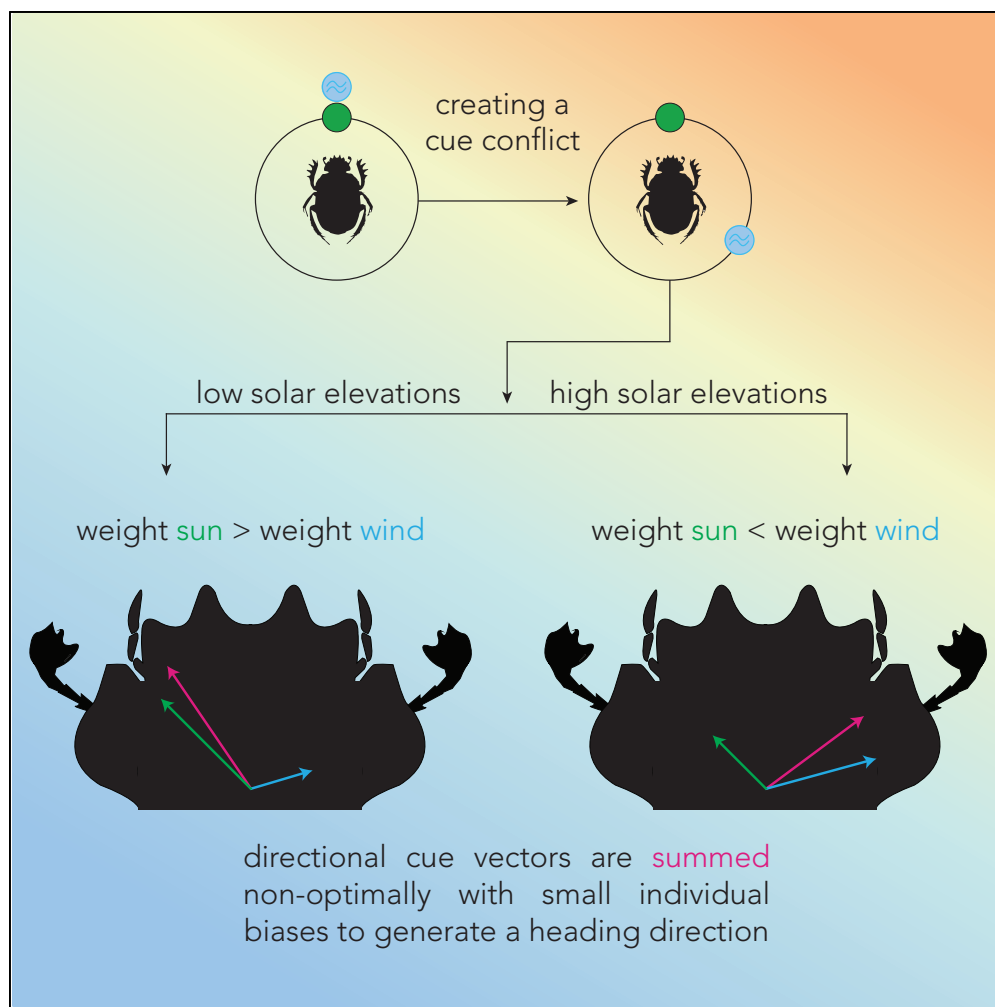
behaviourally observed cue weights.

Note This work was a collaboration with Shahrzad Shaverdian and Elin Dirlik, who collected and analysed the behavioural data. I have received their permission to include the article in my thesis as-is. iScience is an Open Access journal and the article is available under CC BY 4.0. The supplementary material can be found in Appendix C.

Author contributions Behavioral experiments, Shahrzad Shaverdian and Elin Dirlik. Modelling, Robert Mitchell. Writing - original draft, S.S., E.D., and R.M. Writing review and editing, all authors. Funding acquisition, Marie Dacke and Barbara Webb. Supervision, M.D. and B.W. (Given on page 9 of the article.)

Article

Weighted cue integration for straight-line orientation



Shahrazad Shaverdian, Elin Dirlik, Robert Mitchell, Claudia Tocco, Barbara Webb, Marie Dacke

elin.dirlik@biol.lu.se

Highlights

The weight of a sun cue for orientation decreases with increasing elevation

The weight of wind as an orientation cue increases with increasing speed

Behaviorally, dung beetles integrate cues using a vector summation strategy

The insect brain is well suited to encode vector summation

Shaverdian et al., iScience 25, 105207
October 21, 2022 © 2022 The Author(s).
<https://doi.org/10.1016/j.isci.2022.105207>



Article

Weighted cue integration
for straight-line orientationShahrazad Shaverdian,^{1,3} Elin Dirlik,^{1,3,4,*} Robert Mitchell,^{2,3} Claudia Tocco,¹ Barbara Webb,² and Marie Dacke¹

SUMMARY

Animals commonly integrate multiple sources of information to guide their behavior. Among insects, previous studies have suggested that the relative reliability of cues affects their weighting in behavior, but have not systematically explored how well alternative integration strategies can account for the observed directional choices. Here, we characterize the directional reliability of an ersatz sun at different elevations and wind at different speeds as guiding cues for a species of ball-rolling dung beetle. The relative reliability is then shown to determine which cue dominates when the cues are put in conflict. We further show through modeling that the results are best explained by continuous integration of the cues as a vector-sum (rather than switching between them) but with non-optimal weighting and small individual biases. The neural circuitry in the insect central complex appears to provide an ideal substrate for this type of vector-sum-based integration mechanism.

INTRODUCTION

Cue integration is the process whereby an agent combines multiple sensory estimates of the world to perform a single action (Ernst and Bühlhoff, 2004). Most explorations of cue integration stem from human psychophysical experiments (Ernst and Banks, 2002; Alais and Burr, 2004), but many examples of cue integration have been found throughout the animal kingdom from monkeys (Fetsch et al., 2012), to rodents (Knight et al., 2014), to insects (Legge et al., 2014; Wystrach et al., 2015; Hoinville and Wehner, 2018; Dreyer et al., 2018; Dacke et al., 2014, 2019; Müller and Wehner, 2007; Collett, 2012; Khaldy et al., 2021; Wehner et al., 2016).

While there is an abundance of evidence suggesting that insects integrate multiple cues when performing navigation behaviors, few propose concrete models which describe the integration process. Compass cue integration presents a direct case study for directional cue integration, and for this, the ball-rolling dung beetle *Kheper lamarcki* (MacLeay, 1821) provides an ideal model organism. Upon finding a dung pat, these beetles break off a piece of dung, shape it into a ball, climb on top of it, and rotate about their own vertical axis. During this “orientation dance”, a snapshot of all available cues is taken and then used to support a directed and efficient escape from the competition at the dung pile (Baird et al., 2012; Byrne et al., 2003; el Jundi et al., 2016). The natural environment provides a plethora of cues that are used by the beetles to sustain this straight-line orientation. Known orientation cues include the position of the sun (Byrne et al., 2003), moon (Dacke et al., 2004), spectral gradients (Jundi et al., 2015), the intensity pattern of the milky way (Dacke et al., 2013; Foster et al., 2017), and wind direction (Dacke et al., 2019). Previous studies have also shown that the beetles will interpret an artificial green light spot as an ersatz sun and will use it to orient with equal accuracy as under natural conditions (El Jundi et al., 2015).

Dacke et al. (2019) previously demonstrated that the influence of a wind cue is dependent on the elevation of the accompanying sun cue. Sun cues above an elevation threshold were seemingly ignored in the presence of a wind cue. This appears to be consistent with the notion that cues are integrated with weights determined by their reliability; cues which are subject to greater perceptual noise (e.g. the azimuth of the sun at a high elevation) are less reliable. Indeed, statistically optimal cue integration models, which seek to maximize the reliability of the combined cue, predict cue weights which are directly proportional to their respective reliabilities (Ernst and Bühlhoff, 2004; Murray and Morgenstern, 2010).

Here, we extend this line of work by providing an in-depth exploration of multimodal (ersatz sun and wind) cue integration in the dung beetle compass with an indoor setup which completely isolates the cues under

¹Lund Vision Group, Department of Biology, Lund University, 223 62 Lund, Sweden

²School of Informatics, The University of Edinburgh, Edinburgh EH8 9AB, UK

³These authors contributed equally

⁴Lead contact

*Correspondence: elin.dirlik@biol.lu.se

<https://doi.org/10.1016/j.isci.2022.105207>






Figure 1. Mean menotactic orientation precision at each elevation and wind speed. Error bars indicate standard error of the mean.

The mean precision at 60° and 75° elevation is highlighted to allow comparison to the wind (dashed lines). The mean precision at 60° is close to that at 2.5 m/s wind speed, and similarly the mean precision at 75° is less than that at 2.5 m/s wind speed but much greater than would be expected at 1.25 m/s. This relationship matches our cue conflict results. The line fits are those used in Equation 2, excluding the additive constants (which are only required for the κ estimation stage, see STAR Methods, Simulated cue representation). Circular insets illustrate ten paths traveled by a highly directed (left, R-value = 0.91) and a weakly directed (right, R-value = 0.21) beetle.

study. Our approach combines behavioral data with computational modeling to describe the cue integration method employed by the South African ball-rolling dung beetle *K. lamarcki*. Our descriptive modeling points to a vector-sum-based integration mechanism, for which the neural circuitry in the insect central complex appears to provide an ideal substrate.

RESULTS AND DISCUSSION

Reliability of ersatz sun and wind cues

Dung beetle compass orientation is characterized by menotaxis with respect to a given cue-oriented behavior but without a specific directional preference. As a measure of cue reliability, we analyzed beetles' orientation precision at solar elevations and wind speeds where menotactic behavior was observed with respect to the isolated cue. Wind cues are provided by custom-built wind generators and a solar cue is provided by a green LED (ersatz sun; see STAR Methods, Experimental setup). Cue reliability (variability from the perspective of the animal) is not something we can analyze directly. We therefore use orientation precision as a proxy (see STAR Methods, Simulated cue representation). It is important to note that precision changes could be brought about by changes in a cue's usefulness, i.e. how good of an orientation cue it actually is from the perspective of the animal, as opposed to its variability. Orientation precision was calculated from ten exit bearings for each beetle and reported as mean vector lengths; the closer to 1 the more clustered the exits, where more clustered exits are interpreted as higher precision (R-values; see STAR Methods, Quantification and statistical analysis and Figure 1). To show menotactic behavior statistically, we use Rayleigh tests, where a significant result indicates a directed population (i.e. a preferred direction for all beetles with respect to the cue) and a uniform distribution indicates menotaxis (i.e. arbitrary directions selected by each beetle). These Rayleigh tests were performed on the mean directions taken by beetles where each mean was calculated from the aforementioned ten exit bearings.

Orientation to the ersatz sun at different elevations

The bearings traveled across the population of beetles at solar elevations of 5°, 20°, 45°, 60°, 80°, 82°, 84°, 86°, and 88° did not differ from a uniform distribution ($p = 0.201$ ($n = 18$), 0.756 ($n = 17$), 0.534 ($n = 16$), 0.477 ($n = 17$), 0.772 ($n = 10$), 0.237 ($n = 15$), 0.709 ($n = 9$), 0.184 ($n = 6$), and 0.167 ($n = 4$), respectively, Rayleigh tests, see Table S1). Similar results have previously been found in dung beetles orienting under natural conditions (Baird et al., 2010). In contrast, at 75°, the bearings traveled by the beetles were significantly different from uniformity ($p = 0.02$, Rayleigh test, $n = 13$); but as this bearing preference was not present at either lower or higher solar elevations, we decided to treat it as an outlier and proceeded to perform orientation precision analysis for all solar elevations.

The beetles' ability to maintain their bearings remained stable until an elevation of 75°, beyond which it decreased rapidly with increasing elevation (5° = 0.78 [0.71, 0.84], 20° = 0.78 [0.69, 0.86], 45° = 0.74 [0.64, 0.84], 60° = 0.70 [0.61, 0.88], 75° = 0.69 [0.48, 0.80], 80° = 0.55 [0.40, 0.70], 82° = 0.61 [0.40, 0.66],

$84^\circ = 0.50$ [0.34, 0.67], $86^\circ = 0.41$ [0.27, 0.60], $88^\circ = 0.33$ [0.19, 0.48], and $90^\circ = 0.37$ [0.22, 0.49], median R -values [IQR], $n = 20$, see [Figure 1](#)). These results suggest that higher solar elevations provide a less reliable cue, which can be attributed to the decrease in directional information given by a visual cue as it approaches zenith. It should be noted that direct comparisons between the indoor setup and outdoor conditions may not be applicable as the light intensity is lower in the artificial setup. Despite this, a similar decrease in orientation precision at high solar elevations can also be observed outdoors for the same species of dung beetles ([Dacke et al., 2014](#)), as well as for other animals (for example, equatorial sandhoppers and desert ants ([Ugolini, 2001, 2002](#); [Müller and Wehner, 2007](#))).

Orientation to the wind at different speeds

Next, we analyzed the population bearing preference in relationship to the wind cue at different wind speeds. At 0.5, 0.8, 3.0, and 4.0 m/s , the mean bearings differed significantly from a uniform distribution ($p < 0.05$, Rayleigh tests, $n = 10, 9, 13$, and 14 for wind speeds 0.5, 0.8, 3.0, and 4.0, respectively, see [Table S1](#)), suggesting anemotaxis. At 1.0, 1.5, 1.9, and 2.5 m/s , the mean bearings did not differ significantly from a uniform distribution ($p = 0.108$ ($n = 7$), 0.611 ($n = 14$), 0.734 ($n = 14$), and 0.865 ($n = 14$), respectively, Rayleigh tests, see [Table S1](#)), suggesting menotaxis.

Under the conditions where menotaxis was observed, we found that as the speed increased, so did the beetles' ability to maintain their bearings ($1.0 = 0.47$ [0.36, 0.72], $1.5 = 0.61$ [0.45, 0.80], $1.9 = 0.63$ [0.54, 0.81], and $2.5 = 0.68$ [0.51, 0.84], median R -values [IQR], $n = 20$, see [Figure 1](#)). Taken together, the results suggest that higher wind speeds provide a more reliable cue for menotaxis (changing to anemotaxis beyond 3.0 m/s). Previous studies have demonstrated that the antennae of dung beetles are wind sensitive ([Linsenmair, 1972](#)), and [Okubo et al. \(2020\)](#) have shown that with increasing wind speed, the antennae of fruit flies are subject to larger displacements. Our results are therefore in line with previous research as one would expect greater deflections of the dung beetle antennae to provide a clearer perception of wind direction. Note the similarity in reliability between a wind cue of 2.5 m/s and an ersatz sun at a 60° elevation (see [Figure 1](#)).

Cue conflict between an ersatz sun and wind

The effect of reliability on the integration and weighting of a visual sun cue and a mechanosensory wind cue was studied in a cue conflict experiment. The reliability was manipulated by changing the elevation of the ersatz sun or the speed of the wind current and the conflict was introduced by shifting the azimuthal direction of the wind (see [STAR Methods, Behavioural experiments - Cue conflict between an ersatz sun and wind](#)). Changes in heading direction were calculated using the angular difference between two consecutive exits (see [Figure 2](#)). All beetles included in the analysis were able to recover their initial bearing following each cue conflict run (see [STAR Methods, Quantification and statistical analysis - Cue conflict between an ersatz sun and wind](#)).

Dung beetles perform a weighted integration of a wind cue and sun cue

As can be expected from previous studies outdoors ([Dacke et al., 2019](#)), when presented with a wind cue of 2.5 m/s and a simulated solar elevation of 45° , 60° , 75° , or 86° , beetles were able to maintain their bearing between two consecutive exits when the directional information from the two cues remained unchanged ($\mu \pm SD$: $-3^\circ \pm 28^\circ$ ($n = 30$), $-11^\circ \pm 28^\circ$ ($n = 27$), $-5^\circ \pm 45^\circ$ ($n = 26$), and $3^\circ \pm 47^\circ$ ($n = 23$); $p < 0.001$, Rayleigh tests, see [Figure 2B](#)). However, when the directional information from the sun and wind cues were put in conflict, by altering the wind direction between two consecutive exits, the behavior changed depending on solar elevation. At a solar elevation of 45° and the wind current set to 2.5 m/s , the beetles did not change their direction at either a 60° or a 120° cue conflict ($\mu \pm SD$: $-9^\circ \pm 43^\circ$ and $16^\circ \pm 68^\circ$; $p < 0.001$, Rayleigh tests, $n = 30$, see [Figure 2B](#)). This suggests that at this elevation, the ersatz sun has a greater weight compared to the wind cue. In contrast, for solar elevations of 75° and 86° , the beetles updated their bearing in accordance with the 60° or 120° azimuthal change of the wind current ($\mu \pm SD$: $78^\circ \pm 56^\circ$ and $116^\circ \pm 60^\circ$; $p < 0.001$ at 75° elevation ($n = 26$), $74^\circ \pm 80^\circ$ and $123^\circ \pm 62^\circ$; $p = 0.025$ and $p < 0.001$ at 86° elevation ($n = 23$), Rayleigh tests, see [Figure 2B](#)). At a 60° solar elevation, the beetles again responded to the 60° azimuthal shift of the wind ($\mu \pm SD$: $55^\circ \pm 76^\circ$; $p < 0.01$, Rayleigh test, $n = 27$); interestingly, when presented with a conflict of 120° , the changes in bearing did not differ significantly from a uniform distribution ($p = 0.837$, Rayleigh test, $n = 27$). Thus, the beetles did not seem to keep their relative bearing to either the ersatz sun or the wind.

Together, our results suggest that the relative weight between the sun cue and the wind cue is affected by elevation and that the critical elevation (at which the ersatz sun becomes less reliable than a wind cue of 2.5

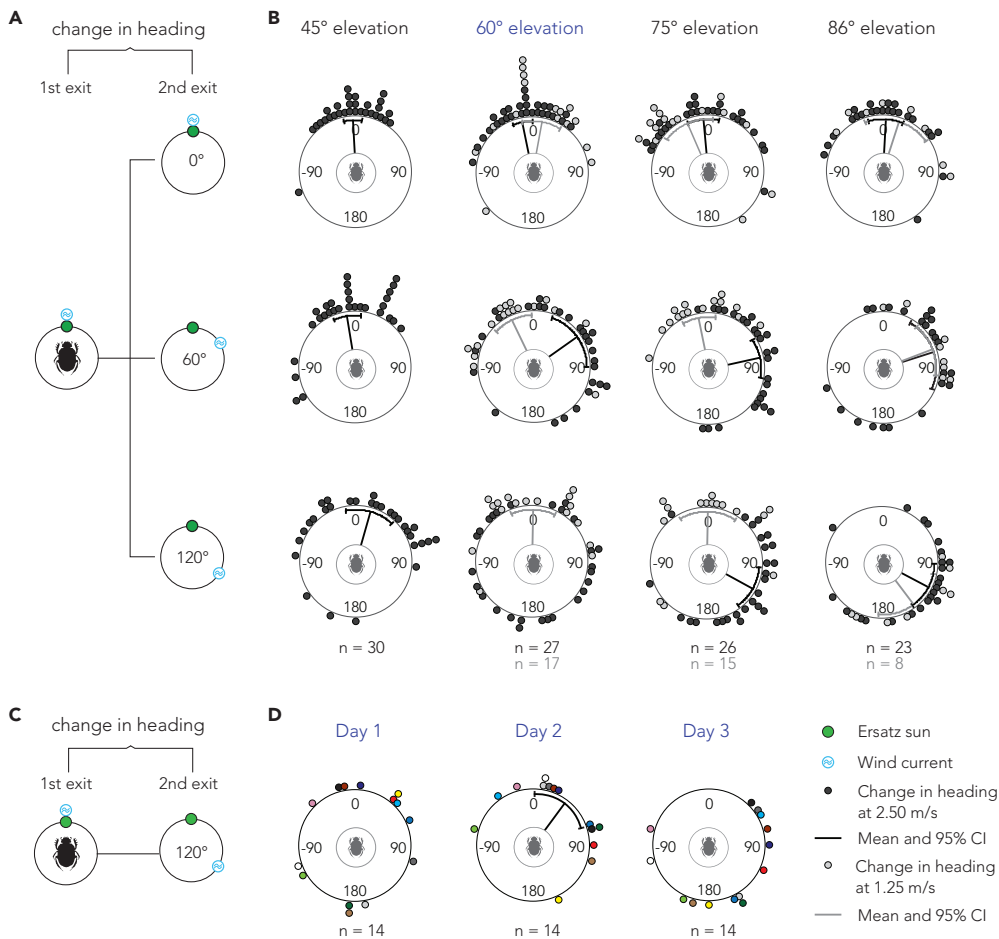


Figure 2. Behavioral results from a cue conflict experiment

(A) Schematic procedure of the cue conflict experiment. Change in heading was calculated for individual beetles between two consecutive exits; initial condition (1st exit where the initial bearing is established) to conflict condition (2nd exit where the wind had changed direction by 0°, 60°, or 120° relative to the ersatz sun).

(B) The changes in headings at wind speed 2.5 m/s are illustrated as black circles and at 1.25 m/s as gray circles. Lines extending from the centers indicate mean vectors, black lines for 2.5 m/s and gray lines for 1.25 m/s, and end in a 95% confidence interval of the spread.

(C) Schematic procedure of the experiment where a 2.5 m/s wind cue was subjected to a 120° azimuthal shift in the presence of a sun cue at 60° elevation.

(D) The changes in headings at three different days. Each colored data point illustrates the change in heading of the same individual across days.

m/s) lies between 45° and 75°. This result matches previous observations in dung beetles (Dacke et al., 2019) and ants (Müller and Wehner, 2007); the higher the solar elevation, the lower the influence of the cue. From our results, it may appear as if the beetles are simply following the more reliable cue, which has previously been a suggested orientation strategy among dung beetles when presented with sun and wind (Dacke et al., 2019), sun and polarized light (Jundi et al., 2014), or sun and other skylight cues (Dacke et al., 2014). However, considering the uniform distribution found at a solar elevation of 60° together with a 120° cue conflict, the beetles are not always able to follow one cue over the other. This could indicate that at this elevation, the cue reliabilities intersect. Furthermore, the population spread increases with conflict; this effect is consistent across almost all test conditions when the wind speed is set to 2.5 m/s. This suggests that the beetles are not following a simple winner-take-all strategy, as under strict winner-take-all (see STAR Methods, Integration models) the population dispersion should be unaffected by cue conflict. The pseudo winner-take-all behavior and increasing variance that we observe at elevations of 45°, 75°, and 86° could be explained by a circular integration model (Murray and Morgenstern, 2010) with

non-optimal weights (see [STAR Methods, Integration models](#)), suggesting both cues are contributing to behavior, even when one appears to be followed and the other ignored.

Orientation behavior varies when cue weights are similar

Due to the random distribution of beetles' changes in heading when the 60° elevation sun cue and the 2.5 m/s wind cue were put in conflict by 120°, it appeared that the population failed to orient under this condition. However, an additional experiment that focused on individual precision, in which each beetle was permitted to exit the arena ten times in the presence of the cue conflict conditions, showed that the beetles did not fail to orient. Instead, we found that the beetles oriented along a new random bearing that they then successfully maintained (0.91 [0.86, 0.95], median *R*-value [IQR], *n* = 11). Beetles were also able to maintain their bearings at conflicts of 0° and 60° (0.76 [0.66, 0.85] and 0.89 [0.82, 0.91], respectively, median *R*-value [IQR], *n* = 11).

Furthermore, when the two cues were returned to their original positions, the beetles recovered their initial bearing (see [STAR Methods, Quantification and statistical analysis - Cue conflict between an ersatz sun and wind](#)), suggesting that this new random bearing we see is actually an effect of the integration strategy and not a permanent re-set of the bearing. Similarly, [Khaldy et al. \(2021\)](#) showed that the ball-rolling dung beetles *Garreta unicolor* and *Garreta nitens* appeared disoriented when subjected to a conflict produced by simultaneous manipulation of a sun cue and the pattern of polarized light. However, upon returning the cues to their original positions, these animals recovered their initial bearings. Likewise, bogong moths fly in a seemingly disoriented manner when presented with a conflict between the magnetic field and visual landmarks; when the cues are returned to their original positions the moths, too, recover their initial bearings ([Dreyer et al., 2018](#)).

To explore these apparent new bearings taken by beetles at a 120° conflict, we tested whether the individual change in bearing was consistent over different days or if it was prone to change. We employed the previously described cue conflict assay, focusing on a 60° solar elevation with a wind speed of 2.5 m/s and tested individual beetles over three consecutive days. Similar to our previous cue conflict experiments, the beetles were able to maintain their bearings each day when the cues were kept in their original positions ($\mu \pm SD$: $-1^\circ \pm 23^\circ$, $9^\circ \pm 41^\circ$, and $-4^\circ \pm 44^\circ$ for day one, two, and three, respectively; $p < 0.001$, Rayleigh tests, *n* = 14, see [Figure S3](#)). However, upon changing the azimuthal position of the wind by 120°, the individual change in bearing differed across the three days (*n* = 14, see [Figure 2D](#)). This shows that the apparent new bearing taken by beetles at a 60° solar elevation is not consistent over days.

The results from this three day experiment reinforce the idea that a weighted integration is taking place as both cues must be considered to generate the variability we see in the mean vector (see [Figure 2D](#)); if this were a winner-take-all (or biased winner-take-all), we would not expect to see any beetles in the lower left-hand quadrant. Furthermore, the inconsistency in population response suggests an additional source of noise in the integration process. A potential explanation for this noise would be individual variation or "preference". A very small (random) individual bias could cause an increase in spread in the population of responses where cue weights are near equal (see [STAR Methods, Integration models; Results and Discussion, Modeling](#)).

The weight given to a sun cue and a wind cue is dictated by their relative reliability

In previous experiments, the reliability of the ersatz sun was altered. Here, the wind speed was reduced to 1.25 m/s to study the effect of decreased wind reliability on the relative weighting of the sun and wind cues. We again employed the previously described cue conflict assay at solar elevations 60°, 75°, and 86°.

At solar elevations of 60° and 75°, the directional changes of the wind current had no effect on the beetles' orientation behavior, regardless of conflict angle ($\mu \pm SD$: for conflicts 0°, 60°, and 120°: $12^\circ \pm 46^\circ$, $-25^\circ \pm 46^\circ$, and $1^\circ \pm 51^\circ$ for 60° elevation (*n* = 17), $-24^\circ \pm 55^\circ$, $-11^\circ \pm 33^\circ$, and $1^\circ \pm 58^\circ$ for 75° elevation (*n* = 15); $p < 0.01$, Rayleigh tests, see [Figure 2](#)). However, at an 86° solar elevation, the beetles changed their bearings in accordance to the azimuthal shift of the wind ($\mu \pm SD$: $17^\circ \pm 51^\circ$, $68^\circ \pm 35^\circ$, and $143^\circ \pm 48^\circ$; $p = 0.023$, $p < 0.01$ and $p = 0.013$, with increasing conflict, Rayleigh tests, *n* = 8, see [Figure 2](#)). These findings stand in contrast to the behavior observed in the presence of a 2.5 m/s wind current at the same solar elevations (recall, at a wind speed of 2.5 m/s and solar elevation of 75° the beetles attributed a higher

relative weight to the wind). These results further reinforce the conclusion that relative reliability dictates the weight given to each individual cue.

Behavior indicating a weighted integration strategy

In all, our behavioral results show that the dung beetle compass is dynamic and that relative cue reliability dictates which cue is favored; this occurs as a continuous integration process, rather than a winner-take-all. The cue which is perceived to be less noisy, and thus more reliable, is predominantly used to guide straight-line orientation. This holds true until the relative reliabilities are similar and the beetles initially appear to be unable to utilize the provided cues for orientation. However, when investigating individual precision, we found that beetles were able to maintain their apparent new bearing, as well as recover their initial bearing when the cues were returned to their original positions. Furthermore, individual orientation behavior differed across three days. Together, these results suggest that the observed randomness is an effect of a weighted integration strategy, but the integration may be inconsistent across days. To attempt to characterize the integration strategy, we performed simulations to experiment with different strategies and weight relationships.

Modeling

Cue integration studies typically compare winner-take-all (WTA) to “optimal” cue integration, defined as the linear weighted arithmetic mean, WAM (Ernst and Banks, 2002; Ernst and Bühlhoff, 2004), or the circular weighted vector sum, WVS (Murray and Morgenstern, 2010). Previous work on dung beetles has assumed WTA (Dacke et al., 2019) but, as already noted, several aspects of our current results do not appear to match what this model would predict. Some studies on ant behavior have been taken to demonstrate WAM (Legge et al., 2014; Sun et al., 2018; Wystrach et al., 2015), but WAM is inappropriate for circular inputs. Instead, cue integration in the circular domain can be represented by a weighted vector sum (STAR Methods, Integration models), for which the optimal weights (which minimize the variance of the combined cue) are given by the concentrations of the von Mises noise distributions which characterize each cue. WVS has been used in a model of ant navigation (Hoinville and Wehner, 2018), and has the interesting property that it can resemble WAM at small conflicts and WTA at large conflicts (see Figure S6). This has inspired us to consider two further alternatives. The first is a “non-optimal weighted” vector sum (NVS) which exaggerates the pseudo-WTA region of the WVS, such that the stronger cue usually dominates the response but both contribute, unlike true WTA. The second is a biased (non-optimal weighted) vector sum (BVS), which simulates small individual biases toward one or the other cue, creating a variety of behavior when the cues are near-equally balanced. Illustrative model outputs for the same input distributions are given in Figure 3. For completeness, we compare all five models (WTA, WAM, WVS, NVS, and BVS, defined explicitly in STAR Methods, Integration models) to the behavioral data to calculate their relative likelihoods. The results are given in Table 1.

Our modeling results indicate that all circular models (WVS, NVS, and BVS) better account for our behavioral data than the classically considered WAM and WTA. The extremely similar outcomes between WTA and WVS were unexpected, but it is likely that these models both account for different portions of the data (i.e. they are equally bad at capturing the full range of behavior). The non-optimal circular model (NVS), which takes advantage of the pseudo-WTA property of a circular integration model (a vector sum, see STAR Methods, Integration models), performs substantially better than either optimal circular integration or a winner-take-all as this model should capture the small influence of the secondary cue. Finally, the inclusion of individual bias in BVS increases the population-level noise where cue weights are near equal, which should capture the behavioral variability at the critical elevation conditions (60° elevation, 2.5 m/s wind speed), leading to the best overall fit.

The key take-away from our simulations is that a weighted circular model best accounts for the data. WVS, NVS, and BVS make different behavioral predictions. WVS predicts mostly intermediate courses with the more reliable cue dominating at large conflicts. NVS predicts the more reliable cue will dominate most of the time with intermediate courses where cue weights are similar. Finally, BVS predicts mostly dominance, with some intermediate courses where cue weights are similar, and increased population variance where they are near equal. Despite their differences, these three models use the same basic vector sum (WTA can also be represented as an extreme case of NVS with very large a , see STAR Methods, Integration models). They differ only in how their weights are computed.

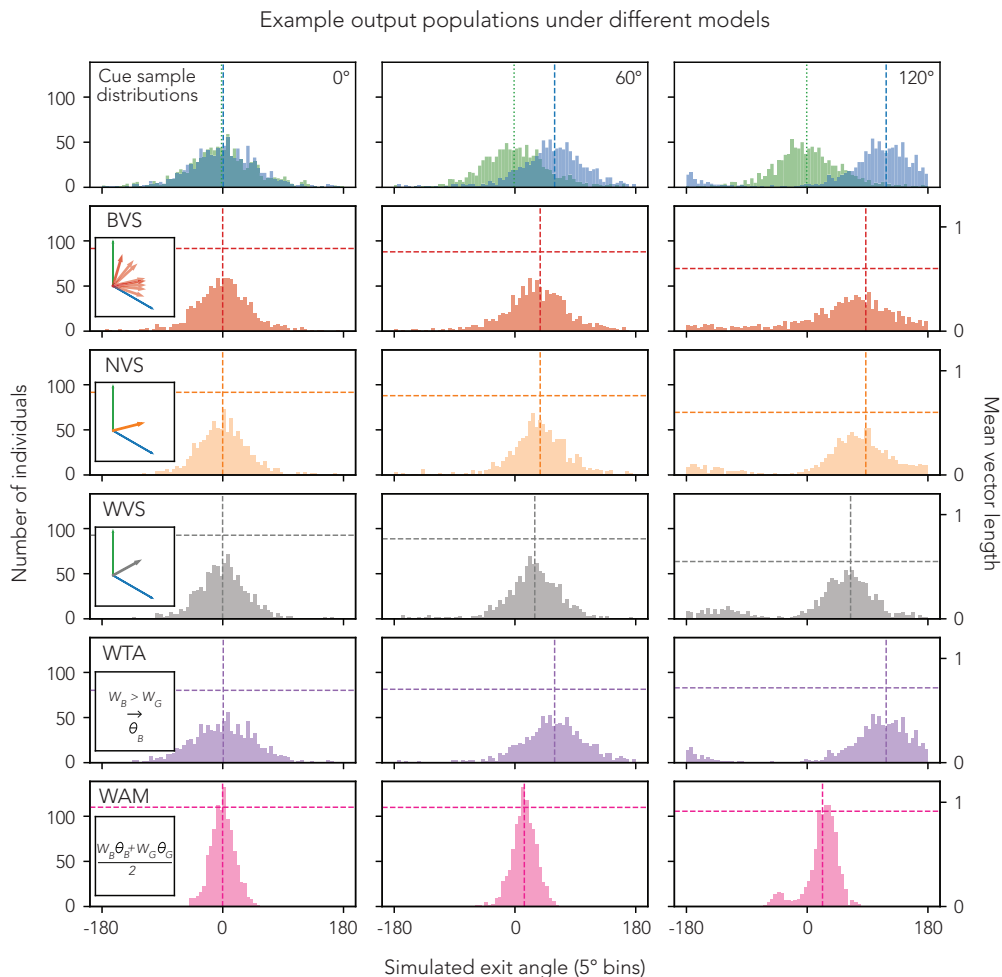


Figure 3. An illustration of the different model outcomes for two sets of input samples

Top row: The two sample distributions used as input to the models. Each cue is described by a noise distribution which is sampled to generate behavior. The noise distributions are von Mises with $\kappa_{Blue} = 2.05$, $\kappa_{Green} = 2$, $\mu_{Green} = 0$, and $\mu_{Blue} \in \{0^\circ, 60^\circ, 120^\circ\}$ (columns). Weights are computed from κ . **BVS** (Biased Vector Sum): Noise is added to the weights (vector magnitudes) which are then passed through an adjustment function. This strategy can generate different outputs for the same inputs due to the added noise. **NVS** (Non-optimal Vector Sum): Weights are adjusted and then the vectors are summed. **WVS** (Weighted Vector Sum): Angular samples are converted to vectors and then summed. **WTA** (Winner-take-all): Weights are compared and the cue with the greatest weight wins complete influence. **WAM** (Weighted Arithmetic Mean): A simple weighted average of the angles. For model definitions, please see [STAR Methods](#), [Integration models](#).

To tie our results to physiology, the neural circuitry in the insect central complex appears to provide an ideal substrate for performing a vector sum calculation. In the insect brain, head direction is maintained by a ring-attractor circuit (Seelig and Jayaraman, 2015; Kim et al., 2017; Heinze, 2017). This circuit is fed by the ring neurons, which seem to cluster into groups which are sensitive to different orientation cues (Okubo et al., 2020; Hardcastle et al., 2021). The two layers are linked by plastic all-to-all connections—every ring neuron connects to every compass (head direction) neuron (Kim et al., 2019; Fisher et al., 2019)—which should allow the network to learn relationships between different cues, forming a single integrated snapshot for orientation. In beetles, the relationship between the different available cues could be learned during the dung beetle “dance” (Baird et al., 2012), which is thought to be the point at which their orientation snapshot is taken (el Jundi et al., 2016). Similar neural models of angular cue integration have been proposed by Page et al. (2014) and Jeffery et al. (2016) in rodents, and Sun et al. (2018) in insects. Interestingly, Sun et al. (2018) implement WAM within a ring-attractor model and observe a strategy switch (where the output switches from WAM to WTA), which looks strikingly similar to the integration curves shown by Murray and Morgenstern (2010) (Figure S6, Right). Working with rodents, Knight et al. (2014) also note a

Table 1. Cross-model comparison results in order of likelihood/fit

Model	Parameters	Log likelihood ratio	AIC ratio	BIC ratio
BVS	2	0.00000	1.000000	1.000000
NVS	1	-17.21685	1.007461	1.006451
WVS	0	-50.45796	1.022295	1.020260
WTA	0	-50.73313	1.022421	1.020386
WAM	0	-199.68191	1.090951	1.088779

All circular models (WVS, NVS, and BVS) outperform the weighted arithmetic mean (WAM). Non-optimal weighted circular models (NVS and BVS) do better than both the statistically optimal WVS and previously hypothesized WTA. As parameter counts are small, the Akaike information criterion (AIC) and Bayesian information criterion (BIC) do not sufficiently penalize either BVS or NVS to affect the order of the results. Likelihoods are best at 0 and decrease as models become less likely. The AIC and BIC ratios are best at 1 and increase as goodness-of-fit decreases.

“switch” between apparently optimal integration behavior and cue selection behavior and finally, [Sjolund et al. \(2018\)](#) note similar results for spatial (distance and direction) cue integration in humans. Not only is a weighted circular model a likely candidate for cue integration in our case and in ants ([Hoinville and Wehner, 2018](#)), but the orientation center of the insect brain would seem to be well suited to encode the underlying vector sum. Overall behavior is then governed by the weights used, which need not be consistent across different species.

In summary, our behavioral data point to a weighted integration of wind and solar cues. Subsequent computer modeling suggests that the integration is most likely a form of vector summation, which seems to be well encoded by the insect head direction circuit. Vector summation can produce a variety of different integration outcomes depending on the weights used; a neural circuit which supports vector summation could produce different behavior depending on how an agent computes these weights. Thus, a single core model (vector summation) with different peripheral processing stages (weight-adjustment and/or bias etc.) may explain a wide range of cue integration behaviors across different insect species, despite the highly conserved neuroanatomy.

Limitations of the study

To fully isolate the orientation cues in question (sun and wind cues), we performed our experiments in an indoor setup which allowed us to control all possible cue parameters. Consequently, the indoor setup is limited in its representation of the real world. One major constraint of our behavioral setup is motion parallax, which becomes more severe at higher solar elevations.

The modeling is based on beetle exit angles which do not fully characterize the strategy in use by the beetles (especially with the aforementioned motion parallax). Modeling based on full tracks for each individual would be more informative but these data are not practically available. This means that applications of the specific model instances presented are limited.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- [KEY RESOURCES TABLE](#)
- [RESOURCE AVAILABILITY](#)
 - Lead contact
 - Materials availability
 - Data and code availability
- [EXPERIMENTAL MODEL AND SUBJECT DETAILS](#)
 - Ethical statement
 - Animal collection and experimental sites
- [METHOD DETAILS](#)
 - Experimental setup
 - Behavioural experiments
 - Simulation overview

- Simulated cue representation
- Integration models
- Evaluation process
- **QUANTIFICATION AND STATISTICAL ANALYSIS**
- Reliability of sun and wind cues
- Cue conflict between an ersatz sun and wind

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.isci.2022.105207>.

ACKNOWLEDGMENTS

We thank Lars Fredriksson and Nils Sundqvist for constructing our indoor behavioral setup and Yakir Gagnon for making it all work. This work was supported by the Swedish Research Council (2020-04046, MD), the European Research Council (ERC, 817535-Ultimate-COMPASS, MD), and the UK Engineering and Physical Sciences Research Council (EPSRC, DTA UoE 2019).

AUTHOR CONTRIBUTIONS

Behavioral experiments, S.S. and E.D. Modeling, R.M. Writing - original draft, S.S., E.D., and R.M. Writing - review and editing, all authors. Funding acquisition, M.D. and B.W. Supervision, M.D. and B.W.

DECLARATION OF INTERESTS

The authors declare no competing interests.

INCLUSION AND DIVERSITY

We support inclusive, diverse, and equitable conduct of research.

Received: July 12, 2022

Revised: September 5, 2022

Accepted: September 22, 2022

Published: October 21, 2022

REFERENCES

- Alais, D., and Burr, D. (2004). The ventriloquist effect results from near-optimal bimodal integration. *Curr. Biol.* 14, 257–262. <https://doi.org/10.1016/j.cub.2004.01.029>.
- Baird, E., Byrne, M.J., Scholtz, C.H., Warrant, E.J., and Dacke, M. (2010). Bearing selection in ball-rolling dung beetles: is it constant? *J. Comp. Physiol. A Neuroethol. Sens. Neural. Behav. Physiol.* 196, 801–806. <https://doi.org/10.1007/s00359-010-0559-8>.
- Baird, E., Byrne, M.J., Smolka, J., Warrant, E.J., and Dacke, M. (2012). The dung beetle dance: an orientation behaviour? *PLoS One* 7, e30211. <https://doi.org/10.1371/journal.pone.0030211>.
- Batschelet, E. (1981). *Circular Statistics in Biology* (Academic Press Inc.).
- Byrne, M., Dacke, M., Nordström, P., Scholtz, C., and Warrant, E. (2003). Visual cues used by ball-rolling dung beetles for orientation. *J. Comp. Physiol. A Neuroethol. Sens. Neural. Behav. Physiol.* 189, 411–418. <https://doi.org/10.1007/s00359-003-0415-1>.
- Cheng, K., Shettleworth, S.J., Huttenlocher, J., and Rieser, J.J. (2007). Bayesian integration of spatial information. *Psychol. Bull.* 133, 625–637. <https://doi.org/10.1037/0033-2909.133.4.625>.
- Collett, M. (2012). How navigational guidance systems are combined in a desert ant. *Curr. Biol.* 22, 927–932.
- Dacke, M., Baird, E., Byrne, M., Scholtz, C.H., and Warrant, E.J. (2013). Dung beetles use the milky way for orientation. *Curr. Biol.* 23, 298–300. <https://doi.org/10.1016/j.cub.2012.12.034>.
- Dacke, M., Bell, A.T.A., Foster, J.J., Baird, E.J., Strube-Bloss, M.F., Byrne, M.J., and El Jundi, B. (2019). Multimodal cue integration in the dung beetle compass. *Proc. Natl. Acad. Sci. USA* 116, 14248–14253. <https://doi.org/10.1073/pnas.1904308116>.
- Dacke, M., Byrne, M.J., Scholtz, C.H., and Warrant, E.J. (2004). Lunar orientation in a beetle. *Proc. Biol. Sci.* 271, 361–365. <https://doi.org/10.1098/rspb.2003.2594>.
- Dacke, M., el Jundi, B., Smolka, J., Byrne, M., and Baird, E. (2014). The role of the sun in the celestial compass of dung beetles. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 369, 20130036. <https://doi.org/10.1098/rstb.2013.0036>.
- Dreyer, D., Frost, B., Mouritsen, H., Günther, A., Green, K., Whitehouse, M., Johnsen, S., Heinze, S., and Warrant, E. (2018). The earth's magnetic field and visual landmarks steer migratory flight behavior in the nocturnal Australian bogong moth. *Curr. Biol.* 28, 2160–2166.e5. <https://doi.org/10.1016/j.cub.2018.05.030>.
- el Jundi, B., Foster, J.J.J., Khaldy, L., Byrne, M.J.J., Dacke, M., and Baird, E. (2016). A snapshot-based mechanism for celestial orientation. *Curr. Biol.* 26, 1456–1462. <https://doi.org/10.1016/j.cub.2016.03.030>.
- El Jundi, B., Warrant, E.J., Byrne, M.J., Khaldy, L., Baird, E., Smolka, J., and Dacke, M. (2015). Neural coding underlying the cue preference for celestial orientation. *Proc. Natl. Acad. Sci. USA* 112, 11395–11400. <https://doi.org/10.1073/pnas.1501272112>.
- Ernst, M.O., and Banks, M.S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* 415, 429–433.
- Ernst, M.O., and Bühlhoff, H.H. (2004). Merging the senses into a robust percept. *Trends Cognit. Sci.* 8, 162–169. <https://doi.org/10.1016/j.tics.2004.02.002>.
- Fetsch, C.R., Pouget, A., Deangelis, G.C., and Angelaki, D.E. (2012). Neural correlates of reliability-based cue weighting during multisensory integration. *Nat. Neurosci.* 15, 146–154. <https://doi.org/10.1038/nn.2983>.

- Fisher, Y.E., Lu, J., D'Alessandro, I., and Wilson, R.I. (2019). Sensorimotor experience remaps visual input to a heading-direction network. *Nature* 576, 121–125. <https://doi.org/10.1038/s41586-019-1772-4>.
- Foster, J.J., el Jundi, B., Smolka, J., Khaldy, L., Nilsson, D.E., Byrne, M.J., and Dacke, M. (2017). Stellar performance: mechanisms underlying milky way orientation in dung beetles. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 372, 20160079. <https://doi.org/10.1098/rstb.2016.0079>.
- Hardcastle, B.J., Omoto, J.J., Kandimalla, P., Nguyen, B.C.M., Keleş, M.F., Boyd, N.K., Hartenstein, V., and Frye, M.A. (2021). A visual pathway for skylight polarization processing in drosophila. *Elife* 10, e63225–e63257. <https://doi.org/10.7554/eLife.63225>.
- Heinze, S. (2017). Unraveling the neural basis of insect navigation. *Curr. Opin. Insect Sci.* 24, 58–67.
- Hoinville, T., and Wehner, R. (2018). Optimal multiguideance integration in insect navigation. *Proc. Natl. Acad. Sci. USA* 115, 2824–2829. <https://doi.org/10.1073/pnas.1721668115>.
- Jeffery, K.J., Page, H.J.I., and Stringer, S.M. (2016). Optimal cue combination and landmark-stability learning in the head direction system. *J. Physiol.* 594, 6527–6534. <https://doi.org/10.1113/JP272945>.
- El Jundi, B., Foster, J.J., Byrne, M.J., Baird, E., and Dacke, M. (2015). Spectral information as an orientation cue in dung beetles. *Biol. Lett.* 11, 20150656–20150665. <https://doi.org/10.1098/rsbl.2015.0656>.
- el Jundi, B., Smolka, J., Baird, E., Byrne, M.J., and Dacke, M. (2014). Diurnal dung beetles use the intensity gradient and the polarization pattern of the sky for orientation. *J. Exp. Biol.* 217, 2422–2429. <https://doi.org/10.1242/jeb.101154>.
- Khaldy, L., Tocco, C., Byrne, M., and Dacke, M. (2021). Compass cue integration and its relation to the visual ecology of three tribes of ball-rolling dung beetles. *Insects* 12, 526. <https://doi.org/10.3390/insects12060526>.
- Kim, S.S., Hermundstad, A.M., Romani, S., Abbott, L.F., and Jayaraman, V. (2019). Generation of stable heading representations in diverse visual scenes. *Nature* 576, 126–131. <https://doi.org/10.1038/s41586-019-1767-1>.
- Kim, S.S., Rouault, H., Druckmann, S., and Jayaraman, V. (2017). Ring attractor dynamics in the *Drosophila* central brain. *Science* 356, 849–853. <https://doi.org/10.1126/science.aal4835>.
- Knight, R., Piette, C.E., Page, H., Walters, D., Marozzi, E., Nardini, M., Stringer, S., and Jeffery, K.J. (2014). Weighted cue integration in the rodent head direction system. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 369, 20120512. <https://doi.org/10.1098/rstb.2012.0512>.
- Knill, D.C., and Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci.* 27, 712–719. <https://doi.org/10.1016/j.tins.2004.10.007>.
- Legge, E.L.G., Wystrach, A., Spetch, M.L., and Cheng, K. (2014). Combining sky and earth: desert ants (*Melophorus bagoti*) show weighted integration of celestial and terrestrial cues. *J. Exp. Biol.* 217, 4159–4166. <https://doi.org/10.1242/jeb.107862>.
- Linsenmair, K.E. (1972). *Anemomenotactic Orientation in Beetles and Scorpions* (NASA, Washington Animal Orientation and Navigation), pp. 501–510.
- Mardia, K.V., and Jupp, P.E. (2009). *Directional Statistics* (John Wiley & Sons).
- Müller, M., and Wehner, R. (2007). Wind and sky as compass cues in desert ant navigation. *Naturwissenschaften* 94, 589–594. <https://doi.org/10.1007/s00114-007-0232-4>.
- Murray, R.F., and Morgenstern, Y. (2010). Cue combination on the circle and the sphere. *J. Vis.* 10, 15. <https://doi.org/10.1167/10.11.15>.
- Nardini, M., Jones, P., Bedford, R., and Braddick, O. (2008). Development of cue integration in human navigation. *Curr. Biol.* 18, 689–693. <https://doi.org/10.1016/j.cub.2008.04.021>.
- Okubo, T.S., Patella, P., D'Alessandro, I., and Wilson, R.I. (2020). A neural network for wind-guided compass navigation. *Neuron* 107, 924–940.e18. <https://doi.org/10.1016/j.neuron.2020.06.022>.
- Page, H.J.I., Walters, D.M., Knight, R., Piette, C.E., Jeffery, K.J., and Stringer, S.M. (2014). A theoretical account of cue averaging in the rodent head direction system. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 369, 20130283. <https://doi.org/10.1098/rstb.2013.0283>.
- Seelig, J.D., and Jayaraman, V. (2015). Neural dynamics for landmark orientation and angular path integration. *Nature* 521, 186–191.
- Sjolund, L.A., Kelly, J.W., and McNamara, T.P. (2018). Optimal combination of environmental cues and path integration during navigation. *Mem. Cognit.* 46, 89–99.
- Sun, X., Mangan, M., and Yue, S. (2018). An analysis of a ring attractor model for cue integration. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 10928 LNAI, pp. 459–470. https://doi.org/10.1007/978-3-319-95972-6_49.
- Ugolini, A. (2001). Relationship between compass systems of orientation in equatorial sandhoppers. *Anim. Behav.* 62, 193–199.
- Ugolini, A. (2002). The orientation of equatorial sandhoppers during the zenithal culmination of the sun. *Ethol. Ecol. Evol.* 14, 269–273.
- Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., et al.; SciPy 1.0 Contributors (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Methods* 17, 261–272. <https://doi.org/10.1038/s41592-019-0686-2>.
- Wehner, R., Hoinville, T., Cruse, H., and Cheng, K. (2016). Steering intermediate courses: desert ants combine information from various navigational routines. *J. Comp. Physiol. A Neuroethol. Sens. Neural Behav. Physiol.* 202, 459–472. <https://doi.org/10.1007/s00359-016-1094-z>.
- Wystrach, A., Mangan, M., and Webb, B. (2015). Optimal cue integration in ants. *Proc. Biol. Sci.* 282, 20151484. <https://doi.org/10.1098/rspb.2015.1484>.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Deposited data		
Exit angles/behavioral data	This study	Zenodo: 10.5281/zenodo.5724225
Model simulated data	This study	Zenodo: 10.5281/zenodo.5724225
Experimental models: Organisms/strains		
Diurnal <i>Kheper lamarcki</i>	Wild caught	MacLeay, 1821
Software and algorithms		
Python	Open source	Version 3.8.10
Oriana	Kovach Computing services, UK	Version 3.21
Other		
Handycam	Sony, Japan	FDR-AX53

RESOURCE AVAILABILITY

Lead contact

Further questions should be directed to and will be answered by the lead contact, Elin Dirlik (elin.dirlik@biol.lu.se).

Materials availability

This study did not generate new unique reagents.

Data and code availability

- Behavioral data have been deposited at Zenodo (hosted on GitHub) and are publicly available as of the date of publication. DOIs are listed in the [key resources table](#).
- All original code has been deposited at Zenodo (hosted on GitHub) and is publicly available as of the date of publication. DOIs are listed in the [key resources table](#).
- Any additional information required to reanalyze the data reported in this paper is available from the [lead contact](#) upon request.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Ethical statement

All applicable international, national and institutional guidelines for the care and use of animals were followed. Animal care was in accordance with the EU Directive 2010/63/EU and the South African National Standard for The Care and Use of Animals for Scientific Purposes.

Animal collection and experimental sites

Ball rolling dung beetles of the species *Kheper lamarcki* were collected using dung baited pit-fall traps at Stonehenge game farm, South Africa (26.39° S, 24.32° E) during November 2019, as well as February and November 2020. Behavioural experiments were conducted indoors at Bergsig Eco Estate game farm near Bela Bela and at University of the Witwatersrand, South Africa, as well as at Lund University, Sweden. Beetles were kept in plastic containers filled with sand and provided with fresh dung 2–4 times per week.

METHOD DETAILS

Experimental setup

The setup consisted of an artificial sky constructed of two metal arches crossed over to create a hemisphere (1.5 m radius). Each arch was lined with 141 LEDs (520 nm, DotStar; Adafruit Industries, New York, USA) approximately 1.3° apart. A single LED served as an ersatz sun with an intensity of 2×10^{11} photons/cm²/sec as

measured from the centre of the setup at a height of about 7 cm (corresponding to the height of a beetle on top of its dung ball) using a spectrometer (QE65000; Ocean Optics). Serving as wind cues, four wind generators were positioned on the floor 1.3 m from the centre of the setup. The first wind generator was aligned with one of the LED-lined arches and the remaining three were placed at an angle of 60°, 120° and 180° relative to the first. Each wind generator was constructed from three fans (PFR0912XHEE, 4.50A; Delta Electronics Inc., Taipei City, Taiwan) separated by 0.25 m and was powered by a Mean Well RSP-320-12, 26.7A power supply. Measures of wind speeds were obtained by the use of a hot wire anemometer (HHF-SD1; Omega) placed 7 cm above the centre of the arena (see [Figure S4](#)). A sand-painted circular arena (0.3 m radius) was placed in the centre of the setup, with the arena perimeter labelled from 0 to 355° in 5° increments and with 0° aligned with magnetic north. To control solar elevation and wind speed, custom-built software was used with an Arduino Uno (experiments conducted in South Africa), or a Raspberry Pi 4 Model B (experiments conducted in Sweden). All experiments were filmed from above using a Sony camera (FDRAX53 Handycam) or a Raspberry Pi camera (Camera Module 2 NoIR), supported by infrared illumination (B07DDJ1YDB, 1A; eecoo, Shenzhen, China). To eliminate unwanted cues, the setup was placed inside a tent made out of blackout cloth (see [Figure S1](#)).

Behavioural experiments

Throughout each experimental day beetles were temporarily kept in shallow bins containing fresh dung and given time to construct dung balls. During behavioural experiments, each beetle was placed alongside its dung ball in the centre of the circular arena (semi-randomly in one of four cardinal directions). Following its characteristic orientation dance, the beetle was allowed to roll to the perimeter where its exit bearing was recorded. The beetle and its ball were then placed back into the centre of the arena and the procedure was repeated a number of times that depended on the experimental question (see below). In total, each beetle took between 5 and 15 min to complete an experimental series, after which it was put away for the day. The same beetle was never tested more than once for each experiment and if it performed another experimental series this was always carried out on a different day.

Reliability of sun and wind cues

We used orientation precision of ball rolling beetles as a proxy for reliability under different cue conditions with the assumption that more reliable cues would lead to greater precision and vice versa. Orientation precision under an ersatz sun was tested at elevations of 5°, 20°, 45°, 60°, 75°, 80°, 82°, 84°, 86°, 88°, or 90°. For every elevation, 20 beetles were tested. Each individual was marked to ensure that it was only used once per elevation. Each beetle was placed in the center of the arena and allowed to exit from it five times. Following this, the azimuth of the ersatz sun was shifted by 180° and the beetle exited the arena an additional five times.

The same procedure was used to test the beetles' ability to perform straight-line orientation at wind speeds of 0.5, 0.8, 1.0, 1.5, 1.9, 2.5, 3.0 and 4.0 m/s, with the direction of the wind current shifted by 180° after five exits. For every wind speed 20 beetles were tested and each individual was marked to ensure that it was only tested once per wind speed. To sustain the beetles' motivation in the presence of wind, this experiment was performed with an ersatz sun positioned in zenith.

Cue conflict between an ersatz sun and wind

Cue conflict experiments. Based on the results gathered from the reliability experiments described above, a cue conflict experiment was conducted using an ersatz sun at elevations of 45°, 60°, 75°, and 86° in the presence of a wind current of 2.5 m/s. In a separate cue conflict assay, solar elevations of 60°, 75° and 86° were presented together with a wind speed of 1.25 m/s. All conflicts were achieved by shifting the azimuthal direction of the wind current while keeping the ersatz sun stationary.

Each beetle exited the arena a total of eight times: three times with the directional information from the ersatz sun and wind current in congruence, once with a conflict of 60° (or 120°), once with the cues in their original position (congruent), once with a conflict of 120° (or 60°, respectively), and finally two exits with the cues returned to their original positions (congruent). The purpose of the repeated congruent exits was to ensure that the beetles strived to adhere to the same bearing throughout its experimental series (see [STAR Methods, Quantification and statistical analysis - Cue conflict between an ersatz sun and wind](#)). The order in which the conflicts were presented was pseudo-randomised. Thus, each beetle performed both test conditions where the azimuthal directional information of the two cues were put in conflict by 60° or 120°

between two consecutive rolls, as well as a control condition where the directional information remained unchanged (0° conflict).

Further cue conflicts at 60° elevation. The same cue conflict assay (congruent (×3) - 60°/120° conflict - congruent - 120°/60° conflict - congruent (×2)) was replicated in another experimental series. These spanned over three days with the ersatz sun at an elevation of 60° and the wind speed set to 2.5 m/ s. Each day, data was collected from the same population of individually marked beetles.

Individual precision. Individual precision was studied at an elevation of 60° and a wind speed of 2.5 m/ s. In this experimental setting, the beetles exited the arena a total of 36 times: ten times with an ersatz sun and a wind current in congruence, ten times with a conflict of 60° (or 120°), three times with the cues in their original positions, ten times with a conflict of 120° (or 60°, respectively), and finally three times with the cues returned to their original positions.

Simulation overview

The software performs a simplified simulation of the cue conflict paradigm above. We are interested in the change in the value of an integration of two angular inputs with von Mises noise. We define a von Mises distribution for each cue (using the precision data described in [STAR Methods, Behavioural experiments - Reliability of sun and wind cues](#)) and an angle is sampled from each. We then compute the integration of these angles when the cues are aligned (their distributions have the same mean), and when the cues are in conflict (the distributions have different means). The difference between the two integrations is the change in the integration, which can be interpreted as a change in bearing. We compared five different cue integration models and assessed their ability to produce our behavioral data by comparing how likely the data would be under any candidate model (see [STAR Methods, Evaluation process](#)).

Simulated cue representation

In order to capture sensory noise in a circular context, cues are treated as independent von Mises random variables ([Murray and Morgenstern, 2010](#)). The von Mises probability density function is given by:

$$f_{VM}(x; \mu, \kappa) = \frac{e^{\kappa \cos(x - \mu)}}{2\pi I_0(\kappa)} \quad (\text{Equation 1})$$

where μ is the mean angle of the distribution, κ is the concentration (equivalent to σ^{-2} for the normal distribution, often called “reliability” ([Ernst and Bühlhoff, 2004](#); [Murray and Morgenstern, 2010](#))), and $I_0(a)$ is the modified Bessel function of the first kind of order zero ([Batschelet, 1981](#); [Murray and Morgenstern, 2010](#)). This is analogous to using normal distributions to simulate Gaussian noise when working with linear data (e.g. time taken by an animal to exit an arena), rather than angular data (e.g. angle at which the animal exits the arena). To sample from these distributions we need to estimate parameters μ_{Wind} , κ_{Wind} , μ_{Light} , κ_{Light} , such that the distributions $f_{Light}(x; \mu_{Light}, \kappa_{Light})$ and $f_{Wind}(x; \mu_{Wind}, \kappa_{Wind})$ are those which can produce simulations which match the observed behavior under light (an ersatz sun) or wind respectively.

The estimates for the means are the input azimuths of each cue; it is reasonable to assume that the average perceived cue position is the true cue position. The concentration parameter estimates can be approximated from the mean vector length of a random sample from a parent distribution ([Mardia and Jupp, 2009](#)). The best available proxy for such a random sample is the data collected to examine the reliability of sun and wind cues (see [Results and Discussion, Reliability of ersatz sun and wind cues](#)). We tried to model the reliability data using linear and split-linear fits respectively (performed using SciPy curve-fitting utilities ([Virtanen et al., 2020](#))); these fits can be seen in [Figure 1](#). However, if we try to approximate the κ -values from these directly the resultant populations are less precise than they should be; to fix this, we included small additive constants which augment the mean vector lengths, improving the final κ approximation with respect to the observed data. The final estimators are:

$$R_{Wind} = (0.11s + 0.43) + c_{Wind} \quad (\text{Equation 2})$$

$$R_{Light} = \begin{cases} (-0.07\varphi + 0.80) + c_{Light}, & \text{if } \varphi \leq 75^\circ, \\ (-1.26\varphi + 2.31) + c_{Light} & \text{otherwise.} \end{cases} \quad (\text{Equation 3})$$

where s is wind speed, ϕ is light elevation with additive constants $c_{Light} = 0.135$, and $c_{Wind} = 0.133$. c_{Light} and c_{Wind} were tuned by hand. Note that line parameters given in Equations 2 and 3 are rounded for brevity, and precision will significantly affect the outcome of the estimation. Parameters at full precision can be recovered from our codebase or by using the same SciPy curve-fitting utilities on our precision data.

We can now estimate an augmented mean vector length R for any sensible light elevation or wind speed which we can use to approximate $\hat{\kappa}$, from Mardia and Jupp (2009) (pg. 85, 86):

$$\hat{\kappa} \approx \begin{cases} 2R + R^3 + \frac{5}{6}R^5, & \text{if } R < 0.53 \\ \frac{1}{2(1 - R) - (1 - R)^2 - (1 - R)^3}, & \text{if } R \geq 0.85 \\ -0.4 + 1.39R + \frac{0.43}{(1 - R)}, & \text{otherwise} \end{cases} \quad (\text{Equation 4})$$

The quality of this approximation can be seen in Figure S7; the approximation is slightly faster to compute, saving some time when running larger simulations. We can test the κ -values by simulating the precision experiments used to estimate them. Including the additive corrections, this method allows us to simulate beetle populations which approximately match real beetles under single-cue conditions.

Integration models

With the above cue representation, we compared five different simple models to evaluate how likely they are to have produced the experimental data. Each integration is computed twice per simulated individual; once for the initial condition and once for the conflict condition.

Winner-take-all (WTA)

Under winner-take-all, we compute weights for each cue and the integration is simply the cue azimuth of the cue with the greatest weight. Weights and integration are given by:

$$W_{Wind} \leftarrow \kappa_{Wind} \quad (\text{Equation 5})$$

$$W_{Light} \leftarrow \kappa_{Light} \quad (\text{Equation 6})$$

$$I_{WTA} = \begin{cases} \theta_{Wind} & \text{if } W_{Wind} > W_{Light}, \\ \theta_{Light} & \text{otherwise.} \end{cases} \quad (\text{Equation 7})$$

Note we do not check the case where cues have equal weights because this never occurs. In such an instance you could break the tie randomly.

Weighted arithmetic mean (WAM)

WAM is the standard (statistically optimal) weighted average model which arises throughout cue integration literature (Ernst and Banks, 2002; Ernst and Bühlhoff, 2004; Knill and Pouget, 2004). A weighted arithmetic mean is not appropriate for angular or otherwise cyclic inputs (Batschelet, 1981; Murray and Morgenstern, 2010); a standard example in circular statistics is to consider the average of 0° and 360° . If we assume equal weights, then Equation 10 will give 180° where we would expect 0° . However, this method has previously been used in the context of directional cue integration in ants (Sun et al., 2018; Wystrach et al., 2015), humans (Alais and Burr, 2004), and monkeys (Fetsch et al., 2012). Furthermore, direction and distance can often be mixed and discussed generally as ‘spatial’ cues (Cheng et al., 2007; Nardini et al., 2008; Sjolund et al., 2018) which can lead to difficulty when interpreting integration across two different domains. Thus, due to its widespread application, we included WAM in our comparison. The weights and integration are given by:

$$W_{Wind} \leftarrow \kappa_{Wind} / (\kappa_{Wind} + \kappa_{Light}) \quad (\text{Equation 8})$$

$$W_{Light} \leftarrow \kappa_{Light} / (\kappa_{Wind} + \kappa_{Light}) \quad (\text{Equation 9})$$

$$I_{WAM} = W_{Wind}\theta_{Wind} + W_{Light}\theta_{Light} \quad (\text{Equation 10})$$

Weighted vector sum (WVS)

This method (due to [Murray and Morgenstern \(2010\)](#)) is derived from a Bayesian integration of angular cues with von Mises noise. Its function is best understood by considering a vector sum heuristic; if we interpret each cue as a polar vector, say $\vec{l} = (\theta_{Light}, W_{Light})$ and $\vec{w} = (\theta_{Wind}, W_{Wind})$, then [Equation 13](#) gives the angular component of $\vec{l} + \vec{w}$ ([Murray and Morgenstern, 2010](#)). This method has been used previously to model the integration of Path Integration and Landmark Guidance cues in ants ([Hoinville and Wehner, 2018](#)). The weights are:

$$W_{Wind} \leftarrow \kappa_{Wind} / (\kappa_{Wind} + \kappa_{Light}) \quad (\text{Equation 11})$$

$$W_{Light} \leftarrow \kappa_{Light} / (\kappa_{Wind} + \kappa_{Light}) \quad (\text{Equation 12})$$

The final integration is:

$$l_{WVS} = \theta_{Wind} + \text{atan2}(\sin(\theta_{Light} - \theta_{Wind}), (W_{Wind} / W_{Light}) + \cos(\theta_{Light} - \theta_{Wind})) \quad (\text{Equation 13})$$

So long as $W_{Wind}/W_{Light} = \kappa_{Wind}/\kappa_{Light}$, this integration is considered optimal ([Murray and Morgenstern, 2010](#); [Hoinville and Wehner, 2018](#)). The following methods are variations on this vector sum which differ only in how the weights are computed (i.e. the magnitude components of the vector arguments).

Non-optimal weighted sum (NVS)

Here we compute the normalised weights as in WAM and WVS, then pass them through a sigmoid adjustment function g . The adjustment function has the effect of minimising the area in weight-space where both cues have a significant impact on the integration (see [Figure S5](#)). Thus, both cues are still considered but it is quite easy for one to dominate the integration resulting in pseudo-WTA behavior. This is distinguished from a true WTA by the fact that increasing cue conflict could still have an effect on overall population spread (as both cues are generally considered, even if one has very little weight).

$$W_{Wind} \leftarrow g(\kappa_{Wind} / (\kappa_{Wind} + \kappa_{Light}); a) \quad (\text{Equation 14})$$

$$W_{Light} \leftarrow g(\kappa_{Light} / (\kappa_{Wind} + \kappa_{Light}); a) \quad (\text{Equation 15})$$

where:

$$g(x; a) = \frac{1}{1 + e^{-a(x-0.5)}} \quad (\text{Equation 16})$$

with $a = 53$ determined as most likely (see [STAR Methods, Evaluation process; Figure S8](#)). The final integration is given by [Equation 13](#).

Biased non-optimal weighted sum (BVS)

Our final model introduces small individual biases which give each individual a random preference for the wind or the ersatz sun. These biases are drawn from a very narrow Gaussian distribution ($\sigma_{bias}^2 = 0.000303$) with $\mu = 0$. As such, we expect that most individuals have no bias. In biological terms, this could be thought of as a weak preference based on prior experience; such a preference would only become apparent where cues are very close in weight. The weights are now computed as:

$$W_{Wind} \leftarrow g((\kappa_{Wind} / (\kappa_{Wind} + \kappa_{Light}) - b); a) \quad (\text{Equation 17})$$

$$W_{Light} \leftarrow g((\kappa_{Light} / (\kappa_{Wind} + \kappa_{Light}) + b); a) \quad (\text{Equation 18})$$

where $b \sim N(0, \sigma_{bias} = 0.017\dots)$ is the bias for this individual. Note that the same bias is used on both the initial and conflict steps. Again, the final integration is given by [Equation 13](#).

Evaluation process

To evaluate the different models, we used each to generate large simulated populations ($n_{sim} = 1,000,000$) under each set of conflict conditions (for each model, the same von Mises random samples were used to minimise the effect of random sampling when comparing models). These populations can be interpreted as probability mass functions (grouped into 5° bins and then normalised) and used to assess how likely each model is to have produced the experimental data above. We cannot simply fit von Mises distributions to our data as: (1) the integration of two von Mises distributions does not necessarily produce a von Mises distribution

(our models do not necessarily produce von Mises distributions) (Murray and Morgenstern, 2010), and (2) three of our test conditions did not produce significantly oriented populations (Figure 2B, 60° elevation/120° conflict, and Figure 2D Day 1 and Day 3) meaning we cannot assume they are von Mises.

Formally, given data for a set of conditions c , and a model M we want to know:

$$P(M|c) \propto P(c|M)P(M) \quad \text{(Equation 19)}$$

$$P(M|c) \propto P(c|M) \quad \text{(Equation 20)}$$

as $P(M)$ is assumed to be uniform. We have 21 different conditions, each of which is conditionally independent, thus:

$$P(c|M) = \prod_{j=1}^{21} P(c_j|M) \quad \text{(Equation 21)}$$

$$\ln(P(c|M)) = \ln\left(\prod_{j=1}^{21} P(c_j|M)\right) \quad \text{(Equation 22)}$$

$$= \sum_{j=1}^{21} \ln(P(c_j|M)) \quad \text{(Equation 23)}$$

For each condition c_j , each data-point is also conditionally independent. By the same reasoning:

$$\ln(P(c_j|M)) = \sum_{i=1}^n \ln(P(d_i|M)) \quad \text{(Equation 24)}$$

$P(d_i|M)$ is drawn from the p.m.f. generated for the model M (i.e. the likelihood of data point d_i occurring given model M). The model with the greatest log likelihood wins. The results in Table 1 are presented as log likelihood ratios given by:

$$\begin{aligned} LR &= \ln \left[\frac{P(c|M)}{P(c|\hat{M})} \right] \\ &= \ln[P(c|M)] - \ln[P(c|\hat{M})] \end{aligned} \quad \text{(Equation 25)}$$

where $P(c|M)$ is the (proportional measure of) likelihood of a model with respect to the data and $P(c|\hat{M})$ is the likelihood of the maximally likely model. The maximally likely model gets a score of 0 and scores decrease as candidates become less likely. Akaike Information Criterion (AIC) ratios are also given to compare model fit while penalising parameter counts:

$$\text{AIC Ratio} = \frac{AIC(M)}{AIC(\hat{M})} \quad \text{(Equation 26)}$$

with

$$AIC(M) = 2k - 2\ln(P(c|M)) \quad \text{(Equation 27)}$$

where k is the number of parameters in the model M . The best AIC ratio score is 1 and scores will increase as fit worsens.

As the AIC did not affect the order of our results, we also examined the Bayesian Information Criterion (BIC) as this more heavily penalises large parameter counts. The BIC is defined as:

$$BIC(M) = k\ln(n) - 2\ln(P(c|M)) \quad \text{(Equation 28)}$$

with $n = 564$ being the total number of samples against which each model is evaluated (over all conflict conditions), and k the number of parameters in the model M . These are also reported as ratios in Table 1, given by:

$$\text{BIC Ratio} = \frac{BIC(M)}{BIC(\hat{M})} \quad \text{(Equation 29)}$$

with M and \hat{M} as above.

QUANTIFICATION AND STATISTICAL ANALYSIS

Analysis of circular data was carried out in Oriana 3.21 (Kovach Computing Services, Anglesey, UK) and all presented bearings are shown as $\mu \pm$ *Circular Standard Deviation (SD)*. All statistical details may be found in the results and discussion. The n-numbers indicate number of individual beetles tested at each experimental condition.

Reliability of sun and wind cues

To investigate (i) the beetles' orientation precision and, (ii) their directional preferences in the presence of a single cue, the ten exit bearings recorded for each beetle were normalised to the azimuthal position of the orientation cue. Beetles whose normalised exits were not significantly different from a uniform distribution ($p \leq 0.05$, Rayleigh test) were deemed unable to orient and thus excluded from analysis for tactic behavior. For each experimental group, i.e. elevation of the ersatz sun (except for 90° elevation) or wind speed, Rayleigh tests were conducted on the population of mean bearings. We define menotaxis as mean bearings taken at any angle with respect to the cue (uniform distribution), while a population showing a directional preference towards or away from a directional cue is defined as performing taxis. Orientation precision was then investigated for the experimental groups that performed menotaxis (i.e. beetles that were able to use the sun and wind stimuli as compass cues).

This was done by calculating the mean vector length (R) from the normalised bearings of each beetle, including the individuals that were previously excluded. The R -value extends from 0 to 1, where a higher value suggests greater precision.

Cue conflict between an ersatz sun and wind

To study the effect of cue reliability on the integration and weighting of directional information given by an ersatz sun and wind in conflict, an exclusion criterion was implemented. The criterion stated that, if the six headings were not significantly different from a uniform distribution when the two cues were in their original positions (congruent) ($p \geq 0.1$, Rayleigh test, see [Figure S2](#) for justification), then the beetle was eliminated from further analysis. This ensured that the remaining beetles were able to return to their original heading consistently and thus able to orient.

Changes in heading direction were calculated using the angular difference between two consecutive exits (see [Figure 2](#)). For changes in heading at the 0° cue conflict, the angular difference was calculated between the first and second exit where the directional information of the two cues remained unchanged (congruent). For changes in heading at a 60° and 120° cue conflict the difference was calculated between an exit where the cues were in congruence and the following exit where the wind cue had been shifted. The population mean change in heading, together with Rayleigh tests ($p \geq 0.05$), were used to determine the behavioural response to the azimuthal shift of the wind. This was carried out for all conflict conditions.

To determine individual precision when presented with a conflict between an ersatz sun at a 60° elevation and a 2.5 m/s wind current, the ten bearings recorded at each conflict were tested for uniformity ($p \geq 0.05$, Rayleigh test).

Conclusion

With hindsight, claims about reliability and (non-)optimality are not well justified (see Chapter 3). Nevertheless, the behavioural data and modelling do point to a subtly weighted integration of orientation cues rather than the previously suggested winner-take-all (Dacke et al., 2019). This can be argued without any appeal to cue reliability as we can interpret the adjusted weights used by NVS and BVS as arbitrary weights. (That is, we can conserve our claims in light of the limitations of Bayesian modelling in insects as discussed in Chapter 3.) The necessity of the weight adjustment also implies that the ‘Bayesian integration’ (in this context WVS or WAM) suggested by Khaldy et al. (2022) is also not entirely correct. Rather, with regards to weighting, the true solution probably lies somewhere in between.

The modelling supports the broad claim that, whatever weighting strategy the beetles are using, the underlying computation appears to be a form of vector sum. Of note is that the vector sum can produce a wide variety of behaviours depending on how the weights are set (and by extension, with respect to what property of the world). This is arguably so general as to be unhelpful, however, it is critical to place this conclusion into the context of the brain. In fruit flies, the putative neural substrate for this form of cue integration is head direction circuit, specifically the interface between R and EPG neurons (Fisher, 2022; Wilson, 2023). This region of the brain is sufficiently conserved across species that it seems likely that the same is true in dung beetles and other insects. Therefore, conserved neuroarchitecture (conserved computation) could support a variety of behaviours depending on the state of the inputs. The following chapter attempts to place the vector sum into the context of the head direction circuit.

Chapter 5

Paper 2: A model of cue integration as vector summation in the insect brain

Introduction

The work in the previous chapter suggests that dung beetle cue integration likely takes the form of a vector sum, though this can be parameterised in different ways to generate different behaviours. The insect head direction circuit provides a potential substrate for this computation and this paper provides an implementation of this circuit.

The model proposes that different R neuron populations represent different cue vectors as sinusoidal signals. These sinusoids are summed in the EPG neurons, which should represent a vector sum if the connections between R and EPG neurons (R \rightarrow EPG mappings) are properly formed. In fruit flies, these connections are plastic (Fisher et al., 2019; Kim et al., 2019), and thought to be updated using Hebbian learning (Fisher, 2022; Wilson, 2023). This flexibility could allow the connections to store information about physical cue offsets as well as input strength and relative variability Wilson (2023).

Note This work is Open Access, available under CC BY 4.0. All behavioural work was conducted by Shahrzad Shaverdian. The complete supplementary material can be found in Appendix D.

Author contributions Robert Mitchell: conceptualization, investigation, methodology, software, visualization, writing—original draft, writing—review and editing; Shahrzad Shaverdian: conceptualization, investigation, methodology, visualization, writing—original draft; Marie

Dacke: conceptualization, funding acquisition, project administration, writing—review and editing; Barbara Webb: conceptualization, funding acquisition, project administration, supervision, writing—review and editing. (Given on page 11 of the article.)

Research



Cite this article: Mitchell R, Shaverdian S, Dacke M, Webb B. 2023 A model of cue integration as vector summation in the insect brain. *Proc. R. Soc. B* **290**: 20230767. <https://doi.org/10.1098/rspb.2023.0767>

Received: 31 March 2023

Accepted: 30 May 2023

Subject Category:

Behaviour

Subject Areas:

behaviour, computational biology

Keywords:

vector, cue integration, plasticity, contrast, reliability, neural model

Author for correspondence:

Robert Mitchell

e-mail: r.mitchell@ed.ac.uk

Electronic supplementary material is available online at <https://doi.org/10.6084/m9.figshare.c.6688818>.

A model of cue integration as vector summation in the insect brain

Robert Mitchell¹, Shahrzad Shaverdian², Marie Dacke² and Barbara Webb¹

¹Institute for Perception, Action, and Behaviour, The University of Edinburgh School of Informatics, Edinburgh, Edinburgh EH8 9AB, UK

²Lund Vision Group, Department of Biology, Lund University, Lund SE-223 62, Sweden

RM, 0000-0002-6396-0995; SS, 0000-0001-6076-5310; MD, 0000-0001-6444-7483; BW, 0000-0002-8336-6926

Ball-rolling dung beetles are known to integrate multiple cues in order to facilitate their straight-line orientation behaviour. Recent work has suggested that orientation cues are integrated according to a vector sum, that is, compass cues are represented by vectors and summed to give a combined orientation estimate. Further, cue weight (vector magnitude) appears to be set according to cue reliability. This is consistent with the popular Bayesian view of cue integration: cues are integrated to reduce or minimize an agent's uncertainty about the external world. Integration of orientation cues is believed to occur at the input to the insect central complex. Here, we demonstrate that a model of the head direction circuit of the central complex, including plasticity in input synapses, can act as a substrate for cue integration as vector summation. Further, we show that cue influence is not necessarily driven by cue reliability. Finally, we present a dung beetle behavioural experiment which, in combination with simulation, strongly suggests that these beetles do not weight cues according to reliability. We suggest an alternative strategy whereby cues are weighted according to relative contrast, which can also explain previous results.

1. Introduction

Cue integration is the process of combining multiple redundant sources of information to form a single estimate of a property of the world [1]. A widespread assumption in the cue integration literature is that brains integrate cues for the *purpose* of maximizing certainty about the estimate [2], from which it follows that cues should be weighted according to their reliability [1,3]. Integration of angular cues (e.g. for orientation) has previously been formalized as a vector sum [4]—if different directional cues are represented as vectors with magnitudes reflecting their weighting, the integration is given by their sum.

Recent models of insect navigation have suggested that vector computation plays a major role in their behavioural capabilities [5,6] and more specifically that they perform reliability-weighted vector-based cue integration [7–9]. This idea has been bolstered by direct evidence that the insect central complex has the necessary circuit properties to support vector computations by representing vectors as sinusoid curves (vector phasor representation) in activity across a neural array [10]. The central complex (CX) is a collection of midline neuropils which is highly conserved across insect species [11], of which the key components are the protocerebral bridge (PB), ellipsoid body (EB), fan-shaped body and a pair of noduli [12,13]. The function and structure of the CX has so far been mapped in most detail in *Drosophila*. The head direction of the insect with respect to external cues is tracked by the activity of a set of EB neurons known as E-PGs [13,14], also colloquially referred to as 'compass neurons'. E-PGs receive input from structures known as the bulbs which sit laterally to the CX [13,15–17] and the bulbs house a class of neurons known as ring neurons

(named for their ring-like arborizations in the EB [12]). A subset of ring neurons, the R neurons, are directionally tuned, responding to the position of a cue around the animal [18]. Each R neuron connects to all E-PGs [13,19] and these connections are plastic, enabling flexible remapping of the outside world onto the compass neurons [16,20,21]. Ring neurons also cluster into different types which seem to encode the position of different orientation cues [15,17]. Where known, the functional role of different CX structures appears to be preserved in other insects. Given that R neurons are directionally tuned, we propose that their population activity profile could be roughly sinusoidal in shape, with the consequence that the E-PG's activity could be a pointwise sum of sinusoids, representing a vector sum of different cue modalities. The encoding of different cue modalities as well as the ability to create flexible relationships between available cues make the R-to-E-PG interface an ideal substrate for flexible multimodal cue integration [9,13,19].

Multimodal cue integration has recently been discovered in the ball-rolling dung beetle *Kheper lamarcki* (MacLeay, 1821) [22]. These beetles have long been known to have a variety of orientation cues at their disposal [23] which enable them to maintain a straight path of arbitrary direction when rolling a ball away from the dung pile. However, it was previously thought that these beetles followed a cue hierarchy [24], i.e. a single cue would dominate at any time, depending on circumstances such as availability, evolutionary history, etc. A collection of new studies are showing that these beetles can in fact use multiple cues simultaneously [9,22,25,26]. The most recent study which examined the cue integration of *K. lamarcki* suggested that these beetles followed a vector summation strategy for the integration of (the directional information provided by) a visual sun cue and a mechanosensory wind cue [9]. Both visual and mechanosensory cues are known to reach the ring neurons in flies [17], with visual input also being confirmed in dung beetles [27]. *Kheper lamarcki* has also been shown to be able to synchronize information from different compass cues [22]. Plasticity between the ring neurons and E-PGs could enable such synchronization [16,17,19,20]. Angular velocity appears to be the signal that regulates plasticity between the ring neurons and E-PG neurons [21] (also modelled by [28]). It is therefore striking to note that *K. lamarcki* perform a stereotyped 'orientation dance' before beginning ball-rolling behaviour [29] which takes the form of a rotation on top of the ball, and ignore orientation cues not present during the dance [30]. Thus, dung beetle behaviour aligns well with known and plausible functionality of the insect head direction circuit in the CX.

In this paper, we extend the model of the insect head direction circuit from [31,32] by including the R neurons and their plastic connections onto E-PG neurons [15–17,20]. We show that if plasticity is gated by angular velocity, the connections form coherently and allow the circuit to perform cue integration by vector summation. We then explore whether and how the relative reliability of cues determines their influence on behaviour, in simulations of and experiments on dung beetles. Together, these results strongly suggest that dung beetles do not weight cues according to actual reliability. We propose that they instead use cue contrast as a proxy estimate for reliability.

2. Material and methods

(a) Conceptual framework

In figure 1 we provide clarification of how several otherwise ambiguous terms will be used in this paper.

(b) Circular cue integration as a vector sum

Consider two von Mises random variables C and H . For samples c_i and h_i , their integration is given by [4]:

$$l_i = h_i + \text{atan2}\left(\sin(c_i - h_i), \left(\frac{\kappa_H}{\kappa_C}\right) + \cos(c_i - h_i)\right), \quad (2.1)$$

with the joint reliability given by

$$\kappa_L = \sqrt{\kappa_H^2 + \kappa_C^2 + 2\kappa_H\kappa_C\cos(c_i - h_i)}, \quad (2.2)$$

κ_H , κ_C could be arbitrary positive weights, but for 'optimal' integration are set as the concentrations (reliabilities) of the von Mises distributions which describe C and H . Note, however, this 'optimality' is based on approximation to the classic maximum-likelihood model [1] for specific weight combinations and conflicts (see appendix of [4]).

This model for integration is equivalent to polar vector addition. Given two cue vectors $\vec{c} = (c_i, \kappa_C)$ and $\vec{h} = (h_i, \kappa_H)$, the integration is simply $\vec{l} = \vec{c} + \vec{h} = (l_i, \kappa_L)$, where l_i gives the angle and κ_L the joint reliability. Our aim in the following is to construct a neural circuit that results in equivalent integration. More specifically, it should produce an output corresponding to the angle, l_i , which can be used to guide steering behaviour; we do not require the magnitude κ_L in our neural model to match that of the pure vector model.

(c) Computer modelling

(i) Ring model overview

Here we provide a textual overview of our model; the included neurons and how they interact. For implementation details, please refer to our electronic supplementary material or code-base. We use fruit fly nomenclature throughout, but non-fly homologues [33] are given for each neuron class.

Our model is a hybrid of the biologically constrained models of the head direction circuits of the fruit fly *Drosophila melanogaster* and the desert locust *Schistocerca gregaria* presented in a previous comparative study [31]; similar to another presented in the context of a model of ant orientation [32]. Specifically, this means we use the 'locust' 8-fold columnar structure (consistent with the description of the PB of *K. lamarcki* [34]) but the 'fly' uniform inhibition [31,35] between compass cells. While it is known that intra compass cell inhibition in beetles more closely reflects the locust anatomy [36], we use a hybrid structure to motivate general plausibility across insect species. The differences in connectivity do not create any functional differences in the circuit which are relevant to the cue integration problem [31]. We extend previous models [31,32] by including the R neurons and their plastic all-to-all connections onto the E-PGs. Our R neuron representation is based on descriptions given by [15–17,20]. Our neurons use the same basic firing rate encoding as [5].

R (TL) neurons. R neurons function as the cue inputs and are split into two groups, one for each cue: R_1 and R_2 (note that these are just indices, we are not explicitly modelling specific classes of ring neuron). Each R neuron has a receptive field centred on its preferred direction. The distribution of receptive fields results in a sinusoidal activity pattern across each R neuron group; the phase of the sinusoid encodes the angle to the cue and the amplitude gives its weight. In the following work we use eight neurons per group to represent the cardinal and ordinal directions (this also makes visualizing the $R \rightarrow E$ -PG mappings

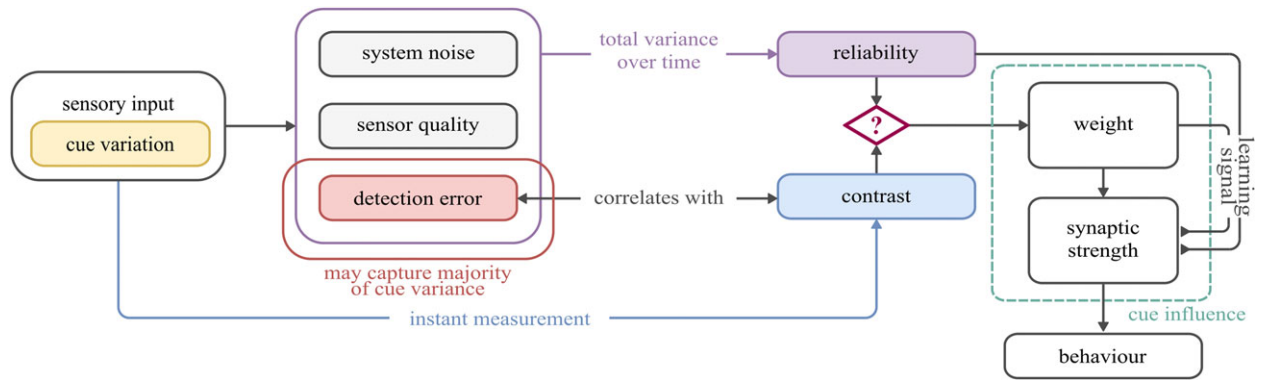


Figure 1. Concepts in cue integration. *Reliability* describes the variance of a cue estimate over time, reflecting external cue variation, but also noise or limitations in the sensory processing system. *Contrast* of a cue with the background influences its detection (and hence reliability) but is an instantaneous property. *Weight*, in this paper, strictly refers to the amplitude of the sinusoid that represents each directional cue at input to the integration circuit. *Synaptic strength* describes the strength of the connections between neurons. Synaptic strength is modified during learning and (as we will show) can be affected differently by cue weight and cue reliability. Cue influence describes the combined effect of weight and the synaptic strengths on final behavioural output.

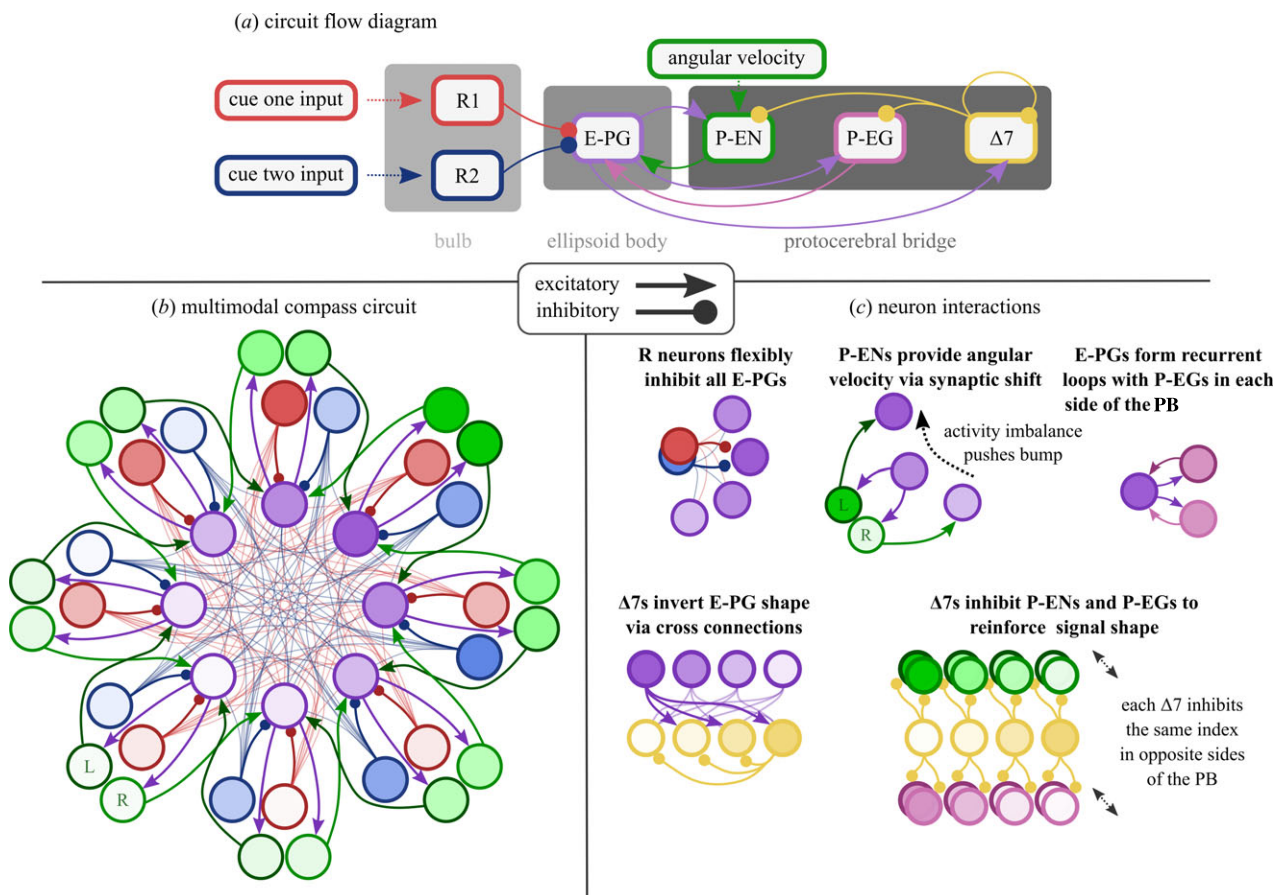


Figure 2. Model illustration. (a) High-level view of information flow through the network (also see the electronic supplementary material, figure S1). (b) The multimodal compass circuit which is the focus of this paper. Each R neuron population (blue, red) encodes a different cue modality as a sinusoidal bump (indicated by shading). R neurons provide input to E-PGs (purple) via plastic all-to-all connections which are formed in combination with angular velocity from P-EN neurons (green). (c) Illustrative microcircuits describing how the different neuron populations interact throughout the model. PB, protocerebral bridge.

easier—see below). As the exact number of R neurons per modality is still uncertain [13,37] we also provide a supplementary exploration of R neuron population size and how this affects the proposed function. In brief, we found that there is no significant difference in network function where R neuron groups are of equal size; however, where group sizes differed, the larger group received more influence in the resulting integration (see the electronic supplementary material).

P-EN (CL2) neurons. P-ENs function as self-motion inputs and are a source of local excitation for the E-PG neurons. The P-ENs are spread across the 16 columns of the PB, dividing into two sub-populations of eight, which represent the same

directional signal (bump) twice over [31,38]; from our observations, this bump appears to emerge owing to input from the E-PGs and $\Delta 7$ s. The different sub-populations drive the E-PG bump in different directions via a single-column synaptic shift. The difference in activity between the two sub-populations is thought to encode idiothetic angular velocity, though it is not currently known if this is proprioception-based or motor efference [21,38]. Total P-EN activity is described as relatively constant across all 16 neurons [38]. We found that the relative constancy of the overall activity was critical in ensuring that the E-PGs are not saturated or starved when the network experiences different angular velocities.

E-PG (CL1a) neurons. E-PGs receive cue azimuth information from the R neurons, self-motion information and local excitation from the P-ENs, and recurrent excitation from the P-EGs. The E-PGs are typically described as dividing into 16 wedges which form a single bump of activity [14]. For ease of development and presentation we instead use eight neurons, however, the number should not significantly affect the accuracy of the angle stored.

In order to account for the results of [14] (E-PG bump retention in darkness and E-PG visual capture) as well as our own proposed learning routine (where E-PGs are driven primarily by angular velocity input), input to the E-PGs is context dependent. If learning is taking place, then R neuron input to the E-PGs is reduced. This can be thought of as an inhibition of R neurons (or enhancement of P-EN neurons) which occurs when the agent/animal is attempting to learn about the cues (see Discussion).

$\Delta 7$ (TB1) neurons. $\Delta 7$ s only receive input from the E-PGs. They: (i) invert the bump from the E-PGs and feed this back into the PB, and (ii) inhibit other $\Delta 7$ s. Both features appear to aid the stability of the bump in the E-PGs and as a result, the rest of the circuit.

P-EG (CL1b) neurons. Finally, the P-EGs provide a simple recurrent loop with the E-PGs, moderated by the $\Delta 7$ s. Together with the P-ENs, the P-EGs maintain activity in the E-PGs in the absence of external input.

(ii) R-to-E-PG connections

Our model is initialised using a default R-to-E-PG connection pattern (figure 4*b*; electronic supplementary material); once the model is initialised, R-to-E-PG connections are all set to an equal value and learned from scratch. In our model, plasticity between R and E-PG neurons is only enabled during specific learning events. R \rightarrow E-PG connections are learned using the following anti-Hebbian update rule:

$$\Delta w_{i,j} = -\eta \cdot (r_{R_i} - \theta_{R_i}) \cdot (r_{E-PG_j} - \theta_E), \quad (2.3)$$

where $w_{i,j}$ is the synaptic strength from from $R_{i,j}$ onto E-PG $_j$, $\theta_E = 0.9$ is the threshold for E-PG activity, and $\eta = 0.1$ is the learning rate. θ_{R_i} is an adaptive threshold on R neuron activity with $\theta_{R_i} = 0.7 \cdot \max(r_{R_i})$ (a fraction of the maximum R neuron activation for group l). Note that if the sinusoidal shape of the R neuron population code can be assumed, then this adaptive threshold could equivalently be a fraction of total activity. Connections are then normalized such that the total synaptic strength onto each E-PG sums to 1. Hebbian learning requires a rapid compensatory mechanism in order to function correctly, and the post-synaptic normalization we are using (keeping total input constant) is one of a few potential regulatory mechanisms (see [39]).

(iii) Behavioural simulations

To provide additional insight into behavioural results, we also constructed a basic behavioural simulation. Conceptually, an agent is placed in the centre of a circular arena and tasked with walking to the edge (in an arbitrary direction with respect to external cues). At the start of each walk (or ‘roll’) the agent performs a learning rotation (or dance). On the first roll of an experiment, the dance will establish R \rightarrow E-PG mappings from a blank mapping (all Rs connect to all E-PGs equally). Subsequent dances will update the existing map. Note that, while cue noise is included, self-motion noise is excluded during the dance. We assume that, over a short rotation, the agent’s perceived motion is correct. Simulation configuration information is given alongside the relevant results. For full details of our simulation environment, please see our electronic supplementary material or codebase.

(d) Animal behaviour

(i) Beetle collection and husbandry

Using dung-baited pitfall traps, the diurnal, ball-rolling beetle *K. lamarcki* was collected at the game farm ‘Stonehenge’ in Vryburg, South Africa (24.32° E, 26.39° S) in November 2021, February 2022, and November 2022. Behavioural experiments were conducted in Lund, Sweden from December 2021 to May 2022, and during December 2022. All beetles were stored in opaque, plastic bins filled with sand of a consistency similar to their natural soil, and fed with horse dung 2–4 times per week. Prior to each experiment, beetles were removed from the sand filled bins and placed in a separate box containing fresh dung for them to construct into balls. Beetles that began to roll their balls of dung were used for experimentation.

(ii) Behavioural experiments

Experimental set-up. All experiments were conducted in an indoor set-up consisting of two metal arches which were crossed to form the skeleton of a hemisphere ($r = 1.5$ m). Each arch was lined with 141 LEDs (520 nm, DotStar; Adafruit Industries, New York, USA) where individual LEDs served as ersatz sun cues with an intensity of 2×10^{11} photons $\text{cm}^{-2} \text{s}^{-1}$ (QE65000; Ocean Optics) as measured from the centre of the set-up at a height of about 7 cm (corresponding to the approximate height of a beetle when on top of its dung ball). Under the arches, in the centre of the set-up, there was a circular, sand-painted arena (radius = 0.3 m); the perimeter of the arena was marked in five degree increments (from 0–355°) and 0° was aligned with magnetic north. Finally, a wind generator was positioned 1.3 m from the centre of the arena and aligned with one arm of the hemisphere. The generator consisted of three fans (PFR0912XHEE, 4.50 A; Delta Electronics Inc., Taipei City, Taiwan) distributed evenly over 1.0 m, and was configured to create an air current with a speed of 2.5 m s^{-1} when measured from the centre of the arena. The elevation of the ersatz sun and the wind speed were controlled using custom-built software in conjunction with a Raspberry Pi 4 Model B. All experiments were filmed using a Sony camera (FDRA53 Handycam) mounted on a tripod above the arena. The set-up was constructed inside a 3×3 m tent constructed from blackout cloth (figure 3*a*).

Experimental procedure. A beetle was placed with its ball in the centre of the circular arena and allowed to roll the ball to the perimeter where its exit angle was noted. The beetle was placed back in the centre and the procedure was repeated to obtain 10 exits in the presence of either an isolated sun or wind cue, or both cues in alignment. When a beetle started with a single cue in isolation, the unseen cue would be added for the subsequent 10 rolls (i.e. if the beetle starts with a sun cue, we add a wind cue). When a beetle started with both cues, we removed one of the cues (figure 3*b* for schematic overview). In total, each beetle rolled its ball 20 times. This experimental procedure was carried out at solar elevations of 45° and 75°, with a wind speed of 2.5 m s^{-1} (based on [9]). The initial condition was presented in a pseudo-randomized order.

Statistical analysis. To assess the orientation precision of an individual beetle, the mean vector length (r -value) was calculated from 10 exit bearings. To test for significant differences between populations of r -values, paired Wilcoxon (Wilcoxon signed-rank) tests were used. Statistical analyses were performed using Oriana 3.21 (Kovach Computing Services, Anglesey, UK) and RStudio 4.1.0. [40]. All p -values presented are unadjusted.

3. Results

(a) The network approximates the angular component of a vector sum

A previous study [9] suggested that the insect head direction circuit could act as a substrate for vector summation, making




Figure 3. Experimental set-up and experimental flow.

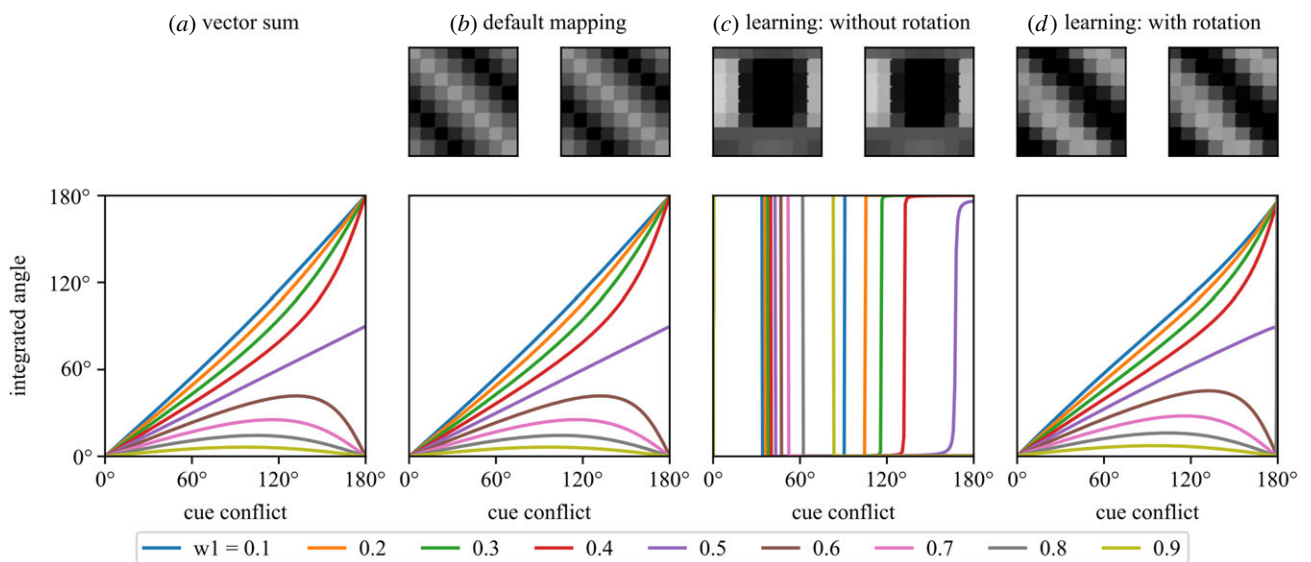


Figure 4. Integrated angular outputs for varying degrees of cue conflict. The pure vector sum [4] is compared against three instances of the ring model with different $R \rightarrow E$ -PG mappings (adjacency matrices shown in the top row). (a) Vector sum (equation (2.1)). (b) Default mapping which implements a pointwise sum-of-sinusoids. (c) The result of Hebbian plasticity alone (equation (2.3)). (d) The result of Hebbian plasticity combined with rotation; the agent rotates through 360° during learning.

it a likely location for compass cue integration. It is clear that our model circuit can encode the angular component of a vector sum to produce an integrated estimate from two directional cues as the amount of conflict between the directions they indicate, and their relative weight, is varied (compare figure 4a, b, and d). Setting a default diagonal connection pattern closely approximates the pure vector sum and our learning rule (equation (2.3)), in combination with a dance rotation, produces a similar diagonal pattern (with arbitrary offset) which generates qualitatively similar output. Learning without experiencing any rotation does not generate a diagonal connection pattern and clearly does not approximate the vector sum (figure 4c); rotation is required to form useful connections (see [21]).

(b) Weight and reliability may independently affect cue influence

(i) Information about cue-state is stored in the $R \rightarrow E$ -PG mappings

It is clear from figure 4 that cue state during the learning routine has an effect on the mappings which are learned, which

will ultimately affect behaviour. To investigate this further, we examined four different learning scenarios: cues separated, one cue useful, effect of weight, and effect of reliability (recall our definitions, figure 1).

Cues separated. One very good reason to have all-to-all plastic connections between the Rs and E-PGs is that it potentially allows an agent to learn about the spatial relationship of physically separated cues [17,30]. Orientation cues may be azimuthally distant without being ‘in conflict’. We found that our model was able to encode cue offsets in the $R \rightarrow E$ -PG maps (figure 5a). The mapping for the second cue is offset from that of the first, meaning that the peak in each R group will be mapped to the same E-PG neuron. This ‘mental’ alignment indicates that the circuit could account for the modality transfer results seen by [22] (see behavioural simulations in the electronic supplementary material).

One cue useful. Here, we wanted to see what would happen in the case that one cue was given significant weight ($w_{\text{blue}} = 0.5$), but provided no directional information. A usable (diagonal) mapping is learned for the useful cue (red—figure 5b) but not for the other (blue). Note that nevertheless the synaptic strengths have a similar range. This indicates that a cue could provide significant input to the

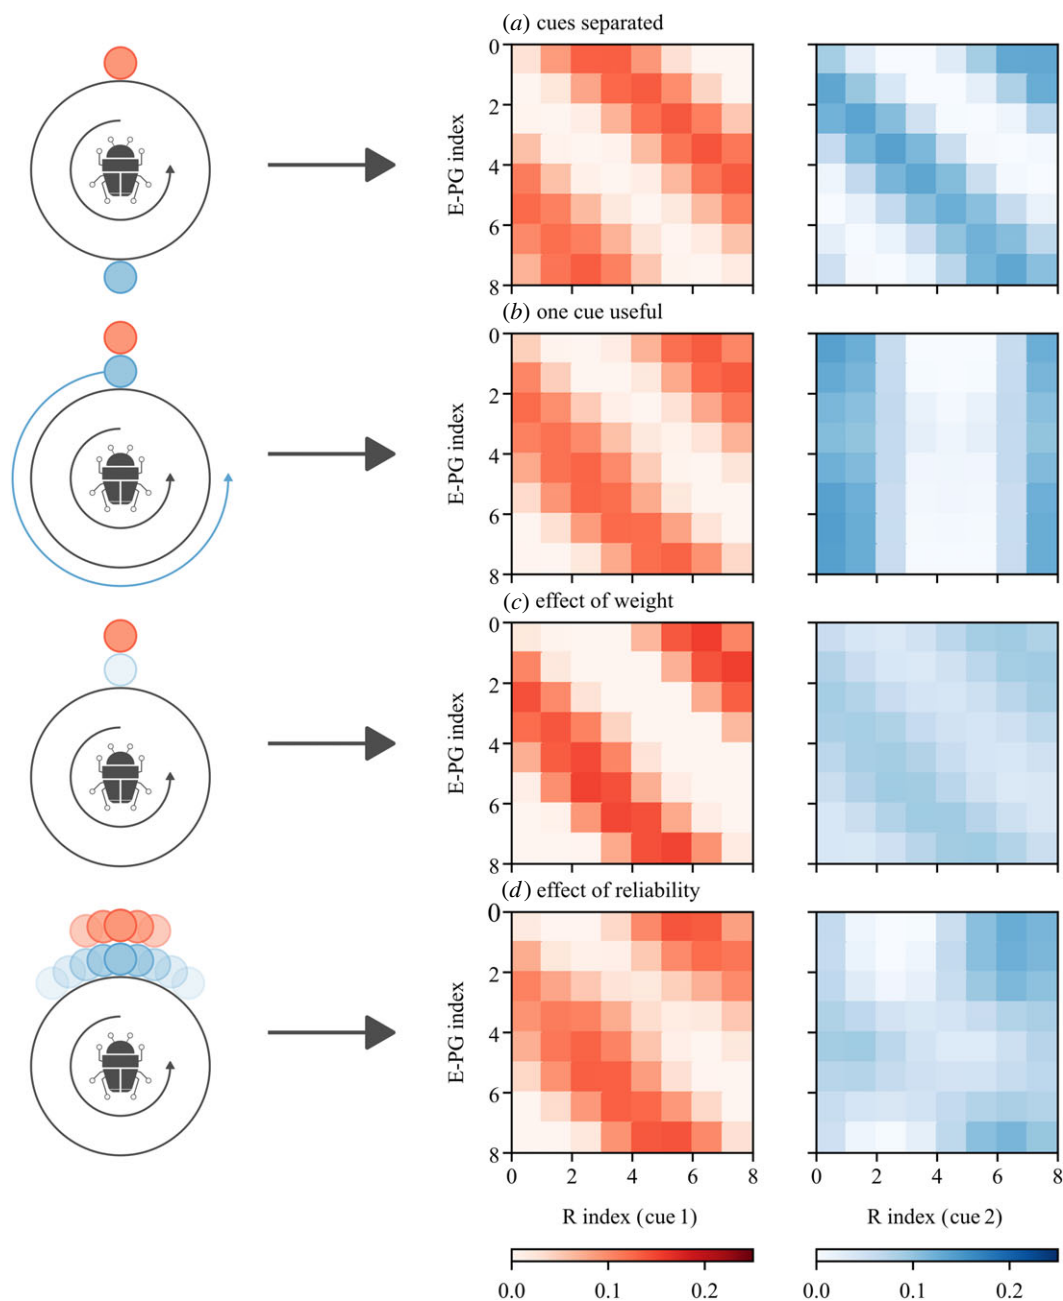


Figure 5. A selection of mappings learned for different cue configurations. (a) Each cue is equally weighted but they are physically separated by 179° . (b) Each cue is equally weighted, however the blue cue rotates with the agent (does not provide directional information). (c) In this case both cues are aligned, however the blue cue is one-quarter the weight of the red. (d) Cues are aligned and of equal weight but have different additive noise distributions (both von Mises). The red cue has concentration $\kappa_{\text{red}} = 2$ and the blue has $\kappa_{\text{blue}} = 0.5$; i.e. the cues have different reliabilities.

integration, even if it does not provide useful orientation information at the time it is learned.

Effect of weight. Here we see that cue weight at the time of learning is encoded in the mapping; greater weight leads to greater synaptic strength. In the extreme, a cue which was not present during the dance should have no influence on behaviour, as seen in the dung beetles [30].

Effect of reliability. In previous experiments, cues have been noiseless (perfectly reliable). This is useful for exploration but may not reflect reality. Variation in the relationship between perceived cue motion and angular velocity should degrade the diagonal pattern of the mapping; this is evident from figure 4c and figure 5b. To explore the effect of this variation, we added von Mises noise to both cues with zero mean and concentrations $k_{\text{red}} = 2$ and $k_{\text{blue}} = 0.5$. These concentrations were chosen so that the effect of noise would be clearly visible.

For higher concentration, the effect of noise on the mapping is minimal. For lower concentrations, it is clear that noise corrupts the mapping but does not appear to have a significant effect on the strength of the connections. Despite being noisy, it appears that the blue cue may have significant influence.

(ii) Cue influence is primarily governed by weight

The mapping results above indicated that both weight and reliability may affect synaptic strength and that the resulting influence on behaviour may not be easy to predict. We therefore simulated a series of cue conflict experiments to investigate the effects of weight and reliability on behavioural output. In each scenario (figure 6a,c), each agent ($n = 100$) exits the arena four times. Note that the agent will update its $R \rightarrow E$ -PG mappings each time it is placed in the arena

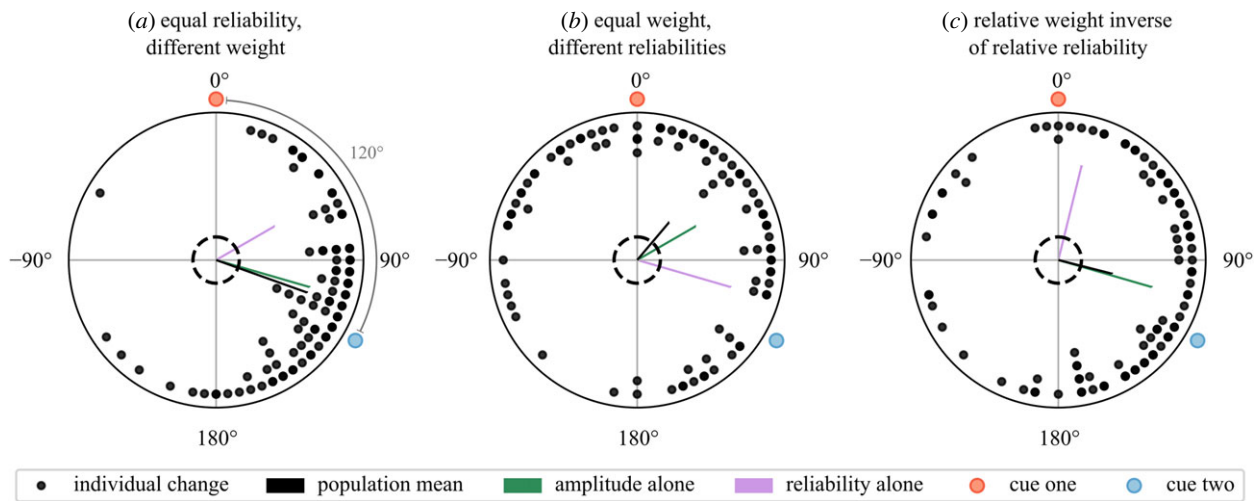


Figure 6. Effect of weight and reliability on cue influence in behavioural simulation of cue conflict. The population mean vector (black) falls closer to the theoretical vector sum where magnitudes are determined by cue weight (green) than where magnitudes are determined by reliability (magenta). Dashed rings indicate the threshold for significance using a Rayleigh test ($p < 0.05$).

centre (§2c(iii)). During the first three rolls the cues are aligned, and on the fourth they are set in conflict. Figure 6 shows the change in bearing between rolls three and four for each agent (as a black dot). Each change is indicative of cue influence for that agent and the population mean gives the average effect (as in the conflict experiment from [9]). The simulated population mean is compared to means predicted by pure vector summation where vector magnitudes are set based on cue reliability (figure 6, magenta vectors), or the arbitrary weight (sinusoidal amplitude) given to the cue (green vectors).

It is clear that the simulated behaviour more closely aligns with the results given by the vector sum where magnitudes are set according to cue weight; that is, sinusoidal amplitude primarily governs cue influence in the circuit. Reliability can affect influence on an individual level and we generally found these simulations to be quite variable (hence the large number of agents, also see the electronic supplementary material); however, this effect does not appear consistent enough to have a population-level effect. This strongly suggests that, if the R → E-PG interface does provide a substrate for cue integration as we propose, the amplitude of the R neuron response should change in response to some property of the cue in order to change cue influence (figure 1)

Previous studies in dung beetles have suggested that cue influence is governed by reliability, however, our simulated results indicate that this is not an emergent property of the head direction circuit combined with plasticity in R neuron connections. Rather, cue weight needs to be set explicitly, meaning it could be decoupled from reliability altogether.

(c) Cue reliability does not determine cue weight in dung beetles

If dung beetles do weight orientation cues according to reliability then, according to the modelling provided by [9], beetle orientation precision with two cues should always be at least as good as orientation precision with the most reliable cue in isolation. We tested this prediction

behaviourally using the same species, and used the neural model to perform a simulated version of the experiment to provide further insight.

(i) Animal behaviour

Beetles were tested in four scenarios:

- (i) sun → sun and wind;
- (ii) wind → sun and wind;
- (iii) sun and wind → sun; and
- (iv) sun and wind → wind.

In each scenario, we have an *initial* and a *test* condition. The beetles rolled their balls 10 times under the initial condition then 10 times under the test condition. The mean vector length is computed for each individual in each scenario for initial and test conditions and box plots are given in figure 7*a*. Each scenario was presented with a solar elevation of 45° or 75° (leading to eight scenarios overall). Note that when the initial condition consists of a single cue, this initial condition will replicate one of the single-cue conditions from [9] (specifically, figure 7*a*(i), (ii), (v) and (vi)). The same is not true when the single cue is presented in the test condition, as we assume that the order of presentation matters (see simulations below). Our key behavioural contribution is the quantitative comparison of single- and multi-cue precision, which was not discussed in [9] (and not explored in detail in [22]).

At a 45° solar elevation, when adding a wind cue to a sun cue, there was no significant difference in orientation precision between the two populations of r -values ($p = 0.11$, Wilcoxon signed-rank test; figure 7*a*(i), $n = 20$). By contrast, when a sun cue was added to a wind cue, there was a significant increase in the beetles' orientation precision ($p < 0.001$, Wilcoxon signed-rank test; figure 7*a*(ii), $n = 20$). When the beetles began with both a sun and wind cue and the wind was removed, there was no significant change in the orientation precision ($p = 0.27$, Wilcoxon signed-rank test; figure 7*a*(iii), $n = 20$). However, when the sun was removed, there was a significant decline in orientation precision ($p < 0.05$, Wilcoxon signed-rank test; figure 7*a*(iv), $n = 30$). In summary, at a 45° solar elevation, the orientation precision of the

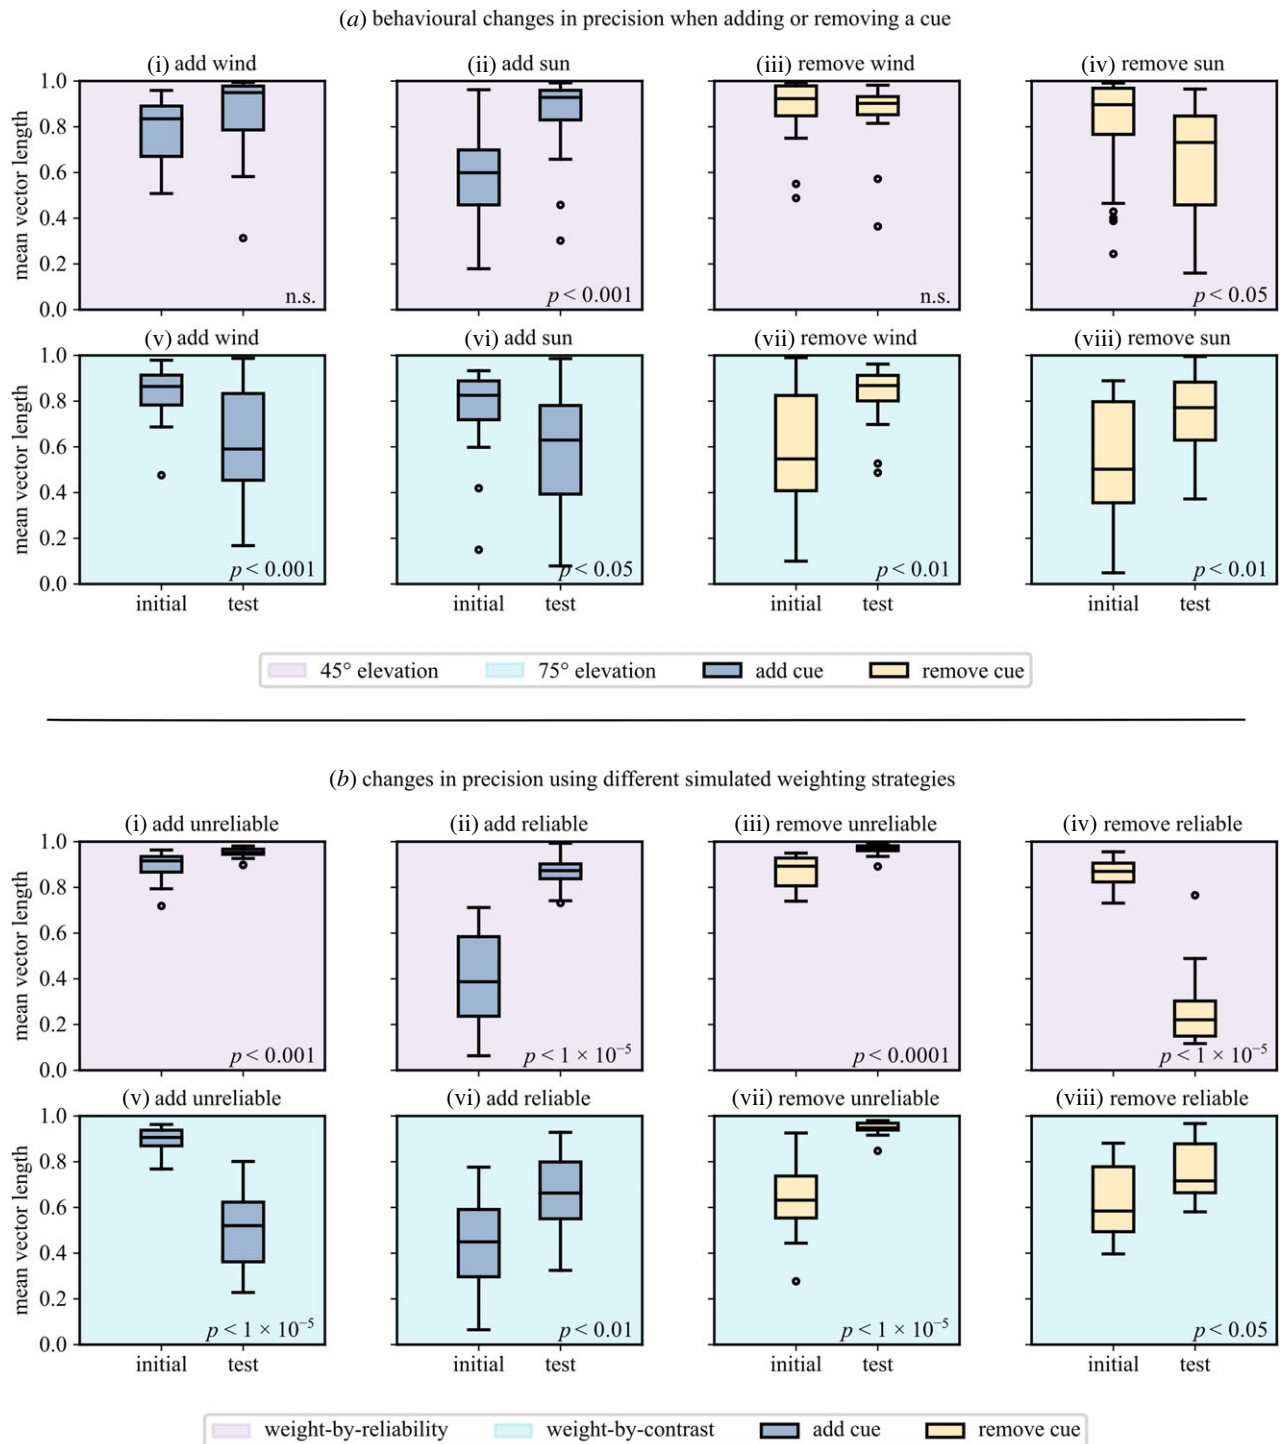


Figure 7. Changes in precision for single or multi-cue conditions. Bars give the median for each group, boxes give first and third interquartile range, and whiskers give the maximum and minimum values. Outliers are marked as separate data points. (a) Dung beetle orientation precision at 45° (top row, magenta) and 75° solar elevation (bottom row, cyan). (b) Simulated model precision using alternative weighting strategies. The top row (magenta) shows the expected outcome where cues are weighted by reliability, the bottom row (cyan) shows the expected outcome where cues are weighted by contrast (weight is decoupled from reliability). Reliable cue has concentration $\kappa = 4$, unreliable $\kappa = 1$.

population of beetles *increased* when both sun and wind cues were present compared to a wind cue alone. Orientation precision with a sun cue alone was not significantly different from that with both cues present.

At a 75° solar elevation, when adding a wind cue to a sun cue, there was a significant decline in orientation precision ($p < 0.01$, Wilcoxon signed-rank test; figure 7a(v), $n = 20$). Similarly, orientation precision was reduced when a sun cue was added in combination with an existing wind cue ($p < 0.05$, Wilcoxon signed-rank test; figure 7a(vi), $n = 20$). By

contrast, when removing the wind cue and leaving the sun cue in place—or vice versa—a significant increase in orientation precision was observed ($p < 0.01$, Wilcoxon signed-rank test; figure 7a(vii), a(viii), $n = 20$). In summary, at a 75° solar elevation, the orientation precision of the population of beetles *decreased* when both sun and wind cues were present. Two cues reduced orientation precision compared to one.

By comparing the sun-only conditions in figure 7 (a(i) and a(v), initial), we can see that the orientation precision does not

change across the different elevations. This indicates that the reliability of the sun cue is not affected by this change in elevation. This is largely consistent with the findings of [9]; while they reported a slight reduction in precision from 45° to 75°, performance does not degrade significantly until elevations of 80° and above. Given that: (i) changing the elevation does not appear to change the reliability of the sun cue, (ii) changing the elevation from 45° to 75° changes the influence of a sun cue [9], and (iii) orientation precision decreases with multiple cues when the sun is at high elevations, it would appear that reliability is not the key component of cue influence for these beetles.

We suggest that the beetles use cue contrast in order to determine cue weight. Loosely, we reason that the contrast of a sun cue (intensity contrast between the solar and anti-solar hemispheres) drops as elevation increases which changes the relative weight compared to a wind cue with constant speed. As the relative contrast of the sun cue drops, the weight relationship changes but the reliabilities do not, meaning that the wind will receive a higher relative weight than the sun. As the wind is probably a less reliable cue (difficulty in reliable detection and turbulence), this weight change leads to a decrease in orientation performance. Note that our ersatz sun stimulus is monochromatic and its intensity is not varied over different elevations.

(ii) Simulated behaviour

Using our neural circuit, we simulated two weighting strategies: weight-by-reliability and weight-by-contrast. Under weight-by-reliability, relative weight (relative sinusoidal amplitude) is given by the relative reliability of the cues:

$$w_1 = \frac{\kappa_1}{\kappa_1 + \kappa_2} \quad w_2 = \frac{\kappa_2}{\kappa_1 + \kappa_2} \quad (3.1)$$

with $\kappa_1 = 4$ and $\kappa_2 = 1$ ($w_1 = 0.8$, $w_2 = 0.2$). Under weight-by-contrast the relative weight of the cue is set arbitrarily as the inverse of the relative reliability. We are not proposing a specific relationship between the two, only that they are not necessarily coupled.

In each case we simulate four scenarios, mimicking the behavioural assay:

- (i) single reliable cue → add unreliable;
- (ii) single unreliable cue → add reliable;
- (iii) both cues → remove unreliable; and
- (iv) both cues → remove reliable.

The agents roll 10 times under the initial condition and 10 times under the test condition; note that the R → E-PG mapping is not cleared between the initial and test conditions. This means that the order in which conditions are presented will affect the outcome. The results are shown in figure 7b; *p*-values indicate significance levels for a Wilcoxon signed-rank test between the initial and test conditions.

Following our reasoning above, at mid-elevations, the relative contrast and relative reliability of the two cues may be similar (the beetle data should appear to be weighted by reliability). At high elevations, the relative contrast of the light cue will drop but the relative reliability will remain unchanged, leading the unreliable wind cue to get more weight than it should (the beetle data should appear to be weighted by contrast).

Where cue weight is dictated by relative reliability (figure 7b, top row), our model generally matches the prediction that two cues should be better than one (median increase in orientation precision). While the change in simulated results is more extreme, we generally see the same pattern as in the beetles at 45° elevation (figure 7a, top row). Where we weight cues by (hypothetical) relative contrast, figure 7b(v), b(vii), and b(viii) indicate that two cues would be worse than one (decrease in median orientation precision). Again, the simulated mean vector length distributions are more extreme in their differences, however, the general pattern matches the beetle data at 75° elevation (figure 7a, bottom row). This reflects our broad expectation if cues are weighted according to relative contrast, as opposed to relative reliability. We propose that the beetles are weighting cues according to contrast, which aligns with reliability at 45° elevation, however, as elevation rises, contrast and reliability diverge leading to decreased multi-cue performance.

We note two anomalies in our simulated results. In figure 7b(iii), removing a cue makes the agents more precise where we would expect it to cause a decrease in performance. It is not clear why this is the case, but the effect is consistent across different random initialisations. Figure 7b(vi) presents another anomaly, the simulation reports an increase in precision where the analogous experiment from the beetle data (figure 7a(vi)) shows a decrease. Even when weighting cues arbitrarily, we would still expect adding a reliable cue to increase precision (in line with the simulation result). We are not sure why the beetles experience a decrease in orientation precision in this case. One possibility is that the biological data itself is highly variable. If we compare the 'initial' conditions in figure 7a(ii), a(vi) (where the beetles experienced wind in isolation), it is clear that the beetles are far more precise in the wind-only condition in figure 7a(vi), despite the fact that these conditions should be identical.

4. Discussion

The insect head direction circuit, specifically the EB, has been previously suggested to house a flexible multimodal compass [13,21]. Here we have provided a functional model which demonstrates that the EB can perform multimodal cue integration as vector summation. Importantly, the plasticity in mapping different cues (R neuron activity) to the EB during rotation establishes a common frame of reference for the vectors. The plasticity can also influence the relative contribution to the vector sum of the different cues. However, we show that this does not lead to an emergent reliability-based influence of the cues on the behaviour; some additional mechanism would be needed to set the input weights. This prompted an exploration into what determines cue weight in our model species, the dung beetle *K. lamarcki*. Contrary to previous studies [9,22,26], we found that these beetles do not appear to weight cues according to their reliability. We suggest relative cue contrast may provide an alternative explanation of these results.

(a) Modelling landscape

Conceptually, none of the individual elements in the model we present are new [8,13,17,21,41]. Nevertheless, we are the first to provide a relatively complete implementation, and in particular, to combine adaptive synaptic connections for

R-neurons with multimodal cue integration. Most previous models do not investigate multi-cue behaviour [8,20,28,42]. The closest work is provided by [43] (extended in [44]) which includes two populations of neurons with Gaussian activation profiles which input to an integration layer which itself forms a ring-tractor. Plasticity between the input and integration (compass) neurons is not included. This is noted as a strength of the model (conceptual simplicity), however, it does not reflect known phenomena in fruit flies and dung beetles [16,20,22]. Such rigidity also creates functional problems with spatially separated cues which automatically form conflicts unless their relationship is innately encoded.

(b) Hebbian learning facilitates vector summation

Hebbian learning is often suggested for the construction of $R \rightarrow E$ -PG mappings [16,17,19,20,28,41]. We show here that such plasticity does not interfere with the vector summation properties of the circuit provided the learning is structured by being linked to periods of increased angular velocity [21,28]. We also note that, in order to build the relationship between perceived cue motion and angular velocity, cue input cannot drive the compass during learning. We found this to be true during model construction but it is also fairly intuitive; if R input drives the E-PG signal and $R \rightarrow E$ -PG connections are updated in a Hebbian fashion, the active R neurons will continually remap to the active E-PGs (see the electronic supplementary material, continuous learning). Other models have assumed that R neurons do not provide any input to the E-PGs during learning [28,42]. We instead suggest that the balance of influence between the self-motion and cue input in the E-PGs could be flexibly regulated by angular velocity in the same manner as the plasticity between R and E-PG neurons.

(c) Cue contrast as a proxy for reliability

Previous work in insect cue integration has suggested that cue influence is determined according to reliability [7–9,22,26]. Our conflict simulations (figure 6) indicate that cue reliability, while affecting the formation of R-to-E-PG connections, does not have a significant effect on cue influence. Further, our behavioural results and accompanying simulations clearly indicate that cue weight is not determined by reliability for the dung beetle *K. lamarcki*, contrary to previous claims [9,22,26].

In previous dung beetle cue integration studies, it is not always clear that cue reliability was manipulated. Reliability estimation usually requires performing a large number of trials with the same individual, which is often impractical in insects. Khaldy *et al.* [26] manipulated the intensity of a sun cue but did not provide evidence that this significantly affected reliability of orientation. In [22] the assumption that higher sun elevations are less reliable is justified by reference to [24] in which the difference between two exit angles is taken to be a measure of orientation precision. In [9] the orientation precision (estimated from multiple exits) of a beetle under an ersatz sun did not appear to be reduced until very high elevations ($> 80^\circ$). In this paper, sun elevations for which precision did not significantly differ nevertheless had different effects on cue integration: beetles orient to the sun for a 45° elevation but the wind for 75° elevation, which suggests reliability is not determining cue weight.

On the other hand, changing elevation or intensity, as done in these studies, would be expected to change the *contrast* of the sun cue. Indeed it is possible that the directional cue provided from the sun actually comes from intensity gradient it produces across the sky, which beetles are known to use for orientation [45–47]. Cue contrast can contribute to cue reliability by reducing detection error (figure 1), but (unlike reliability) is directly available in the instant. A neural circuit that has evolved to weight cues by contrast would in most natural situations obtain a good proxy of weighting by reliability.

(d) Future work

(i) R neuron representation

A recent connectomic analysis observed that R neurons inhibit all other R neurons in the same group (e.g. all R1 neurons inhibit all other R1 neurons) [13]. These connections were not included as the data was not available when the initial model was constructed. We plan to investigate whether within-group inhibition supports a sinusoidal activity profile which we assumed for R neuron encoding of cue direction (consistent with [19] and [8]) but which has not been directly shown. By analogy to the EB [13] inhibition could encourage a sinusoidal shape for different input profiles. A second possibility is that inhibition could provide a mechanism for reliability-based weighting, keeping the peak of an activity bump relatively stable, while modulating the amplitude as a result of noise (see [8]). R neuron groups also inhibit each other, which could implement a hard-wired, yet flexible cue hierarchy (which sufficient weighting could overcome); alternatively, this could partially implement our R activity normalization scheme.

(ii) Questions in dung beetles

As solar elevation increases, dung beetle orientation precision remains relatively stable then decreases rapidly at high elevations [9]. This is what we might expect if the contrast at lower elevations remains well above a detection threshold. A contrast assay [47] could be used to check what solar to anti-solar contrast corresponds to the minimum usable contrast threshold in *K. lamarcki*. It would then be useful to try and determine how the beetles compute solar to anti-solar contrast. There may also be a way to measure wind cue contrast by examining antennal displacement for different wind speeds (and directions).

While our data suggest that dung beetles do not fully weight cues by reliability, our modelling suggests that reliability could have some effect on cue influence. It would be useful to examine this explicitly in the animal, e.g. by adding azimuthal variation to a sun cue.

5. Conclusion

Cue integration research (our own included [9]) tends to assume Bayesian reasoning: a cue estimate includes information about the variance of the estimate (reliability), which is used to weight the cue [1,3]. However, the meaningfulness of a Bayes-optimal result is questionable [48] and the use of Bayesian modelling is disputed [49–53]. In this small, behavioural/modelling example, the Bayesian assumption was counterproductive; only by leaving reliability behind were we able to make sense of our data. Future behavioural data

may be better understood by first asking what quality of a cue is actually changed by an experimental manipulation, rather than assuming reliability is involved. This could expose a host of efficient, elegant solutions to otherwise computationally intensive problems.

Ethics. Animal collection and behavioural experiments were conducted in accordance with the Swedish and South African guidelines for the use of animals for scientific purposes. Animal care followed the regulations as detailed in the EU Directive 2010/63/EU.

Data accessibility. All raw data and code are available in the public GitHub repository accompanying this paper (https://github.com/refmitchell/CX_cue_integration_model/tree/v1.0.0 [<https://doi.org/10.5281/zenodo.7788258>] [54]). Follow the link to access this repository and download our code/data. Additional experiments in our electronic supplementary material [55] use data previously published in Dacke *et al.* [22] and Shaverdian *et al.* [9] (used with author permission). This data is included in our repository for completeness (to allow others to replicate our analysis) and the sources are clearly marked.

Software implementation details and additional modelling results can be found in the electronic supplementary material [55].

References

- Ernst MO, Bühlhoff HH. 2004 Merging the senses into a robust percept. *Trends Cogn. Sci.* **8**, 162–169. (doi:10.1016/j.tics.2004.02.002)
- Rohde M, Dam Lv, Ernst MO. 2016 Statistically optimal multisensory cue integration: a practical tutorial. *Multisensory Res.* **29**, 279–317. (doi:10.1163/22134808-00002510)
- Knill DC, Pouget A. 2004 The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci.* **27**, 712–719. (doi:10.1016/j.tins.2004.10.007)
- Murray RF, Morgenstern Y. 2010 Cue combination on the circle and the sphere. *J. Vis.* **10**, 1–11.
- Stone T *et al.* 2017 An anatomically constrained model for path integration in the bee brain. *Curr. Biol.* **27**, 3069–3085. (doi:10.1016/j.cub.2017.08.052)
- Le Moël F, Stone T, Lihoreau M, Wystrach A, Webb B. 2019 The central complex as a potential substrate for vector based navigation. *Front. Psychol.* **10**, 690. (doi:10.3389/fpsyg.2019.00690)
- Hoinville T, Wehner R. 2018 Optimal multiguideance integration in insect navigation. *Proc. Natl Acad. Sci. USA* **115**, 2824–2829. (doi:10.1073/pnas.1721668115)
- Kutschireiter A, Basnak MA, Wilson RI, Drugowitsch J. 2023 Bayesian inference in ring attractor networks. *Proc. Natl Acad. Sci. USA* **120**, e2210622120. (doi:10.1073/pnas.2210622120)
- Shaverdian S, Dirlik E, Mitchell R, Tocco C, Webb B, Dacke M. 2022 Weighted cue integration for straight-line orientation. *Isience* **25**, 105207. (doi:10.1016/j.isci.2022.105207)
- Lyu C, Abbott LF, Maimon G. 2022 Building an allocentric travelling direction signal via vector computation. *Nature* **601**, 92–97. (doi:10.1038/s41586-021-04067-0)
- Heinze S. 2017 Unraveling the neural basis of insect navigation. *Curr. Opin. Insect Sci.* **24**, 58–67. (doi:10.1016/j.cois.2017.09.001)
- Hanesch U, Fischbach KF, Heisenberg M. 1989 Neuronal architecture of the central complex in *Drosophila melanogaster*. *Cell Tissue Res.* **257**, 343–366. (doi:10.1007/BF00261838)
- Hulse BK *et al.* 2021 A connectome of the *Drosophila* central complex reveals network motifs suitable for flexible navigation and context-dependent action selection. *Elife* **10**, e66039. (doi:10.7554/eLife.66039)
- Seelig JD, Jayaraman V. 2015 Neural dynamics for landmark orientation and angular path integration. *Nature* **521**, 186–191. (doi:10.1038/nature14446)
- Hardcastle BJ, Omoto JJ, Kandimalla P, Nguyen BC, Keleş MF, Boyd NK, Hartenstein V, Frye MA. 2021 A visual pathway for skylight polarization processing in *Drosophila*. *Elife* **10**, e63225. (doi:10.7554/eLife.63225)
- Fisher YE, Lu J, D'Alessandro I, Wilson RI. 2019 Sensorimotor experience remaps visual input to a heading-direction network. *Nature* **576**, 121–125. (doi:10.1038/s41586-019-1772-4)
- Okubo TS, Patella P, D'Alessandro I, Wilson RI. 2020 A neural network for wind-guided compass navigation. *Neuron* **107**, 924–940. (doi:10.1016/j.neuron.2020.06.022)
- Seelig JD, Jayaraman V. 2013 Feature detection and orientation tuning in the *Drosophila* central complex. *Nature* **503**, 262–266. (doi:10.1038/nature12601)
- Fisher YE. 2022 Flexible navigational computations in the *Drosophila* central complex. *Curr. Opin Neurobiol.* **73**, 102514. (doi:10.1016/j.conb.2021.12.001)
- Kim SS, Hermundstad AM, Romani S, Abbott LF, Jayaraman V. 2019 Generation of stable heading representations in diverse visual scenes. *Nature* **576**, 126–131. (doi:10.1038/s41586-019-1767-1)
- Fisher YE, Marquis M, D'Alessandro I, Wilson RI. 2022 Dopamine promotes head direction plasticity during orienting movements. *Nature* **612**, 316–322. (doi:10.1038/s41586-022-05485-4)
- Dacke M, Bell AT, Foster JJ, Baird EJ, Strube-Bloss MF, Byrne MJ, El Jundi B. 2019 Multimodal cue integration in the dung beetle compass. *Proc. Natl Acad. Sci. USA* **116**, 14 248–14 253. (doi:10.1073/pnas.1904308116)
- Dacke M, El Jundi B. 2018 The dung beetle compass. *Curr. Biol.* **28**, R993–R997. (doi:10.1016/j.cub.2018.04.052)
- Dacke M, El Jundi B, Smolka J, Byrne M, Baird E. 2014 The role of the sun in the celestial compass of dung beetles. *Phil. Trans. R. Soc. B* **369**, 20130036. (doi:10.1098/rstb.2013.0036)
- Khaldy L, Tocco C, Byrne M, Dacke M. 2021 Compass cue integration and its relation to the visual ecology of three tribes of ball-rolling dung beetles. *Insects* **12**, 526. (doi:10.3390/insects12060526)
- Khaldy L, Foster JJ, Yilmaz A, Belušić G, Gagnon Y, Tocco C, Byrne MJ, Dacke M. 2022 The interplay of directional information provided by unpolarised and polarised light in the heading direction network of the diurnal dung beetle *Kheper lamarcki*. *J. Exp. Biol.* **225**, jeb243734. (doi:10.1242/jeb.243734)
- El Jundi B, Warrant EJ, Byrne MJ, Khaldy L, Baird E, Smolka J, Dacke M. 2015 Neural coding underlying the cue preference for celestial orientation. *Proc. Natl Acad. Sci. USA* **112**, 11 395–11 400. (doi:10.1073/pnas.1501272112)
- Dan C, Hulse BK, Jayaraman V, Hermundstad AM. 2021 Flexible control of behavioral variability mediated by an internal representation of head direction. *bioRxiv* 2021-08. See <https://www.biorxiv.org/content/10.1101/2021.08.18.456004v2>.
- Baird E, Byrne MJ, Smolka J, Warrant EJ, Dacke M. 2012 The dung beetle dance: an orientation behaviour? *PLoS ONE* **7**, e30211. (doi:10.1371/journal.pone.0030211)

30. El Jundi B, Foster JJ, Khaldy L, Byrne MJ, Dacke M, Baird E. 2016 A snapshot-based mechanism for celestial orientation. *Curr. Biol.* **26**, 1456–1462. (doi:10.1016/j.cub.2016.03.030)
31. Pisokas I, Heinze S, Webb B. 2020 The head direction circuit of two insect species. *Elife* **9**, e53985. (doi:10.7554/eLife.53985)
32. Goulard R, Buehlmann C, Niven JE, Graham P, Webb B. 2021 A unified mechanism for innate and learned visual landmark guidance in the insect central complex. *PLoS Comput. Biol.* **17**, e1009383. (doi:10.1371/journal.pcbi.1009383)
33. Honkanen A, Adden A, da Silva Freitas J, Heinze S. 2019 The insect central complex and the neural basis of navigational strategies. *J. Exp. Biol.* **222**(Suppl_1), jeb188854. (doi:10.1242/jeb.188854)
34. Immonen EV, Dacke M, Heinze S, El Jundi B. 2017 Anatomical organization of the brain of a diurnal and a nocturnal dung beetle. *J. Comp. Neurol.* **525**, 1879–1908. (doi:10.1002/cne.24169)
35. Kim SS, Rouault H, Druckmann S, Jayaraman V. 2017 Ring attractor dynamics in the *Drosophila* central brain. *Science* **356**, 849–853. (doi:10.1126/science.aal4835)
36. El Jundi B, Warrant EJ, Pfeiffer K, Dacke M. 2018 Neuroarchitecture of the dung beetle central complex. *J. Comp. Neurol.* **526**, 2612–2630. (doi:10.1002/cne.24520)
37. Omoto JJ, Nguyen BC, Kandimalla P, Lovick JK, Donlea JM, Hartenstein V. 2018 Neuronal constituents and putative interactions within the *Drosophila* ellipsoid body neuropil. *Front. Neur. Circ.* **12**, 103. (doi:10.3389/fncir.2018.00103)
38. Turner-Evans D, Wegener S, Rouault H, Franconville R, Wolff T, Seelig JD, Druckmann S, Jayaraman V. 2017 Angular velocity integration in a fly heading circuit. *Elife* **6**, e23496. (doi:10.7554/eLife.23496)
39. Zenke F, Gerstner W, Ganguli S. 2017 The temporal paradox of Hebbian learning and homeostatic plasticity. *Curr. Opin. Neurobiol.* **43**, 166–176. (doi:10.1016/j.conb.2017.03.015)
40. Allaire J. 2012 RStudio: integrated development environment for R. Boston, MA. See <https://www.rstudio.org/>.
41. Green J, Maimon G. 2018 Building a heading signal from anatomically defined neuron types in the *Drosophila* central complex. *Curr. Opin. Neurobiol.* **52**, 156–164. (doi:10.1016/j.conb.2018.06.010)
42. Cope AJ, Sabo C, Vasilaki E, Barron AB, Marshall JA. 2017 A computational model of the integration of landmarks and motion in the insect central complex. *PLoS ONE* **12**, e0172325. (doi:10.1371/journal.pone.0172325)
43. Sun X, Mangan M, Yue S. 2018 An analysis of a ring attractor model for cue integration. In *Biomimetic and biohybrid systems: 7th International Conference, Living Machines 2018, Paris, France, 17–20 July 2018, Proceedings 7 2018*, pp. 459–470. New York, NY: Springer International Publishing.
44. Sun X, Yue S, Mangan M. 2020 A decentralised neural model explaining optimal integration of navigational strategies in insects. *Elife* **9**, e54026. (doi:10.7554/eLife.54026)
45. El Jundi B, Smolka J, Baird E, Byrne MJ, Dacke M. 2014 Diurnal dung beetles use the intensity gradient and the polarization pattern of the sky for orientation. *J. Exp. Biol.* **217**, 2422–2429.
46. El Jundi B, Foster JJ, Byrne MJ, Baird E, Dacke M. 2015 Spectral information as an orientation cue in dung beetles. *Biol. Lett.* **11**, 20150656. (doi:10.1098/rsbl.2015.0656)
47. Foster JJ, El Jundi B, Smolka J, Khaldy L, Nilsson DE, Byrne MJ, Dacke M. 2017 Stellar performance: mechanisms underlying Milky Way orientation in dung beetles. *Phil. Trans. R. Soc. B* **372**, 20160079. (doi:10.1098/rstb.2016.0079)
48. Scarfe P. 2022 Experimentally disambiguating models of sensory cue integration. *J. Vis.* **22**, 5. (doi:10.1167/jov.22.1.5)
49. Bowers JS, Davis CJ. 2012 Bayesian just-so stories in psychology and neuroscience. *Psychol. Bull.* **138**, 389. (doi:10.1037/a0026450)
50. Griffiths TL, Chater N, Norris D, Pouget A. 2012 How the Bayesians got their beliefs (and what those beliefs actually are): comment on Bowers and Davis (2012). *Psychol. Bull.* **138**, 415–422. (doi:10.1037/a0026884)
51. Bowers JS, Davis CJ. 2012 Is that what Bayesians believe? Reply to Griffiths, Chater, Norris, and Pouget (2012). *Psychol. Bull.* **138**, 423–426. (doi:10.1037/a0027750)
52. Rahnev D, Denison RN. 2018 Suboptimality in perceptual decision making. *Behav. Brain Sci.* **41**, e223. (doi:10.1017/S0140525X18000936)
53. Colombo M, Seriès P. 2012 Bayes in the brain—on Bayesian modelling in neuroscience. *Br. J. Philos. Sci.* **63**, 697–723. (doi:10.1093/bjps/axr043)
54. Mitchell R. 2023 refmitchell/CX_cue_integration_model: Initial public release. Zenodo. (doi:10.5281/zenodo.7788258)
55. Mitchell R, Shaverdian S, Dacke M, Webb B. 2023 A model of cue integration as vector summation in the insect brain. Figshare. (doi:10.6084/m9.figshare.c.6688818)

Correction



Cite this article: Mitchell R, Shaverdian S, Dacke M, Webb B. 2023 Correction to: 'A model of cue integration as vector summation in the insect brain' (2023) Mitchell *et al.* *Proc. R. Soc. B* **290**: 20231993.
<https://doi.org/10.1098/rspb.2023.1993>

Received: 6 September 2023

Accepted: 6 September 2023

Subject Category:

Behaviour

Subject Areas:

behaviour, computational biology

Keywords:

vector, cue integration, plasticity, contrast, reliability, neural model

Author for correspondence:

Robert Mitchell

e-mail: r.mitchell@ed.ac.uk

Correction to: 'A model of cue integration as vector summation in the insect brain' (2023) Mitchell *et al.*

Robert Mitchell, Shahrzad Shaverdian, Marie Dacke and Barbara Webb

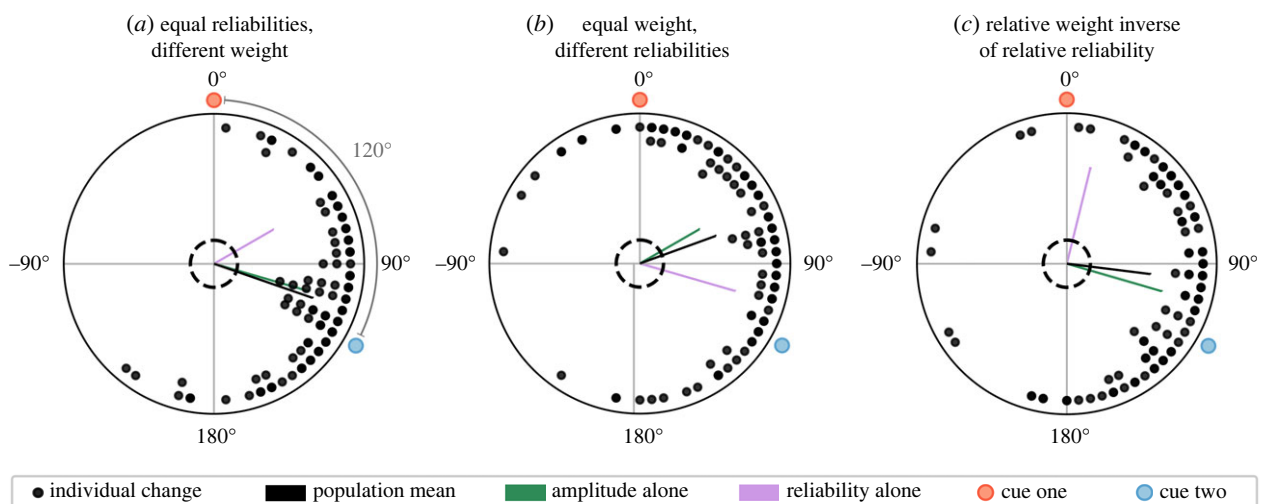
ORCID RM, 0000-0002-6396-0995; SS, 0000-0001-6076-5310; MD, 0000-0001-6444-7483; BW, 0000-0002-8336-6926

Proc. R. Soc. B **290**, 20230767 (Published online 28 June 2023). (doi:10.1098/rspb.2023.0767)

In the original version of this article [1], fig. 6 contained data which were based on a flawed simulation script. This script has been fixed in our code repository, and the corrected version of fig. 6 is included here. The data still support the original point that cue weight primarily determines cue influence. The authors apologize for any confusion which arose as a result of this minor error.

Reference

1. Mitchell R, Shaverdian S, Dacke M, Webb B. 2023 A model of cue integration as vector summation in the insect brain. *Proc. R. Soc. B* **290**, 20230767. (doi:10.1098/rspb.2023.0767)



Replacement for fig. 6 in [1]. Effect of weight and reliability on cue influence in behavioural simulation of cue conflict. The population mean vector (black) falls closer to the theoretical vector sum where magnitudes are determined by cue weight (green) than where magnitudes are determined by reliability (magenta). Dashed rings indicate the threshold for significance using a Rayleigh test ($p < 0.05$).

Downloaded from https://royalsocietypublishing.org/ on 22 October 2023

Conclusion

Separating weight, reliability, and influence

In preparing this work, it became almost immediately apparent that the idea of cue ‘weight’ discussed in (Paper I) gives an oversimplified view.

Cue weight is generally used to describe the final effect on behaviour but also connotes a term in an equation (i.e. a computational hypothesis). For the neural implementation, we had to tie this to something specific. In the vector sum, weight is given by the magnitude of the vector representing the cue (Murray and Morgenstern, 2010), thus the natural correlate was the amplitude of the sinusoidal signal in the R neurons. We therefore we defined cue weight as the amplitude of this sinusoidal population code. While it is often suggested that cue weight is set according to cue reliability, it is worth noting that there is no obvious mechanism for reliability to be encoded in this sinusoidal amplitude. That said, this is not the only factor which can affect behaviour. The synapses between R and EPG neurons are plastic and will be affected by the weight (sinusoidal amplitude) but also by cue variability. For this model, variability may emerge from the cue itself, the sensor, or even the processing chain from sensation to central brain. This variability is captured by the established definition of cue reliability. We also include the idea of cue influence. The interaction of cue weight and synaptic strength (which is also affected by reliability) will change the final behavioural outcome but it is not immediately clear which factors are critical in generating the behavioural change. This conceptual separation is quite useful, as it allows us to consider where, how, and to what extent different factors end up impacting behaviour.

The separability of weight, reliability, and influence is an important conceptual contribution and it is worth emphasising. The idea that R → EPG mappings form a substrate for reliability-based influence of orientation cues has begun to appear in the literature, with the suggestion that this would be ‘ideal’ (Wilson, 2023). While it is obviously true that cue reliability will affect cue influence via corrupted synaptic maps, it is unclear how consistent this effect will be at behaviourally relevant levels of noise. Furthermore, the modelling also shows that this is not the only factor which can affect cue influence. A change in weight (sinusoidal amplitude in an R neuron population) seems to have a greater, more consistent effect on cue influence than changes in reliability.

It is obvious that the model and simulations are limited; this is but one possible parameterisation of this circuit and there is some detail omitted. I am not claiming that reliability will not affect cue influence in insects at all and I think it is quite likely that this will vary across

species. Nevertheless, I think we can say with reasonable confidence that reliability-based cue influence is not an emergent property of the circuit. Moreover, other factors can play a significant role in cue influence and this allows us to explain our perplexing behavioural results. For example, the key bar for reliability-based influence is an increase in orientation precision in a multi-cue scenario (Jones, 2016). Our behavioural work alone shows that the beetles get worse with multiple cues which rules out reliability-driven cue combination. It is also worth referring back to (Paper I) where cue reliability (as measured) was not a good predictor of final cue influence. These results are difficult to explain if we argue that the learning mechanism automatically sets cue influence according to reliability (Wilson, 2023). It is therefore important to consider all factors which contribute to cue influence.

Ideal behaviour

The proposal from Wilson (2023) is attractive from the perspective that it should result in ideal use of noisy information in the sense that it minimises the variance of the final orientation estimate. It is worth reiterating that fixation on this narrow conceptualisation of ‘ideal’ has been unhelpful in human perceptual literature (Rahnev and Denison, 2018) and we should be careful not to make the same mistake in insects. What is ideal depends on the animal, behaviour, and evolutionary history. For a dung beetle, obtaining the best possible orientation precision may not matter (the conferred behavioural advantage may not be great enough to generate selection pressure). Conversely, for a central place forager precision may matter a great deal; for an ant, relatively minor compass errors could mean not finding her nest which means death. This is relevant here as the models and definitions prevalent in the human psychophysical literature have been used as the basis for previous claims in insects (see Chapter 3).

Future work should consider the resultant effect on behaviour with *behaviourally relevant* levels of noise. It is possible that, during normal behaviour, cues simply are not noisy enough to alter the synaptic mappings in a way which produces a consistent behavioural effect (this could depend on the sensory ecology of the species). It is also important that future work considers any other potential contributors to cue influence (e.g. R neuron response amplitude, see above) and how these interact with any changes in the synaptic map. Noise could affect synaptic maps, but this may not be the primary contributor to cue influence (as suggested by the model).

Contrast as a proxy for reliability

In the paper, we argue that cue contrast could act as a functional proxy for cue reliability but it may not be completely clear how cue reliability is affected by manipulating solar elevation or wind speed, and how contrast allows for approximation of reliability. This arose as a way to explain why the dung beetle would give greater influence to an apparently less reliable cue, resulting in a decrease in orientation performance in a multi-cue scenario.

Previous work in dung beetle orientation has argued that, as solar elevation increases, the reliability of the sun as an orientation cue drops (Dacke et al., 2014) and that dung beetles use this change in reliability in order to determine cue influence (Paper I). Dung beetles appear to detect the sun as a wide-field intensity gradient cue (though specifically how this is done is not yet known) (el Jundi et al., 2014). This can be thought of simply as determining which half of your visual field contains the sun; however this requires being able to detect an intensity difference across your field of view (we have good reason to believe the beetles do not perform view-matching, see Chapter 2). As the sun approaches the zenith, this will become more difficult as more light bleeds across the centre of the visual field, reducing the contrast between the solar and anti-solar portions of the visual field. Presumably, the dung beetles do not have perfect sensors and therefore there will be some minimum contrast threshold which is required across their visual field in order to use the sun for orientation (such a threshold has been quantified in the nocturnal *Scarabaeus satyrus* (Foster et al., 2017)). As the contrast approaches this threshold, the sensor may start to experience detection errors (a decrease in reliability), resulting in a rapid deterioration in orientation performance. This appears to be what we saw in (Paper I), orientation precision stays relatively stable until elevations above 80° (Figure 1 in that paper). In (Paper II) we argue that such detection errors are the main contributor to the reliability of the cue and therefore cue contrast could be used to approximate reliability without having to actually measure variance over time.

It should be noted that cue contrast may not directly correlate with cue reliability over the full range of elevations. For example, a 75° sun cue (in theory) generates less contrast than a 45° sun cue, even if the beetles are able to orientate equally well under each of these conditions (see results from (Paper I) and the sun-only precision conditions from (Paper II)). However, *Kheper lamarcki* is much less active during the highest solar elevations (Smolka et al., 2016), so for the sun in particular, cases where cue contrast deviates from cue reliability may not be behaviourally relevant. This could explain why the beetle appears to give

more weight to a less reliable cue in the tent setup. We propose that the animal has ended up with a heuristic which allows it to use multiple cues to improve the precision of its orientation estimate, without actually needing to estimate the degree of noise in each modality. In the tent however, we break this heuristic by exposing the beetle to higher elevations than it would normally experience. This leads to a divergence between cue contrast and cue reliability, and results in an errant up-weighting of a noisy cue, which in turn leads to a decrease in overall orientation precision

While this discussion has focused on solar elevation it is clear that a similar heuristic work well for a wind cue. How beetles detect wind direction is not yet known, however fruit flies appear to compare deflections of their antennae (Suver et al., 2019; Okubo et al., 2020). Higher wind speeds should generate greater antennal displacements, which should make wind direction easier to detect (up to a point). Wind direction as measured at a specific point by a specific sensor may also be variable due to turbulence, this is currently under investigation.

Sinusoidal R neuron activity profile

In this work we assumed that R neuron populations took on a sinusoidal activity profile which encoded the agent’s orientation with respect to a specific cue modality. This assumption is reasonable in the context of the dung beetle and the simulated experiments (with a single visual cue), however it is worth noting that R neuron population dynamics are not currently known. Visual R neurons have previously been described as feature detectors (Seelig and Jayaraman, 2013), giving the impression that an R neuron population may have multiple peaks and the curve encodes the current scene. While this would break proposed link to the vector sum, the proposed cue integration mechanism should still function. Previous work from Kim et al. (2019) has shown that, even where R neurons have a more complex population activity profile, the (anti-)Hebbian learning mechanism can still map each of these peaks to a single peak in the EPG population. The single peak in the EPGs is facilitated by lateral inhibition and local excitation via the $\Delta 7$ neurons and PEN/PEG neurons respectively.

A recent connectomic description of the fruit fly central complex has uncovered similar lateral inhibition in R neuron populations (Hulse et al., 2021). This inhibition could facilitate ‘winner-take-all’ dynamics, here meaning that one prominent feature (or set of features) in the environment could inhibit other R neurons, causing a complex scene to be reduced to a (roughly) sinusoidal encoding.

Sinusoidal amplitude modulation in R neurons

In this work, we assumed that R neuron response amplitude (cue weight) could be regulated by some property of the sensory scene. We suggested cue contrast as this explanation made sense of our (and previous) data, but is not currently known if or how R neuron population responses can change or what property could regulate such change. This seems testable in beetles and could be supplemented by work in fruit flies. In beetles, we could investigate R neuron activities using electrophysiology. For a sun cue, we could theoretically locate the receptive field of a particular cell and then alter the apparent elevation of the stimulus to see if this causes a change in the response of the cell. A similar assay may work for wind cues.

Complementary testing could be done on populations of R neurons using calcium imaging in fruit flies. In addition, the actual response dynamics of EPGs could be observed with respect to cue conflict. We could look to see if responses mimic the curves given in Figure 4 in (Paper II). In fruit flies, it also seems possible to test the effect of stimulus noise on R neuron activity. If we assume that R neuron population codes are sinusoidal, then it seems possible that the amplitude could be modulated by cue noise via fast lateral inhibition discussed in (Paper II).

The orientation snapshot

[el Jundi et al. \(2016\)](#) previously suggested that the R \rightarrow EPG mappings could store offsets between different cues and therefore provide a substrate for the dung beetle orientation snapshot (see Chapter 2). This capability is clearly demonstrated in the model, confirming the anatomy supports such a function. Unfortunately, we removed this discussion from (Paper II) prior to publication (due to space constraints) without realising how significant this demonstration was. The mappings store the angular offset between different cues and also require rotation (e.g. a dance movement) to form ([Fisher et al., 2022](#)), making them an ideal substrate for the snapshot. This is also discussed in Chapters 7 and 8.

Chapter 6

BeetleBot

We have thus far followed the neuroethological train from descriptive behavioural modelling ([Paper I](#)) to a simulated mechanism ([Paper II](#)), with the obvious next step being a physical biorobot ([Webb and Consi, 2001](#)), operating under the same environmental conditions as the beetles. A *BeetleBot* was constructed early in the project, however, it has seen limited use in this project. The neural substrate for cue integration ([Paper II](#)) ended up being more complicated than was initially envisioned and so simulation was the more practical option for development and testing.

Nevertheless, the BeetleBot has demonstrated its usefulness. The robot served as a valuable testing platform for a prototype polarisation sensor proposed by [Gkanias et al. \(2019\)](#). The robot was used to field-test this sensor in Sardinia and South Africa ([Gkanias et al., in press](#)) (see Appendix F). In addition, a modified version of the neural model from ([Paper II](#)) has been put on the robot and used to demonstrate that the model can account for the dung beetle orientation snapshot, operating under similar experimental conditions to the beetles (see Chapter 7).

This chapter describes the robot platform along with relevant justification for design decisions. A full start-up guide (with simulated and real demo scripts) is available in Appendix E. The full (documented) codebase is available on GitHub ([BeetleBot Software](#)).

6.1 System description

The BeetleBot is based on an older TurtleBot3 (Burger) kit (ROBOTIS, see e-manual ([TurtleBot Manual](#))). The TurtleBot3 is a ROS-based platform which consists of a Raspberry Pi 3B+, OpenCR1.0, two Dynamixel servos and a collection of structural elements which can be configured as appropriate for the intended use case. (The kit also includes a LIDAR unit which I did not use.) The platform uses an operating paradigm whereby there is a Host PC which communicates with the on-board Raspberry

Pi via ROS (Quigley et al., 2009), using a local WiFi network. There is therefore a distinction between software which runs on the robot and software which runs on the Host PC (see (TurtleBot Manual)). All software is written for ROS1.

The TurtleBot3 kit is a particularly convenient mobile robot base. Parts compatibility is guaranteed and the kit comes with a number of built-in sensors including motor encoders and a relatively accurate inertial measurement unit (IMU). There is also software available for basic interaction with motors and built-in sensors. This meant that the hardware was essentially ready to use out of the box. Sensory inputs could be piped to a navigational model and model outputs could be piped to the motors.

All of the provided software is open source which means that it can be changed or checked where documentation is unclear. All of the hardware is also open source; precise measurements are available for all components which proved invaluable when constructing custom 3D printed parts. (This also means that some hardware can be 3D printed in the lab as opposed to being purchased.) The stock hardware is also highly configurable. For example, in Figures 6.1A and B, you may be able to see that the robot internals have been re-arranged as appropriate for each experimental situation. (The internals are packed more tightly in B so that the wind sensor and camera can sit lower to the ground; this meant that there were less modifications required for the beetle experimental setup.) The core Raspberry Pi computer is also very widely used which makes for relatively easy use and troubleshooting. In short, the kit is easy-to-use, highly configurable, highly extensible, and well documented which made it an ideal base for the BeetleBot.

6.1.1 Custom hardware

For information on the kit see (TurtleBot Manual). Here I will discuss all custom hardware which was added in order to use the BeetleBot for different experiments.

Wind sensor In order to test the robot with a real wind stimulus, I designed and fabricated a custom wind sensor. While early attempts were made at antennae inspired wind sensors, this proved difficult to scale up to the size of the robot and so a standard vane + anemometer design was chosen instead (Figure 6.1B).

Both the vane and anemometer assemblies are based on the AS5600 magnetic encoder (AMS-Osram, packaged by Seeed Studio). The encoder detects the relative angle of a nearby magnet, which can be read via I2C (the 'Inter-Integrated Circuit' protocol, also called IIC or I²C).




Figure 6.1: **A** The BeetleBot in the field with the polarisation sensor attached. **B** The BeetleBot in the experimental setup from Chapter 7 with the wind sensor attached. **C** The camera unit with wide-angle lens. **D** The custom I2C multiplexer interface board to allow plug-and-play usage of the Adafruit TCA9548A package.

Magnetic encoders were used as they are small, low-cost, and do not introduce friction on the axle. For each assembly, the encoder is fixed inside a 3D printed housing. These encoder housings are designed to be placed in opposition to each other so that the central axles of the vane and the anemometer are aligned (i.e. they both measure from the same point in space). The housing has four mounting holes for fixing the sensor to the TurtleBot, a bearing mount which allows a mechanosensory part (vane or anemometer) to be fixed in place over the encoder, and four mounting holes for metal standoffs which connect to the opposite plate. In addition, two holes are included which allow a camera mounting plate to be fixed to the top of the wind sensor.

Each mechanosensory part (vane and anemometer) is designed to be 3D printed as a single piece with a built-in central axle. This axle is passed through a bearing, and fixed from the opposite side by another 3D printed piece which houses a magnet. The vane also includes a slot for counterweights to balance the nose and tail over the bearing (trimmed M2.5 bolts were used as the counterweight). Each assembly (vane/anemometer, bearing, magnet housing, and magnet) is then glued together and glued into the bearing mount on the encoder housing. Example construction for an anemometer unit is shown in Figure 6.2.

I adapted a pre-existing Arduino software library to allow the magnetic encoders to be read directly from the Raspberry Pi GPIO pins (see [BeetleBot Software](#)).

I2C multiplexer + breakout As the wind sensor used two identical encoders (whose I2C addresses could not be changed), a multiplexer (TCA9548A; Texas Instruments, packaged by Adafruit) was included between the Raspberry Pi and the encoders. To match interfaces with the encoders, I designed a custom breakout board (Figure 6.1D) which maps the pins from the Adafruit TCA9548A package onto the 4-pin connectors used by the Seeed Studio Grove System (used by the encoders). The PCB was fabricated by Garry Ellard in the Informatics workshop. I then assembled the final board with the multiplexer and connectors. Custom software was written to read from this multiplexer directly from the Raspberry Pi.

Camera A camera does not come as standard with the TurtleBot3 (Burger) kit (though any Raspberry Pi compatible camera or USB webcam could be used). I chose the Arducam 8MP Sony IMX219 camera module (Figure 6.4C) as it is directly compatible with the Raspberry Pi and also comes with an M12 lens mount which allows for a wide-angle (LN020, Arducam) lens to be fitted. The camera is mounted on top of the robot




Figure 6.2: Anemometer construction. Parts are shown in the upper left, going clockwise, we have the bearing, magnet housing, magnet, retaining plate, encoder package, anemometer, and finally the housing. The encoder package is placed in the housing and the retaining plate is placed behind to hold it in place. The anemometer design incorporates an axle over which a bearing is fitted. The bearing is then held in place by the magnet housing. This whole assembly then sits inside the bearing mount on the housing which holds the magnet within range of the encoder. Vane construction is identical except for the mechanical part attached to the magnet (a vane instead of an anemometer).




Figure 6.3: Example dummy dynamixel unit from the BeetleBot. **A** Rear view. **B** Side view. **C** front view. The unit is designed such that the relationship between the mounting holes (visible in B) and the centre of the wheel axle (visible in A and C) is identical to the powered Dynamixel servos which drive the robot.

facing the sky so as to extract orientation information from the intensity pattern across the sky.

Caterpillar track system For some of the field work conducted with the polarisation sensor, I used TurtleBot3 track kit to allow the robot to operate on rough terrain. The tracks themselves are available to purchase from ROBOTIS but they required a trailing axle system which I custom built (Figure 6.3). The trailing axle system is made up of two identical ‘dummy dynamixel’ units. These dummy units are designed so that they have identical measurements to the driving dynamixel servos; they can be mounted in the same way, with the same screws, and centre of the axle is in the same place relative to the mounting points. Similarly, the axle allows the standard TurtleBot3 sprocket wheel to be mounted with the same screws as on the real dynamixel unit.

Each dummy unit is made up of five parts: a main housing, two bearings, a through-axle, and a sprocket wheel. The housing accepts one bearing on each side and the axle passes through both and clips in place; this design locks both bearings in place without requiring glue and also ensures that the axle does not wobble about its axis of rotation. Removing the need for glue meant that the bearings could be re-used if a 3D printed part failed, or could be swapped out without having to re-print the whole unit. The sprocket wheel is then directly mounted to the axle (which features a wheel mounting plate at its base).

Polarisation sensor A prototype polarisation sensor was constructed to test the polarisation compass model from [Gkanias et al. \(2019\)](#). This work ([Gkanias et al., in press](#)) is included as Appendix F.

Eight PCBs with four UV sensitive photodiodes each were provided by the Scottish Microelectronics Centre. Photodiode responses were avail-

able via two analogue-to-digital converters (ADCs) which could be read via I2C. I designed and produced 3D printed housings for these PCBs which would hold small polarising filters over the photodiodes at 0° , 90° , $+45^\circ$, and -45° to allow them to act as polarisation analysers. I then designed and produced a complete mounting unit, which uniformly arranged the analysers around a ring, with an elevation of 45° . This unit was designed to include a camera mount-point and the whole assembly could be directly attached to TurtleBot base. I also made adaptor cables for each unit to connect the PCBs to the I2C multiplexer board described above.

I adapted a pre-existing Arduino library in order to read the outputs from each ADC directly from the Raspberry Pi I2C interface (see ([Beetle-Bot Software](#))). All code to read from and analyse the polarisation analyser units was custom written. While the dataset collected for ([Gkanias et al., in press](#)) was analysed off-line, it is possible to decode and use the sensor data in real-time to perform navigation tasks (though the utility of this is limited by unavoidable sensor latency). For any additional information on the polarisation sensor, please see Appendix F.

6.1.2 Software

The software is divided across nine ROS packages which are designed to split the chain from sensation to action into manageable chunks. This section provides an overview of the software architecture. Each package has its own detailed documentation in the codebase which includes information about specific nodes, usage, topics, overrides, etc. The majority of the software is written in C++ with some Python.

bb_sensors contains any ROS nodes which interact directly with the sensor hardware. Nodes should read raw sensor data and publish it on the network in a usable format. All nodes in this package are designed to be built and run on the Raspberry Pi so additional processing is omitted at this stage. Camera data is captured at low resolution as we are loosely trying to mimic insect vision; this also conveniently reduces the overhead in capturing and publishing the images.

bb_improc provides a (very basic) image processing pipeline which is split into different nodes. Nodes are provided for image masking, Gaussian blurring, and colour channel selection. ROS subscribers and publishers can be fully configured at run time so we can choose which processing steps are applied. The order of processing stages can also be configured to some degree.

bb_detection provides nodes which act as cue detectors. Given some suitably processed sensor data, these nodes extract relevant information and package it as a Cue object which can be used by navigation models downstream. This package also includes an implementation of the compass model from [Gkanias et al. \(in press\)](#) to allow real-time navigation using the polarisation sensor prototype described in [Gkanias et al. \(in press\)](#).

bb_computation contains ROS nodes which act as host processes for insect-inspired navigation models. The package also includes the code which defines the models. Specifically, the central complex model from [Stone et al. \(2017\)](#), the vector memory model from [Le Moël et al. \(2019\)](#), and a modified version of the cue integration model from ([Paper II](#)) (which also includes a steering circuit inspired by [Mussells Pires et al. \(2024\)](#); see Chapter 7). Generally, each of the nodes in this package starts an instance of one of the navigation models and can receive orientation information from any topic which provides a Cue (i.e. anything from `bb_detection`). This generalisation can be useful as it means the underlying models are completely agnostic as to what ‘Cue’ they are following. All of the nodes here are expected to output something which can be interpreted as a movement command; for example, the central complex model from [Stone et al. \(2017\)](#) produces an output which indicates which direction the agent should turn to re-align with its home vector.

bb_locomotion contains a single service node which acts as a bridge between the rest of the codebase and the TurtleBot3 core utilities which control movement. The main reason that this bridge exists is that the central complex path integration and vector memory models constantly output a motor command, but whether this motor command should be applied is contextual (‘Am I homing or foraging?’). We can effectively set the context by adding or removing the bridge. Generally, this also means that the robot will not move on its own unless the locomotion service is active, which is a useful safety guarantee.

bb_util provides utilities which are used across multiple packages. The package also provides one convenience node which extracts yaw information from the complete odometry provided by the TurtleBot. There are two libraries included in this package which I did not write; these libraries, their original authors, and licenses are clearly given in the code repository.

bb_dummy contains a dummy cue node. There can be any number of these operating at once and dummy cues can also be set at any azimuth (relative to the IMU zero position). Dummy cues are updated according to the IMU information provided by the TurtleBot3 core, packaged as Cue objects and then published; allowing them to be used by navigation models. Dummy cues are exceptionally useful test constructs, primarily due to the fact that ROBOTIS provides several simulation packages which mimic the topic outputs for the TurtleBot3 core, including the IMU. This means that all navigation models can be tested in simulation prior to performing any experiments with the real robot.

bb_graphics contains nodes which can animate navigation model state in real time. This is useful for debugging but also provides nice demos (see Appendix E).

bb_launchers acts as a central location for all roslaunch launch files on the Host PC. Given the granularity of the codebase, launch files are useful for succinctly starting specific robot configurations.

6.2 Proposed extensions

There are some possible extensions to the current BeetleBot.

Updates and compatibility Much of the software and hardware on the robot is at or nearing its end of life (see Appendix E). While not of scientific interest, it would be worthwhile investing time bringing everything up to date to increase the longevity of the platform. At present, this would mean migrating all software onto ROS2 and upgrading the Raspberry Pi hardware.

Better visual modelling The data was not (and is not yet) available to construct an accurate eye model for a ball-rolling dung beetle. As such, the BeetleBot simply downsamples and blurs camera data. This is probably not a terrible approximation of a superposition compound eye, however, it would be interesting to construct an accurate eye model which could be used to construct a more realistic beetle view.

An antennal wind sensor While the wind sensor works for the purposes of testing orientation cue integration, it would be interesting to construct an antennae inspired wind sensor. [Suver et al. \(2019\)](#) provide a model for computing direction from antennal displacement and I

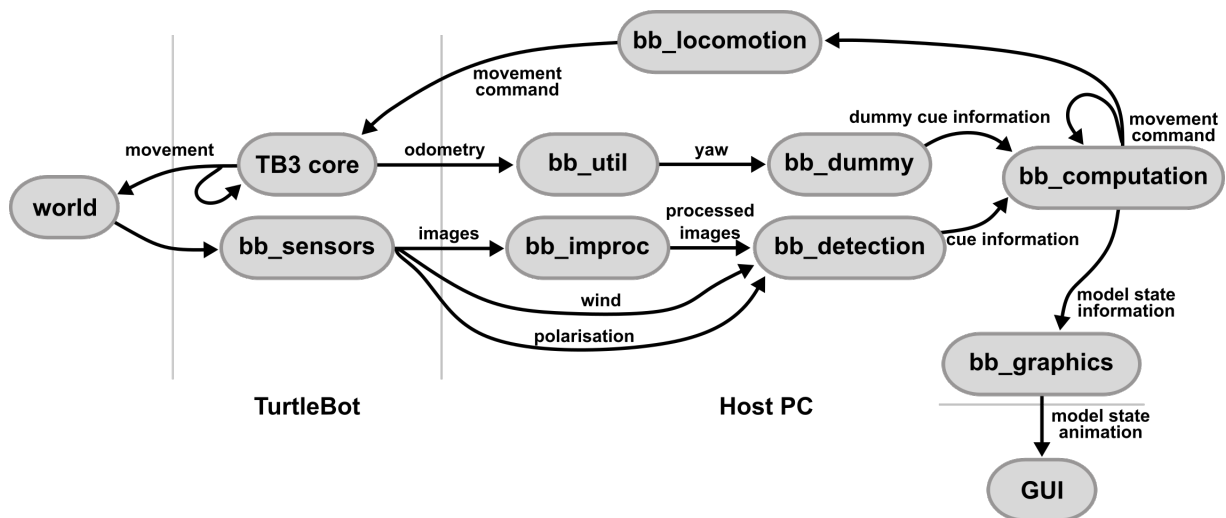


Figure 6.4: Schematic demonstrating how the different BeetleBot software packages fit together. **TB3 core** denotes utilities provided by ROBOTIS for the TurtleBot kit. Sensors are read using **bb_sensors**, raw camera data goes to **bb_improc** for processing (e.g. blurring and channel selection) before being passed on to **bb_detection**. **bb_detection** contains packages for detecting orientation cues in different modalities. **bb_computation** contains insect navigational models and host processes. **bb_locomotion** provides a link between navigational models and motor output. **bb_util** provides shared utilities and message types, as well as yaw information (based on the IMU). **bb_dummy** provides dummy cue utilities based on yaw. **bb_graphics** provides live animated output of neural model state for the [Stone et al. \(2017\)](#) path integration model, [Le Moël et al. \(2019\)](#) vector memory model, and the modified cue integration and steering model from Chapter 7.

had originally wanted to build an antennal sensor. Unfortunately, this was not straightforward and as wind detection was not the focus of the project, I decided to use a more standard design. From a cursory look at the literature, most bio-inspired wind detection systems focus on hair-like sensors which require complex, expensive fabrication methods. An antennal wind sensor which mimicked the Johnston's organ ([Dacke et al., 2019](#)), with a bio-inspired bilateral arrangement ([Suver et al., 2019](#)) could present an interesting contribution.

Miniaturised sensor units While the current polarisation sensor and wind sensor could be made to work together (allowing intensity, wind, and polarisation to be tested in combination on the robot), it would be useful to produce smaller, more robust sensor units. The core components of the polarisation sensor (UV-sensitive photodiodes and linear polarising filters) are reportedly straightforward to miniaturise. Reducing the size of the wind sensor is more difficult. Mechanical wind sensing is generally tricky at a small scale. If we eschew biological plausibility then temperature or air pressure gradients could be used to measure wind speed.

6.3 Summary

One of the original goals of this project was to construct a robot on which cue integration models could be tested. While most of the work for this thesis was conducted in simulation, a BeetleBot was constructed. The platform was used successfully to field-test a prototype polarisation sensor ([Gkaniias et al., in press](#)). In addition, the BeetleBot has been used to test a modified version of the neural model from ([Paper II](#)), demonstrating that the model can account for the dung beetle orientation snapshot in the real world.

Chapter 7

Demonstrating the orientation snapshot using the BeetleBot

Previous work has shown that dung beetles translate orientation cues into the same spatial reference frame (el Jundi et al., 2016; Dacke et al., 2019); this is known as the dung beetle orientation snapshot (though see Chapters 2 and 8 for clarification). Dacke et al. (2019) demonstrated that dung beetles (*Kheper lamarcki*) were able to integrate new cues into their orientation snapshot (see Chapter 2). They allowed beetles to roll their balls to the edge of a circular arena under four conditions in sequence. Beetles were first presented with a single green LED, on the two subsequent rolls a wind cue was added (at $\pm 90^\circ$ to the light cue), and on the final roll the light cue was removed. The change in bearing between the first (LED only) and final (wind only) condition was recorded for $n = 40$ beetles. (Note that they also varied the presentation order, sometimes starting with wind and finishing with the LED.) These changes in bearing were statistically compared to a predicted angle of 0° using a V test (Batschelet, 1981). If the change in bearing was clustered around 0° this was taken to mean that the beetles had integrated the new wind information into their existing orientation snapshot.

Dacke et al. (2019) found that the beetles could indeed maintain their bearing. This multimodal interaction should improve the robustness of the compass system as the animal should be able to maintain its bearing even where cues are intermittent. This is described both as integrating a new cue into the snapshot and a transfer of information between a sun and a wind compass. In light of the work in (Paper II), this can equivalently be seen as a calibration of different cue modalities which compensates for physical offsets between cues, allowing an agent to flexibly integrate new cues into a pre-existing compass.

This experiment was simulated in (Paper II) and it was found that the neural circuit could account for this ability. Here, the neural cir-

cuit is placed on a biorobot (the BeetleBot) and put under comparable experimental conditions to the real beetle to demonstrate that the explanation provided by (Paper II) holds in the real world. The model has had some minor modifications and has also been outfitted with a bio-inspired steering circuit to enable menotaxis. In addition to the biorobotic data, this chapter also provides an exploration of our current understanding of insect goal direction and steering circuitry.

7.1 Methods

7.1.1 Robot behaviour

The robot hardware is described in Chapter 6. Experiments were conducted at Lund University in Sweden during September 2023.

Experimental arena The experimental arena used for the robot was identical (in terms of hardware) to that used for dung beetles in (Paper I; Paper II) and should be comparable to that used by Dacke et al. (2019) (with minor adjustments for the robot). See Figure 7.1 for a picture of the arena.

A circular wooden arena with a textured surface ($r = 0.5m$) and edges marked in 5° increments was placed beneath two crossed metal arches ($r = 1.5m$). These arches were lined with LED strips (Dotstar; Adafruit industries, New York USA). Clusters of three LEDs were used to generate an ersatz sun cue. Intensity and spectral properties were not measured. Light intensity was chosen so that the robot could reliably detect the light cue if it was in view. A transparent plastic board was placed over the textured arena surface so that the robot to operate on the arena, while still allowing the angular increments to be read. Reflections from this plastic board were not visible to the robot. The setup was configured to provide an ersatz sun cue at approximately 45° elevation in two possible locations (180° apart in azimuth).

A wind cue was provided by a custom-built wind generator placed approximately $0.9m$ from the arena centre. The generator consisted of three fans (PFR0912XHEE, 4.50 A; Delta Electronics Inc., Taipei City, Taiwan) distributed evenly over $1m$. Wind speed and (radial) generator position were chosen so that the robot could detect the air current when on the opposite side of the arena. The wind generator was placed such that it could be $\pm 90^\circ$ azimuthally offset from a given sun cue.

The robot was operated from a laptop within the tent environment. Light from the laptop was not visible to the robot. Different *individuals*

are simulated by re-initialising the neural model (see below). Neural model state was recorded for each individual.




Figure 7.1: The BeetleBot in the experimental setup. The circular wooden arena was the same as that used for (Paper I) and (Paper II). The white dots mark the centres of the individual fans which make up the wind generator. A transparent plastic sheet was placed over the textured arena surface to allow the robot to operate whilst keeping the markings visible. A laptop was used to operate the robot when in the tent. No lights from the laptop were visible to the robot. No reflections from the plastic sheet were visible.

Stimulus presentation Each individual exited the arena four times. On the first exit, the robot had a single initial cue available (wind or sun). On the subsequent two exits, the robot had two cues available (wind and sun). On the final exit, the initial cue was removed.

The choice of initial cue was varied and the azimuthal offset between the cues was varied such that it was either 90° or -90° . This was done to ensure that the choice of initial cue or offset direction did not have an effect. Wind-only conditions also had an ersatz sun (cluster of three LEDs) present in the zenith, to mimic previous beetle experiments (Dacke et al., 2019) (Paper I; Paper II). The zenithal light is required for motivation in real beetles and is assumed not to have an effect on their behaviour (see cue detection section below).

Robot handling and behaviour The robot was placed in the centre of the arena and the neural model was initialised (see below). On command, the robot would rotate approximately 360° on the spot (a dance) and select a random goal direction (which was then encoded in the modelled FC2 neurons, see below). The robot would then attempt to steer along this goal direction until it reached the edge of the arena. The robot stopped when it reached a distance of $0.5m$ from its start point. The exit angle was noted and the robot was placed in the centre of the arena.

Each individual exited the arena four times (with the stimuli presented as described above); that is, the neural model persisted over all four exits. For each exit, the robot's initial orientation was randomly chosen to ensure the robot was actually trying to steer along a set direction. After the fourth exit, the neural model was re-initialised and this was treated as a new individual.

Dance rotations were controlled by the IMU to make sure they were consistent. The duration of each menotactic attempt was controlled by idiothetic path integration which also used the IMU.

Robot control During menotaxis, steering was controlled by a rudimentary proportional derivative (PD) controller which operated on differential steering neuron output (Equation 7.7). Large differences in steering neuron output were interpreted as large orientation errors. (This is sufficiently true to be useful, but fails in specific cases where the robot is orientated in the direction opposite to the goal - see PFL3 implementation in Section ??.) Where orientation errors were very large the robot would halt and correct on the spot. Forward movement would resume once the orientation error was reduced below a threshold. This behaviour was chosen to reduce the effect of the control scheme on the results.

Cue detection Cues were either present or absent; wind speed or perceived light elevation did not alter the weight given to a cue. For all data in Figure 7.3, cue weight was normalised such that single cues had a weight of 1; if both cues were present, each had a weight of 0.5 (see [Paper II](#)). For the data in Figure 7.5, this normalisation was disabled to see if it was affecting behaviour.

Wind presence was detected via (any rotation of) the anemometer and direction was then determined from the vane (see Chapter 6). The ersatz sun was detected by looking for the brightest spot in the image and drawing a vector from the image centre to this spot (the 'brightest vector'). The length of this vector indicates perceived elevation and was thresholded so that only lights below a perceived elevation of approximately 70° were detected.

This meant that, for wind-only stimuli, the light in the zenith had no effect (though occasionally it was detected at the arena’s edge). The presence of a dominant light source ensured that no visual noise was detected as an erroneous brightest point. It also streamlined the experimental process and acted as a working light in the blackout tent.

Statistical analysis

Circular statistical analysis was conducted using custom-made python software (Paper II). Each individual’s change in bearing between the first and fourth exit was collected, and the set of all changes was compared to a predicted mean of 0° using the V test (Dacke et al., 2019; Batschelet, 1981). The population mean vector is computed then projected onto a unit vector pointing in a pre-chosen predicted direction. The test statistic V ranges from -1 to 1 , where values approaching 1 indicate that the population aligned with the prediction and approaching -1 indicate the population was orientated in the opposite direction.

7.1.2 Multimodal compass model

The compass model from (Paper II) was used with some modifications. Where terms have been altered, I have kept the symbols the same and these can be treated as drop-in replacements for corresponding terms in Appendix D. Exact values for any tuning constants can be found in the codebase (BeetleBot Software) (see bb_computation documentation).

Learning rate The learning rate is now dynamic and coupled to angular velocity input:

$$\eta = \frac{c}{v_{\max}} \cdot [v_\theta]_{[0, v_{\max}]} \quad (7.1)$$

where c is a tuning constant, v_{\max} is the maximum angular velocity which can be perceived by the network (see Appendix D, Section 1.2 - P-EN neurons), and v_θ is the absolute angular velocity in the instant. $[x]_{[y,z]}$ denotes a clipping operation $[x]_{[y,z]} = \max(y, \min(x, z))$. The learning rule is identical to (Paper II).

Flexible R neuron regulation As learning rate increases, R neuron inputs to EPGs are inhibited so as to reduce (but not eliminate) the influence of allothetic compass inputs (Paper II). This inhibition is also now proportional to instantaneous angular velocity and is given by:

$$S_R = -1.4 \cdot (1 - [\rho \cdot v_\theta]_{[0, 0.8]}) \quad (7.2)$$

where v_θ is the current angular velocity and ρ is a tuning constant. S_R is a scaling term, so the smaller its absolute value, the less input the R neurons get; -1.4 is the basic scalar R neuron weight onto the EPGs which is then reduced (proportional to angular velocity) by the bracketed term.

Tuning In bringing the model into the real world, some issues with network stability were noted. The simulated network did not need to deal with physical handling and inputs were relatively consistent from one timestep to the next. Moreover, cases where cue inputs changed without accompanying self-motion input were rare. In reality, it seemed quite easy to move the physical robot in such a way that resulted in the EPG bump dying off due to rapidly changing cue inputs. If the EPG bump dies off this could be unrecoverable depending on the state of the $R \rightarrow$ EPG mappings.

The network has been re-tuned to increase the relative input from the protocerebral bridge to the EPGs (this amounts to a small up-weighting of the PENs and PEGs). This seems to have greatly improved bump robustness. The exact changes can be found in the beetlebot software repository ([BeetleBot Software](#)).

Activation function noise A small amount of Gaussian noise has been added to the output of the sigmoid activation function which appears to improve robustness in the network. Noise seems to interact favourably with the ring attractor circuitry in a way which encourages the bump back to a stable state. This helps for initialisation but also in cases where abrupt changes in inputs deform the bump. The activation function is now:

$$r = \frac{1}{1 + e^{-(aI+b)}} + \epsilon \quad (7.3)$$

where I is the input, a and b give the slope and bias of the activation function, and $\epsilon \sim N(0, 0.001)$.

Initialisation In ([Paper II](#)), the model was initialised by rapidly rotating it on the spot. In this chapter, the inclusion of Gaussian noise and the upregulating of protocerebral bridge inputs to the EPGs seems to allow bumps to form naturally in the EPGs given enough time. Small imbalances in population activity due to noise get reinforced by lateral inhibition and local excitation which results in a bump of activity forming. This did not happen in the version of the network described in ([Paper II](#)),

in part because no activation noise was present. As such, network initialisation consists of letting the model sit for some time, during which a bump is formed.

Due to this spontaneous bump formation, $R \rightarrow$ EPG mappings can be randomised (and normalised) on model initialisation. I would consider this a slight improvement in terms of plausibility, in that we now do not have set starting assumptions as to activity; noisy neurons and circuitry support bump formation.

7.1.3 Neural steering circuit

A bio-inspired neural steering circuit has been added to the model which is based on insights from [Mussells Pires et al. \(2024\)](#) (see Figure 7.2). This circuit is not anatomically constrained, rather, I took the steering principle described by [Mussells Pires et al. \(2024\)](#) and generalised it to allow the construction of a steering circuit which could be easily integrated into the pre-existing compass model. This revealed a set of general rules which could be used to construct a variety of different steering circuits.

The rules themselves are presented in the discussion below (Section 7.3.2). Here, I only give the implementation details for the specific steering circuit used. The steering circuit requires two additional neural populations which have been named according to their functional biological counterparts. Notation for inputs and outputs is identical to ([Paper II](#)) (Appendix D).

FC2 (goal direction) neurons These neurons encode a goal direction as a sinusoid which inherits its directional meaning from the EPGs. Each FC2 neuron has a specific preferred angle, generating a sinusoidal signal which acts as a population vector encoding.

The input for the i th FC2 neuron is given by:

$$I_{\text{FC2}_i} = \cos(\theta_g - \theta_i) \quad (7.4)$$

where θ_g is the goal direction and θ_i is the preferred angle of the i th FC2 neuron. This formulation is essentially identical to that given by [Mussells Pires et al. \(2024\)](#) (they pass this through a non-linear transformation which is fit to their behavioural data). The model steering circuit uses eight FC2 neurons with evenly spaced preferred directions, however in theory, any number of FC2 neurons should work (with a minimum of three, see discussion below). Note that the ‘preferred angle’ of the FC2 neuron and the goal represented by the FC2 neuron are *not* the same (see Figure 7.2).

PFL3 (steering) neurons PFL3 neurons attempt to place the EPG bump at a specific point on the EPG ring by comparing outputs from differently tuned EPG neurons. This comparison is then gated by FC2 neurons, whose activity determines which PFL3 comparisons drive steering, and therefore what the target head direction is. PFL3s split into two populations, left and right. In this work, these directions describe the sampling direction. That is, *left* PFL3s sample from EPGs which lie to the left of the goal and *right* PFL3s sample from EPGs which lie to the right of the goal. Left PFL3s drive right turns and right PFL3s drive left turns (Figure 7.2). (This left/right labelling ends up being counter-intuitive regardless of whether it represents the sampling or steering direction.) The input to PFL3 neurons is given by:

$$I_{\text{PFL3}_L} = W_{\text{FC2} \rightarrow \text{PFL3}_L} * \bar{r}_{\text{FC2}} + W_{\text{EPG} \rightarrow \text{PFL3}_L} * \bar{r}_{\text{EPG}} \quad (7.5)$$

$$I_{\text{PFL3}_R} = W_{\text{FC2} \rightarrow \text{PFL3}_R} * \bar{r}_{\text{FC2}} + W_{\text{EPG} \rightarrow \text{PFL3}_R} * \bar{r}_{\text{EPG}} \quad (7.6)$$

$W_{\text{EPG} \rightarrow \text{PFL3}_L}$ and $W_{\text{EPG} \rightarrow \text{PFL3}_R}$ are offset from each other by a single column. By scaling the weight matrices, we can set the relative inputs from the EPGs and FC2s (which essentially serves the role of the d parameter given by [Mussells Pires et al. \(2024\)](#)).

A final steering signal is generated by summing the total activity on the left and right and taking the difference ([Mussells Pires et al., 2024](#)) (also see [Stone et al. \(2017\)](#)). The activity imbalance generates a left or a right turn:

$$\text{steering} = \left(\sum_i r_{\text{PFL3}_{L_i}} \right) - \left(\sum_i r_{\text{PFL3}_{R_i}} \right) \quad (7.7)$$

Left/right PFL3 activity will be balanced when the robot is pointing towards its goal direction, and maximally unbalanced when the robot is pointing 90° from its goal direction. Curiously, the PFL3 activity will also be balanced when the robot is pointing 180° from its goal direction. This problem has been described as ‘false nulling’ and it is thought that PFL2 neurons (not modelled) correct for this ([Westeinde et al., 2024](#)). The false null state is inherently unstable as any deviation will steer the robot away from the anti-goal direction and back to the goal and, in practice, the robot almost always corrected for this behaviour.

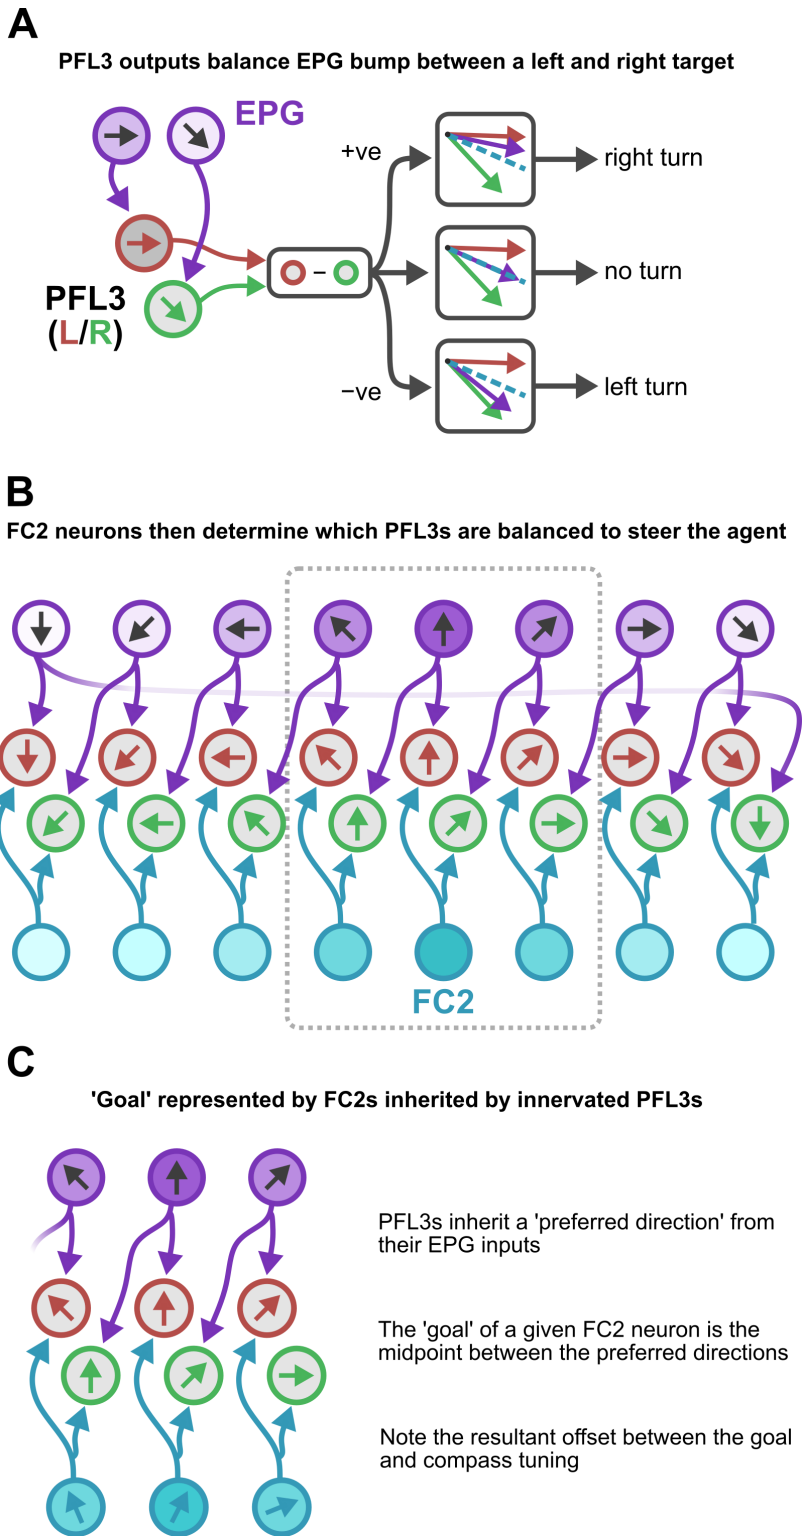


Figure 7.2: A bio-inspired (but not anatomically constrained) steering circuit. **A** Pairs of PFL3 neurons (L and R) receive input from EPGs with distinct directional tuning. L/R PFL3 activity will be balanced when the EPG bump lies at the midpoint between the directions sampled by the L and R PFL3 neurons. Thus, the imbalance in PFL3 generates a steering command towards that midpoint. **B** A complete circuit. Different pairs of PFL3 neurons will try to place the EPG bump at different points in the EPG ring. FC2 neurons allow the circuit to select which PFL3s will have an influence on steering. **C** FC2 neurons can be thought of as representing a 'goal direction'; the goal represented by a given FC2 neuron is inherited from the PFL3 neurons it innervates.

7.2 Results

7.2.1 Modality transfer

The model captures beetle behaviour

The results for the modality transfer experiment are shown in Figure 7.3(C-E). The comparable results from [Dacke et al. \(2019\)](#) are shown in Figure 7.3B. An initial dataset (Figure 7.3C) was collected to see if the robot could maintain its bearing when it was given a comparable behavioural task. The robot passes the same statistical test as the beetles (V-test; Robot: $n = 40, V = 0.6, p < 1 \times 10^{-6}$; Beetles: $n = 40, V = 0.29, p < 0.01$ ([Dacke et al., 2019](#))). Statistically, the robot is actually doing a better job of keeping its bearing than the beetles do. This is true for all datasets collected (Figures 7.3C-E, and Figure 7.5).

Dance direction biases behaviour

While this is a positive result, there is noticeable biasing in Figure 7.3C. On average, the robot is steering to the left of its goal and the relatively tight clustering of this data suggests that this bias is systematic. Each individual in Figure 7.3C only rotated left when dancing. If EPGs did not update in lock-step with R neuron inputs, then the compass would be right-shifted during dances and the robot would over-correct to the left.

To test for this, I repeated the experiment with individuals which only danced to the right (Figure 7.3D). The behaviour is now right-biased which aligns with the explanation given above. I repeated the experiment again with alternating dances to see if this cancelled out the drift effect (Figure 7.3E). This appears to work for exits two and three, but for the fourth exit there is still a substantial drift. This indicates that the biasing is partly caused by circuit lag, however, this does not explain why the deviation is consistently larger on the final exit.

Changes in relative cue weight exacerbate drift

To further investigate, I looked at the angles encoded by each layer of the head direction model over the course of the dance (averaged over the population). For the left-dancing individuals in Figure 7.3C, these are plotted in Figure 7.4 (the average rate of change for each neural population is given in Table 7.1). It can be seen that the EPGs (and all ‘downstream’ populations) drift on every dance, but the effect is clearly more pronounced on the first and fourth dances. This can also be seen if we

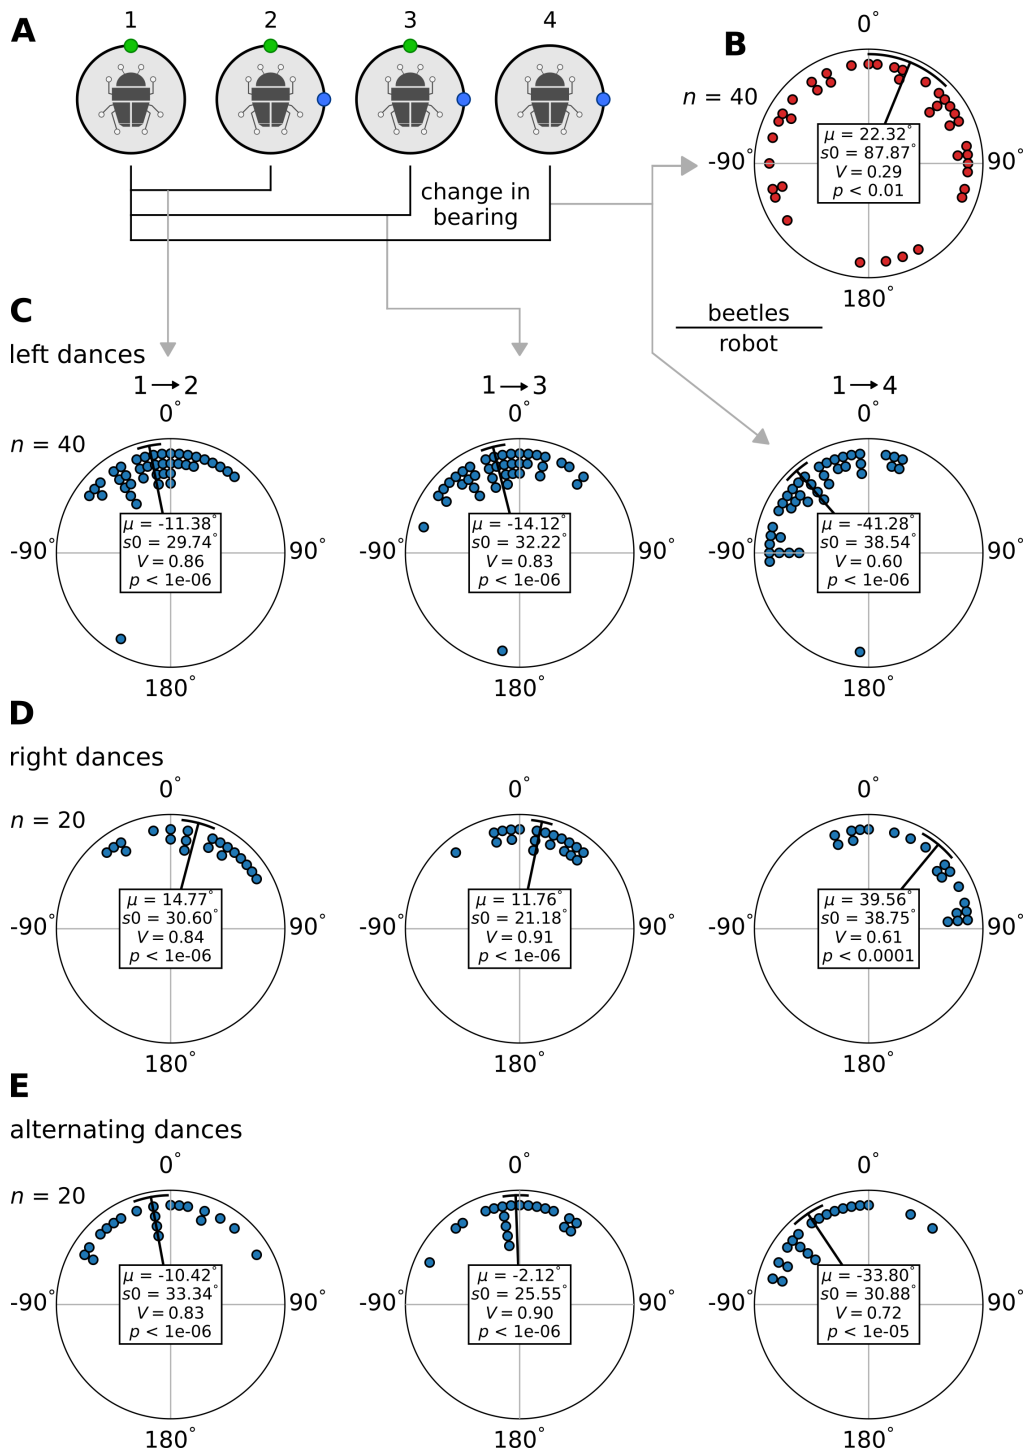


Figure 7.3: Modality transfer experiment results. **A** Behavioural procedure schematic. Each individual performs one exit with an initial cue, two with both cues, and a final exit with the initial cue removed. Cue offsets were varied (either $\pm 90^\circ$), and the choice of initial cue was varied (half of the individuals started with a simulated sun, half started with wind). **B** The beetle data from the original behavioural experiment (Dacke et al., 2019). **C** Changes in bearing for robot individuals which only danced to the left. **D** Changes in bearing for individuals which only danced to the right. **E** Changes in bearing for individuals which performed alternating dances (R, L, R, L). For the robot, changes in bearing were computed between the first and each subsequent exit. In the original experiment, they were only given for the first and fourth exit. Each plot in a given column can be compared. Plots across a row illustrate how the average orientation changed over the course of the experiment.

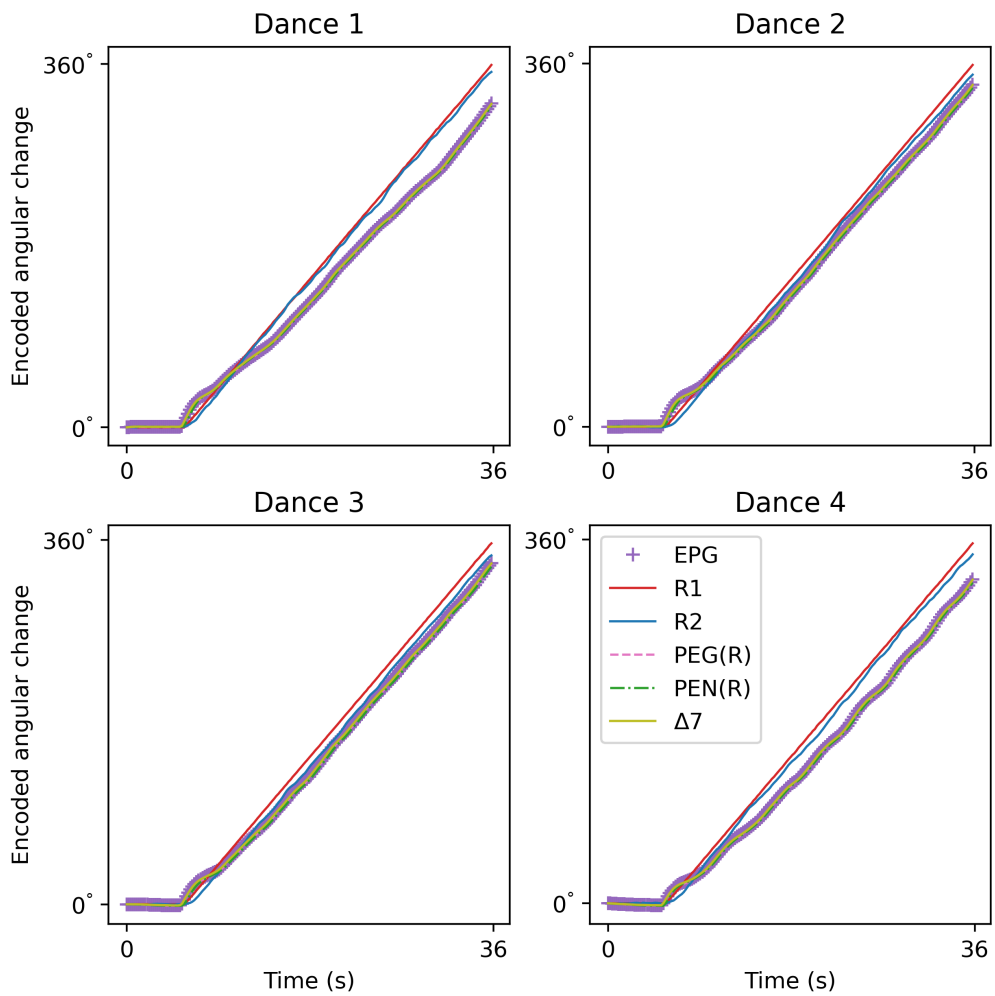


Figure 7.4: Average change (across all individuals in Figure 7.3C) in angle represented by R1 (sun cue, red), R2 (wind cue, blue), and EPG (purple) neurons over the course of each dance. The position of the ersatz sun as measured from the camera (R1 trace) can be considered a ground truth for the robot's true rotation in the world. Note that for dances 1 and 4, the traces for R1 and R2 have only 20 samples because half of the time, the robot started with wind and half of the time it started with the ersatz sun.

look at the average angular velocity experienced by each neuron population over the course of the dance (Table 7.1). On the first dance, $R \rightarrow$ EPG connections are forming from scratch so R input to EPGs is changing rapidly over the course of the rotation. On the second and third dances, the robot should have a relatively stable mapping from the initial cue onto the EPGs, and a weak mapping for the secondary cue. On the final dance, the removal of the first cue results in the loss of a relatively strong input to the EPGs (see Appendix D, Section 1.6). In addition, on the final dance the strengthening of the mapping from the secondary R population onto the EPGs will also cause a shift in input strength.

One way to test for the effect of varying input strengths is to remove the inter R group normalisation. This means that: (1) single cues get less weight which should result in less drastic shifts in input if they are added or removed, and (2) the reduction in weight also means a reduction in learning rate as this is proportional to R and EPG activity. Both of these things together mean that R neuron influence on EPGs should change less severely when cues are added or removed, and shifts in input strength should be more gradual as learning occurs. To this end, I removed inter R group normalisation and repeated the experiment with $n = 20$ individuals. The behavioural results are shown in Figure 7.5, the encoded angles during dancing are shown in Figure 7.6, and the average angular velocities experienced by each population are given in Table 7.2. There is a noticeable reduction in behavioural drift which is also reflected in the rotation experienced by each neuron class (Figure 7.6). The average angular velocities experienced over each dance are also more consistent (even if there is still some drift).

In summary, it seems that there is some natural delay in the EPGs which is exacerbated by large changes in the relative influence of EPG inputs. The source of the natural delay is not clear but it could be due to the fact that EPG inputs are asynchronous. R neurons represent instantaneous allothetic input, PENs represent instantaneous idiothetic input, and then PEGs and PENs also represent past network state. (Looking at Tables 7.1 and 7.2, EPGs usually experience a marginally higher average angular velocity than PEGs and PENs; there is a very small systematic difference between these populations.) It is not clear why large changes in R inputs cause more extreme drifts (beyond the basic fact that large changes in inhibition will cause EPG bump distortion which would change the encoded direction).




Figure 7.5: Changes in bearing for individuals with intra R-group normalisation disabled.

Table 7.1: Average rates of change ($^{\circ}/s$) for each neuron population for each dance, over all individuals in Figure 7.3C (with inter R group normalisation). PEGs and PENs are split into left and right groups as these populations should encode the same bump twice over. R1 rates of change represent a ground truth for the robots true rate of rotation.

	Dance 1	Dance 2	Dance 3	Dance 4
EPG	8.93	9.43	9.48	9.07
R1	10.02	9.97	10	9.99
R2	9.88	9.82	9.72	9.82
PEG(L)	8.9	9.4	9.44	9.06
PEG(R)	8.9	9.4	9.44	9.06
PEN(L)	8.88	9.34	9.4	9.01
PEN(R)	8.87	9.34	9.4	9.01
Delta7	8.93	9.42	9.47	9.06

Table 7.2: Average rates of change without normalisation. Average rates of change ($^{\circ}/s$) for each neuron population for each dance, over all individuals in Figure 7.5. PEGs and PENs are split into left and right groups as these populations should encode the same bump twice over. R1 rates of change represent a ground truth for the robots true rate of rotation.

	Dance 1	Dance 2	Dance 3	Dance 4
EPG	9.66	9.56	9.36	9.45
R1	10.02	10.02	9.97	10.11
R2	10.01	9.96	9.84	9.81
PEG(L)	9.63	9.56	9.33	9.44
PEG(R)	9.63	9.56	9.32	9.44
PEN(L)	9.63	9.49	9.3	9.38
PEN(R)	9.63	9.49	9.29	9.39
Delta7	9.66	9.56	9.35	9.44

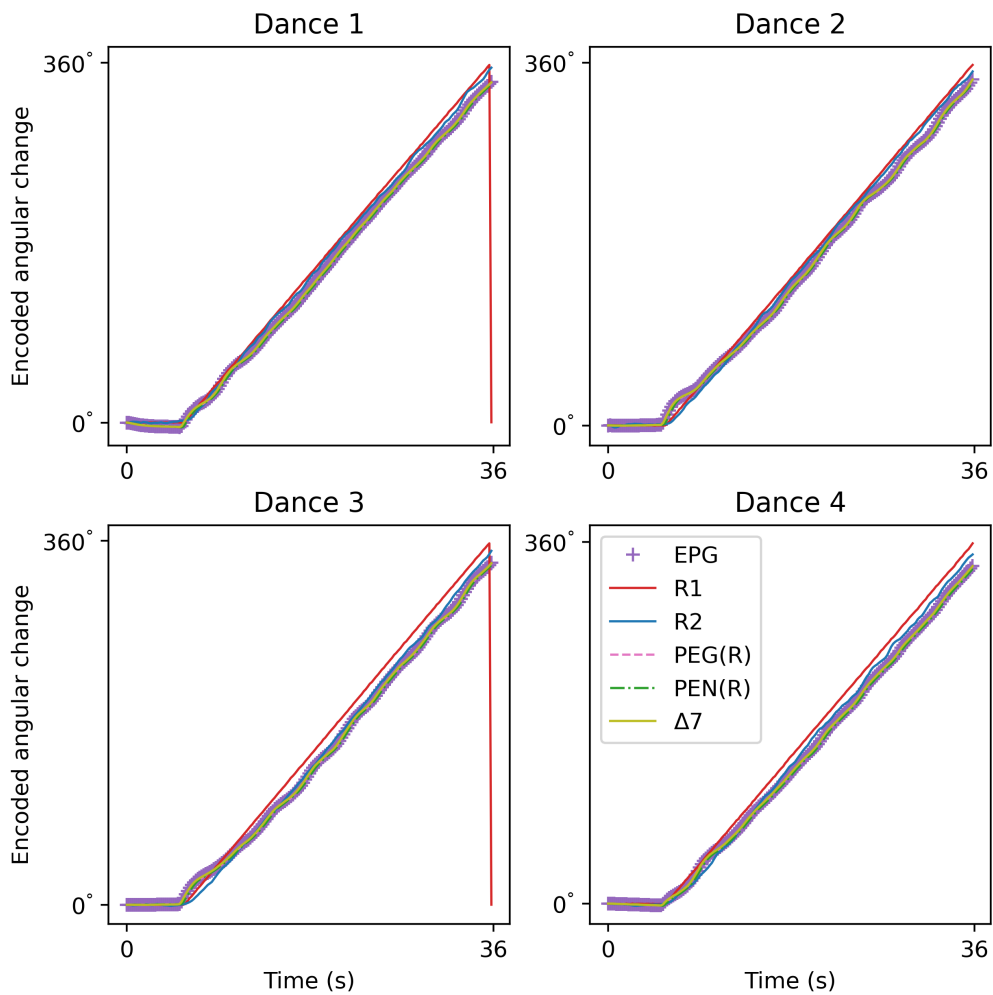


Figure 7.6: Average angular change encoded by each neuron population over the course of each dance with inter R group normalisation removed. These traces are averaged across every individual in Figure 7.5. Only one sub-group is shown for the PENs and PEGs. The large jumps at the end of dances 1 and 3 are due to the encoded angle crossing the 0/360° boundary.

7.3 Discussion

7.3.1 Modality transfer

Here I have demonstrated that the neural model presented in (Paper II) can explain a previous result in dung beetles when tested on a robot, using the same experimental assay in a comparable setup.

Interestingly, there do not seem to be many behavioural examples of this type of cue synchronisation experiment in the literature. Dacke et al. (2019) also demonstrate a similar capability for green and UV point sources, positioned 180° apart (an identical experiment is also given by el Jundi et al. (2015a)). el Jundi et al. (2016) demonstrate that dung beetles are able to learn their bearing with respect to all cues present (the authors used green and UV point sources, and UV polarised light). Even if the beetles responded to a ‘primary cue’, they were able to keep their bearing when this cue was removed. The authors suggest that the (at the time) apparently plastic (Seelig and Jayaraman, 2015) interface between R and EPG neurons could allow this information to be stored (el Jundi et al., 2016). This substrate’s capacity to encode cue offsets has been demonstrated both in simulation (Paper II) and here on a robot. The plausibility of this mechanism allowing an animal to learn arbitrary cue configurations is also discussed by Okubo et al. (2020) and reiterated in recent reviews (Fisher, 2022; Wilson, 2023) (in the context of fruit flies). The model given in (Paper II) is first relatively complete implementation which demonstrates this hypothesis.

The model does experience some systematic error and, while this could be a problem with the implementation (e.g. circuit tuning), it is possible that insects experience similar problems. Seelig and Jayaraman (2015) note that (in fruit flies) the EPG bump did not always respond instantly to cue movement; this could be symptomatic of a delay caused by the disagreement between instantaneous and recurrent inputs. Given that the EPG lag in the robot data was associated with dances, it is also worth considering how dung beetles dance in reality. Similar compass lag could be hidden in a population if dance directions were evenly split, this could be a problem for the beetles which we cannot see due to the form of the available data. From a robot engineering perspective, compass lag seems like a terrible bug which needs to be fixed. From a biological perspective, the animal does not need to be perfectly engineered, it just needs to perform well enough to reproduce.

While there are potential behavioural analyses which could allow us to study compass lag, the circuit’s sensitivity to changes in cue weights is more difficult. Again, this is most likely a limitation of the implementation but it would be good to know what mechanisms allow the real

circuit to remain stable in the face of changing inputs. It is worth considering just how remarkable this capacity is. Neuron influence can be considered as a combination of activity and synaptic strength. PEN influence on EPGs should therefore reflect the animal's angular velocity. Now consider the learning process. As $R \rightarrow$ EPG mappings are formed, R neurons gradually start providing targeted inhibition onto the EPGs. Relative PEN influence therefore decreases. The rotation signalled by R neurons and that signalled by PENs starts to disagree (if relative PEN influence drops, then the perceived angular velocity from self-motion also drops). The extent of this disagreement depends on the relative influence of Rs and PENs, but also depends on PEGs. PEGs provide recurrent input which tries to keep the EPG bump where it is which is another directional influence on the compass bump. The interaction of all these different influences must be very tightly tuned to ensure consistent compass updates, even while the relative influence of each population changes.

To reiterate a point from Chapter 3, this interaction seems to be treated as a solved problem. However, other models have either omitted plastic inputs (Kutschireiter et al., 2023; Pisokas et al., 2020; Goulard et al., 2021; Sun et al., 2018), or omitted R neuron input during learning (Cope et al., 2017; Dan et al., 2021). The issues described in this chapter may have been missed by previous work. It would seem that just controlling the relative inputs from different neuron classes constitutes an interesting line of investigation. This is highly relevant to discussions of allothetic cue influence in cue integration, as well as the relative influence of idiothetic and allothetic cues.

7.3.2 A possible generalisation of insect steering

The steering circuit include above is not anatomically correct but does operate using the same principle as the circuit described in fruit flies (Mussells Pires et al., 2024). The modifications were made for practical reasons but adapting the anatomy revealed a set of rules for building fly-like steering circuits. Moreover, the rules permit a wide variety of circuit architectures which all steer the agent in the same way. These rules are presented here as they may provide a framework with which to analyse future circuit descriptions across species.

Preliminaries

Below I will use vectors as an abstraction for neurons. This discussion is limited to two dimensional vectors with real components (i.e. vectors in \mathbb{R}^2). Linear algebra gives the concept of a *basis*, that is, two vectors whose whose linear combinations can describe any point in 2D space.

For example, if

$$\mathbf{x} = a\mathbf{u} + b\mathbf{v} \quad (7.8)$$

then we can choose a and b such that \mathbf{x} can represent any point in 2D space so long as \mathbf{u} and \mathbf{v} are linearly independent (i.e. not parallel).

A *positive basis* is an analogous concept in which the vector coefficients (a and b above) are restricted to positive values. This means that at least three vectors are required to span the plane (see Figure 7.7). Now we have,

$$\mathbf{x} = a\mathbf{u} + b\mathbf{v} + c\mathbf{w} \quad (7.9)$$

with a, b, c all strictly positive. In order to form a positive basis, the origin must lie within the convex hull of $\mathbf{u}, \mathbf{v}, \mathbf{w}$. (The convex hull intuition is due to a discussion on StackExchange¹.)

The positive basis concept is illustrated in Figure 7.7 and follows from reverse engineering the population vector average commonly used to decode activity bumps in the central complex (e.g. (Seelig and Jayaraman, 2015)). A population of neurons is thought to encode an angle as a vector which is decomposed across a population of neurons, each of which is sensitive to a different preferred direction. The neurons can be thought of as a set of vectors pointing along the preferred directions. The basic requirement for such a set to represent any angle is that it forms a positive basis. Thinking of population codes in this way can be useful, for example, there is no requirement that tuning curves tile space evenly, so long as a positive basis is formed.

Rules for steering circuits

If we assume the steering circuit operates according to the activity balancing principle described by Mussells Pires et al. (2024) (also see Stone et al. (2017)), then we can generalise to a set of core rules for circuit construction. In the following, I will use vectors to represent directionally tuned neurons. I will refer to cells and the vectors representing those cells interchangeably; I find it easiest to think in terms of a loose union of the two concepts (e.g. Figure 7.7D).

A steering circuit is made up of three neural populations: compass, goal, and steering (which split into L and R). We can think of these populations as unit vector sets: \mathbf{C} for compass neurons, \mathbf{G} for goal neurons, and $\mathbf{S}_L, \mathbf{S}_R$ for the two groups of steering neurons. Lower case with a sub-script refers to a specific neuron, e.g. \mathbf{c}_i is the (vector describing the)

¹<https://math.stackexchange.com/questions/4714021/term-for-positive-only-basis>. Last accessed: 26/10/2023.

i th compass neuron. The basic construction of a steering circuit is then governed by the following rules (also see illustrations in Figure 7.8):

1. **C must form a positive basis.** This states that any direction can be represented by some combination of compass neuron activity. This is intuitively true for any set of ‘compass neurons’.
2. **Steering neurons form left/right pairs according to goal neuron innervation.** Steering pair i is defined by the goal neuron \mathbf{g}_i which inputs to those steering neurons. I.e. $\mathbf{s}_{L,i}$ and $\mathbf{s}_{R,i}$ both receive input from \mathbf{g}_i .
3. **Steering neuron direction is a weighted sum of the compass neurons sampled.** A steering neuron $\mathbf{s}_{L,i}$ can sample from multiple compass neurons in a weighted fashion. The direction of that steering neuron is then the weighted sum of all compass neurons sampled: $\mathbf{s}_{L,i} = w_1\mathbf{c}_1 + w_2\mathbf{c}_2 + \dots + w_n\mathbf{c}_n$ (assume $\mathbf{s}_{L,i}$ is then normalised, which also implies that $\|\mathbf{s}_{L,i}\| \neq 0$). The compass neurons do not need to be contiguous and different steering neurons could sample from different numbers of compass neurons (and may have overlapping samples).

In the simplest case, a steering neuron samples from a single compass neuron (as in Figure 7.2).

4. **Left/right sampling of compass neurons must be consistent.** From the previous rule $\mathbf{s}_{L,i}$ and $\mathbf{s}_{R,i}$ both have a direction which is a weighted sum of their inputs. If $\mathbf{s}_{L,i}$ lies to the left of $\mathbf{s}_{R,i}$, then all L steering neuron directions must lie to the left of all R steering neuron directions. If this were not the case then some steering pairs would generate left turns while others generated right turns. By the same logic, paired steering neurons cannot have the same direction: $\mathbf{s}_{L,i} \cdot \mathbf{s}_{R,i} \neq 1$ (recall, these vectors are normalised). This either violates sampling consistency or means that no steering neuron pairs will generate turns.
5. **The ‘goal’ represented by a goal neuron is defined by innervated steering neurons.** The goal represented by a particular goal neuron is the midpoint between the directions of the steering neurons which are innervated by that goal neuron: $\mathbf{g}_i = \mathbf{s}_{L,i} + \mathbf{s}_{R,i}$.

From a modelling perspective, this is distinct from the preferred firing direction. The goal represented by a specific goal neuron is entirely defined by its wiring.

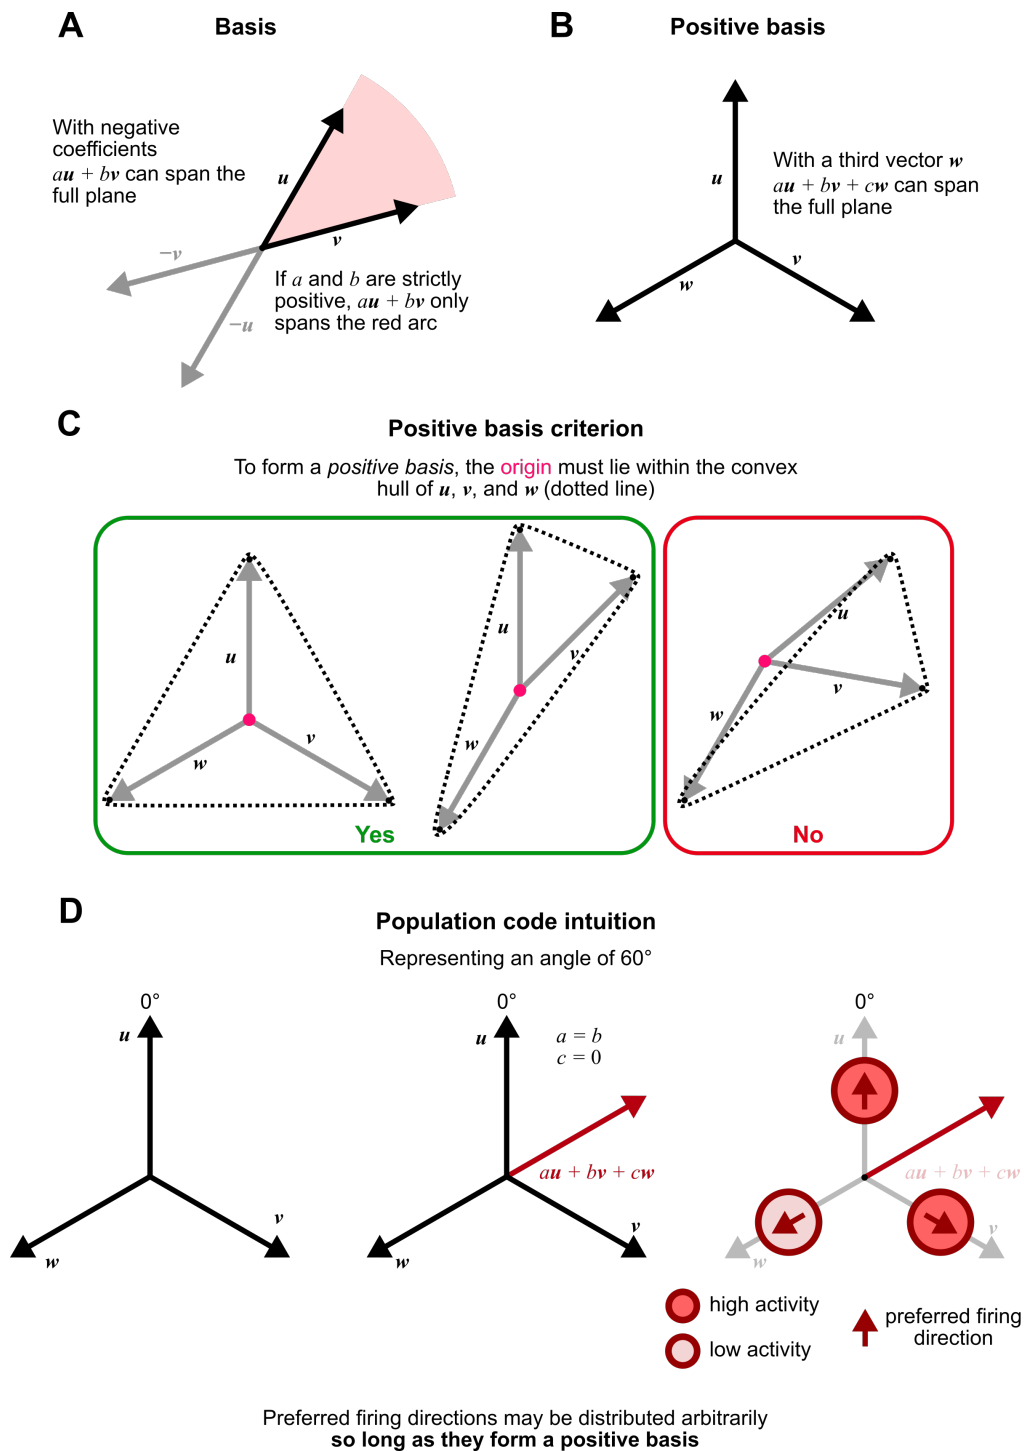


Figure 7.7: Illustrating the concept of a positive basis. **A** A standard 2D basis. **B** An example positive basis. **C** Illustrating the convex hull criterion for a set of vectors to be considered a positive basis. **D** Linking the positive basis concept to population codes. Neurons can be thought of as representing vectors. Firing rates represent magnitudes (or positive coefficients), and angles are represented by preferred firing directions. Different angles can then be represented as linear combinations of these vectors.

6. **G must form a positive basis.** This states that any goal direction can be represented by some combination of goal neuron activity. This is again intuitively required to be able to represent any goal with respect to the compass.

These rules seem to allow a lot of flexibility in circuit construction. For example, there are next to no limitations on how steering neurons sample compass neurons. The resultant angle between steering neuron directions could even be different for different steering pairs, so long as rules 4 and 6 are obeyed. Rule 6 also implies that ‘represented goals’ do not have to be equally spaced, and similarly, rule 1 implies that compass neuron tuning does not have to be equally spaced. There is also no requirement on the number of neurons used in each layer. A minimal circuit could be constructed with three compass neurons, three goal neurons, and six steering neurons. The steering circuit described in Figure 7.2 was constructed by using these rules so as to be easily integrated with the multimodal compass model from (Paper II).


Interestingly, this implies that the steering circuit anatomy described by [Mussells Pires et al. \(2024\)](#) is more complicated than it needs to be to do what it does in the way that it does. This immediately poses the question ‘Why?’. This question is out of scope for this thesis, but this could be very interesting to explore in future. Different forms of steering circuit could be optimised for different tasks or ecologies whilst all operating in the same basic way.

In addition, thinking in terms of these rules might also help in understanding observed anatomy. For example, [Mussells Pires et al. \(2024\)](#) describe a double sampling of specific protocerebral bridge columns and note this as ‘curious feature’ of the anatomy which introduces a small amount of deviation between the goal neuron phase and the compass phase in their model. While [Mussells Pires et al. \(2024\)](#) never claim that the deviation is important, the fact that they mention it suggests that they expected the phase of the head direction and goal signals to be aligned. The fact that it approximately does appears to be an artefact of their modelling, in that they use the same preferred firing angles for goal and compass neurons. Rule 5 above suggests that there is no reason to expect the phase of the goal direction neurons to match the phase of the compass neurons (and raises the question of how we should define ‘preferred direction’ for goal neurons).


As another example, [Mussells Pires et al. \(2024\)](#) use twelve functional PFL3 columns (left/right neuron pairs) based on their anatomical findings, but note that an alternative nine column model produced qualitatively similar results. In this chapter I used eight functional columns and a different PFL3 left/right sampling scheme and obtained a functional steering circuit which operates in the same way as the anatomi-

1. Compass neuron preferred directions must form a positive basis


e.g. for a set of three compass neurons




2. Steering neurons form left/right pairs defined by goal neuron innervation



3. Steering neuron direction is a weighted sum of the compass neurons sampled



4. Left/right sampling of compass neurons must be consistent



5. The 'goal' represented by a goal neuron is defined by innervated steering neurons



6. Goal neuron 'goal directions' must form a positive basis

e.g. for a set of three goal neurons




Figure 7.8: Illustrating rules for steering circuits. **Rule 1** Compass neuron preferred directions must form a positive basis (Figure 7.7) in order to represent any angle in $[0, 360^\circ]$. **Rule 2** Steering neuron pairs are defined by goal neuron innervation. Coloured boxes indicate paired steering neurons. **Rule 3** Steering neuron direction is a weighted sum of the compass neurons sampled. Note that the central steering pair only samples from two compass neurons while the outer pairs sample from three. Equal weighted inputs are used for clarity but this is not required. For larger circuits, this could be more convoluted so long as the remaining rules are obeyed. **Rule 4** The sign of the angle between left and right driving steering neurons must be consistent. In the invalid example, the central steering pair will generate turns in the opposite direction to the others. **Rule 5** The goal direction represented by a specific goal neuron is the mid-point of the directions of the steering neurons it innervates. Note that this is in turn defined by the combination of compass inputs to a given steering neuron as in rule 3. **Rule 6** Resultant goal neuron goal directions must form a positive basis.

cal circuit. It could be interesting then, to perform a comparative study and see if different species' steering circuits were constructed in different ways which all obeyed the same rules. Connectonomic variation could be behaviourally relevant, while conserving the underlying function.

7.3.3 Repeating experiments *in silico*

The experiment presented here acts as a good example as to why robot testing is useful in a literature dominated by computer simulation (where concrete computational models are included). Implementing this work on a physical robot indicated a few shortcomings of the simulation used in (Paper II).

As an example, wind was a far more reliable cue than was assumed in that work, and moreover, it is clear that the error in perceived wind position was not evenly distributed about its true position (i.e. von Mises with zero mean as assumed in (Paper II)). On average, the wind input lags slightly compared to the intensity cue (Figures 7.4 and 7.5). This implies that the $R \rightarrow EPG$ mappings formed by wind would perhaps be slightly skewed, but they would not experience significant disruption. In turn, this makes it extremely unlikely that disruption to the $R \rightarrow EPG$ mapping would actually affect cue influence. This conclusion is obviously influenced by sensor design but it is worth considering this obviousness. In simulation, it is easy and uncontroversial to abstract noisy wind detection to a von Mises process. This effectively hides assumed details of wind sensor construction in the simulation case and does not consider how realistic this is for, for example, a bilateral antennal array.

The simulation also included quite a simplistic notion of time which is made explicit by forcing the model to operate in real time on a robot. Movements which were very fast in simulation ended up being quite slow in reality which highlighted how sensitive the model was to changes in instantaneous PEN influence. Moreover, the assumption that inputs could simply be scaled to operate on arbitrary timescales was incorrect. Low PEN activation over 100 seconds is not the same as high PEN activation over 10 seconds. Low PEN activation can fail to overcome recurrent input from PEGs and therefore fail to update the compass at all.

In simulation, even a stochastic sequence of events is quite tightly controlled. As soon as the model was running on a robot which was being handled in an arena, it became clear just how unintentionally sterile the simulation was. With the benefit of hindsight it is easy to argue that noise and time could have been better represented, and some kind of handling model could have been included. The broader take-home is that there are many things we *could* account for in simulation which a robot gives us for free (Webb, 2000). I am not arguing that simula-

tions are useless; despite its shortcomings, the simulation in (Paper II) provided insight into orientation cue integration in insects. Robotic implementations also have problems but these are usually obvious from the design of the system (e.g. it is clear that the BeetleBot wind sensor will not behave like antennae). The robot, whatever its shortcomings, is forced to operate in the real world.

7.4 Summary

The neural model from (Paper II) has been shown to account for the modality transfer capability in dung beetles in simulation. Here, this capacity was demonstrated on a robot operating under similar behavioural conditions to the beetles in that experiment. While the robot statistically out-performed the beetles, there was some clear biasing in the data which appears to be due to model construction. That said, it is not actually known if animals experience similar biasing; behavioural variation in a population could hide such an effect. Additionally, the head direction circuit was augmented with a bio-inspired (but not anatomical) steering circuit. The construction of this steering circuit led to an interesting generalisation of insect steering which could support a wide variety of circuit architectures which all operate in the same basic way.

7.5 Acknowledgements

I would like to thank Shahrzad Shaverdian and Yakir Gagnon for helping me with the behavioural setup, Elin Dirlik for helping me with photo and video collection, and Marie Dacke for facilitating the trip to Lund University. In addition, my thanks go to all those in the Dacke Lab who spent days 'on-call' as experimental assistants.

Chapter 8

Conclusions

Using dung beetles as a behavioural model, this thesis has explored orientation cue integration in the insect brain. Here I re-state some of the core contributions of the thesis and suggest future work which could shed more light on this problem.

8.1 Cue integration as a vector sum

The work in (Paper I) suggested beetle cue integration behaviour could best be captured by vector summation. In (Paper II), we provided the first relatively complete implementation of a multimodal cue integration circuit as previously described (Okubo et al., 2020; Fisher, 2022; Wilson, 2023) and showed that it could account for the vector sum (as suggested in (Paper I) and by Kutschireiter et al. (2023)). This proposed computation suggests that compass cue integration may vary across species despite conserved anatomy. This argument is supported by the fact that, in our model, cue influence depends on the form of the input rather than circuit structure (Paper II).

The obvious question posed is whether or not the real circuit dynamics match the vector sum prediction. If so, we should be able to observe an apparent shift in compass activity from cue integration to cue capture as cue conflict increases. This hypothesis is directly testable in fruit flies and may also be testable in dung beetles. In flies, it should be possible to image EPG neurons while presenting orientation cues at varying degrees of conflict; we should be able to observe the shift from cue integration to cue capture in the circuit. A similar test is perhaps more difficult in dung beetles but still possible using electrophysiology. As a motivating example, extracellular recordings from rat head direction neurons appear to show a shift from cue integration to cue capture as predicted by the insect model (Knight et al., 2014).

8.2 Reliability does not determine cue influence

Recent work has argued that Hebbian synapses between R and EPG neurons could result compass cues being ‘weighted’ (see (Paper II)) according to their reliability (Wilson, 2023). Reliability has also suggested to account for cue influence in dung beetles (Dacke et al., 2019; Khaldy et al., 2022).

The modelling from (Paper II) strongly suggests that this is not an emergent property of the circuit. The modelling also indicates that cue weight, the amplitude of the sinusoidal R neuron population code, is the primary factor which determines cue influence. It is worth noting that we do not actually know if this amplitude can change, let alone what could cause it to change; indeed, this could vary by species.

Consider fruit flies, in which R neurons within the same group appear to inhibit one another. If this inhibition is fast, then input variance could regulate the amplitude (Hulse et al., 2021), the idea being that high variance leads to scattered inhibition which reduces overall activity. This could plausibly provide a mechanism for the fly circuit to weight cues by reliability. This is also testable in that we could provide cues with angular variance and examine R population amplitude over time. Such evidence would provide an interesting contrast with current dung beetle data. We could also change other aspects of the cue to see if this affects amplitude. Our contrast hypothesis could be tested by examining firing rate as light elevation increases. Additionally, we could present opposing light sources and decrease the contrast between them to demonstrate that increasing elevation and decreasing contrast have the same physiological effect.

There are also other factors which affect cue influence. In the model, it was assumed that R neuron populations were normalised against each other and total R input to a given EPG input was normalised. Changing these assumptions changes how the R neurons interact with the compass (Chapter 7) and it is not entirely clear whether this would present a problem for the real animal. Cue reliability can also have an effect by corrupting R \rightarrow EPG mappings (but again, this does not change cue influence consistently). In all, cue weight appears to be the primary factor which determines cue influence but there are others which can change how R neurons interact with the EPGs.

8.3 Testing models *in silico* provides benefits beyond simulation

The complex interactions which affect cue influence were further elucidated by testing the neural model on a robot. The model was used to recreate a previous behavioural experiment from [Dacke et al. \(2019\)](#), demonstrating that the model can account for the dung beetle orientation snapshot when operating under comparable experimental conditions to the animal. The robot implementation exposed stability issues with the neural model which were hidden by an unintentionally sterile simulation. The shortcomings of the simulation itself were also made clear, particularly the simulated representations of sensor noise and time.

While the problems were minor, the fact that our standard assumptions about noise were inaccurate motivates the testing of models in the real world. Similarly, the simplistic representation of time was assumed to be justified as it seemed we could simply re-scale the inputs to act on different timescales. This ended up not being true, particularly with PEN inputs. In turn, this highlighted that changes in relative influence of different neural populations during learning may cause instability in the compass representation. This is a ‘bug’ which may have been hidden from previous models of the head direction circuit due to lack of plasticity ([Sun et al., 2018](#); [Pisokas et al., 2020](#); [Goulard et al., 2021](#); [Kutschireiter et al., 2023](#)) or artificial gating of allothetic compass input during learning ([Cope et al., 2017](#); [Dan et al., 2021](#)).

The model circuit appears to be relatively brittle in the face of changing inputs and this observation arose from a relatively complete implementation which was forced to operate in the real world. This brittleness immediately poses the question of how the real circuit regulates different inputs to the compass, which is of obvious relevance to discussions on the potential factors involved in cue influence. If different EPG inputs are externally regulated by some gain control mechanism, it is then unclear how the system keeps the relative influence of different EPG inputs consistent.

8.4 The orientation snapshot requires a firm definition

Another benefit of concrete modelling can be seen in the conceptual model of the dung beetle orientation snapshot. A hypothetical snapshot is first described by [Baird et al. \(2012\)](#). The authors propose that beetles

use their first dance to ‘store a compass reading of the celestial cues as they appear along the preferred roll direction.’ They argue that subsequent dances are attempts to re-align with a remembered view of the sky. The authors note that this is strikingly similar to ant view-matching behaviour (Wehner and Wehner, 1990; Lent et al., 2010).

Later, el Jundi et al. (2016) demonstrated that beetles were able to learn arbitrary arrangements of celestial cues and that the beetles learned about non-dominant cues. They proposed that the snapshot could be stored in plastic connections between R and EPG neurons during the dance. Conversely, el Jundi et al. (2018) identified PFN neurons in *Xheper lamarcki* and argued that these neurons could act as a substrate for the goal direction of the animal, based on a (then) recent model of path integration memory in the bee brain (Stone et al., 2017). Different neural substrates are proposed for the same conceptual model and it is easy to see why. The original description from Baird et al. (2012) can be interpreted as encoding relative cue positions and goal direction. The broad nature of the conceptual snapshot model makes it difficult to assign a computational circuit and to determine what behaviour is and is not covered by the ‘snapshot’.

By building a concrete multimodal model with steering circuit, we are forced to consider relative cue position and goal direction as different entities. In addition, it can be seen that only relative cue position depends on rotation for storage, thus it is easier to tie this to the dance. This aligns neatly with the behavioural evidence presented by el Jundi et al. (2016), who showed that beetles only learn the relative position of cues which were present during the dance. Conversely, there is no reason to couple the choosing of a goal direction to the dance behaviour. This ‘choice’ could happen before, during, or after the dance; the point is that there is no clear reason that rotation is required to generate a goal.

While this clearly demonstrates the benefits of a fully-specified model, it should be noted that the original description given by Baird et al. (2012) makes complete sense given the literature available at the time. It is only with the benefit of later studies (Seelig and Jayaraman, 2015; Cope et al., 2017; Stone et al., 2017) that we can pry the different components apart. Nevertheless it would be useful to give the *orientation snapshot* a concrete definition in the literature going forward. The most straightforward definition would be that the snapshot encodes the relative positions of different orientation cues. This aligns with some previous discussion of the snapshot and also provides a firm link to the dance behaviour with which the snapshot is associated.

8.5 General rules for insect steering circuits

The steering circuit generalisation given in Chapter 7 presents an interesting way to contextualise future results. Previous work in the head direction circuit has shown that differences in circuitry do not change overall function (Pisokas et al., 2020).

If we assume steering operates in the same way across species, it would be interesting to generate different steering circuits according to these rules and evaluate them on different navigation tasks. If we assume that new neurons and connections have a cost associated with them, we could perform optimisation to see what kind of circuitry we end up with under reasonable constraints. It could be the case that different circuits perform better or worse on different tasks which could provide predictions as to what circuitry we might see in different insects. Approaching this from a different angle, we could look at non-fly homologues for the circuit described by Mussells Pires et al. (2024). While circuitry may be different, we could ask if the homologous circuits still obey the same rules and thus could operate in the same way. It would again be interesting to examine differences in circuit construction to see what affect they have on behaviour.

8.6 Bayesian cue integration models are of limited utility in insects

Insect cue integration behaviour has been described as Bayesian on a number of occasions (Collett, 2012; Legge et al., 2014; Wystrach et al., 2015; Hoinville and Wehner, 2018; Khaldy et al., 2022). A key conceptual issue addressed in this thesis (Chapter 3) is whether or not this is a meaningful label or useful framework within which to analyse behaviour. Use of Bayes' theorem to provide a benchmark for behaviour is ubiquitous in cue integration literature (Trommershauser et al., 2011) (and human decision making in general (Rahnev and Denison, 2018)). However, it is important that we think about why this is the case. The vast majority of the literature is concerned with human test subjects whose brains are inaccessible. Thus, statistical modelling presents a valuable tool for understanding behaviour.

Statistical modelling is obviously useful in insect behaviour too, and can serve as a valuable starting point for an investigation (i.e. (Paper I)). That said, it is important that we fully understand the framework we are applying, established terminology, and embedded assumptions. Even within the Bayesian decision making literature, standard behavioural adjectives such as 'Bayes-optimal' or 'optimal' are controversial (Bow-

ers and Davis, 2012a,b; Rahnev and Denison, 2018). It should also be noted that the literature in this area is usually concerned with what the nervous system *should* do under specific circumstances (Kording, 2007). The framework assumes test subjects are behaving optimally according to some goal and works to try and find out what that goal is (Griffiths et al., 2012; Rahnev and Denison, 2018). We must take care to specify whether we are interested in what the nervous system should do (and under what assumptions) or what it is actually doing.

This is relevant to previous literature in insects, but also to current hypotheses about cue influence in the insect compass (Khaldy et al., 2022; Wilson, 2023). The idea that insects would have evolved to behave as is mathematically ideal is enticing. But, we must specify what is meant by ideal and why we think this is ethologically relevant for a given insect. Insect brains are relatively accessible and we can produce well-reasoned models to determine whether or not the relevant circuitry supports ideal behaviour.

Epilogue

When studying animal behaviour, humans always have a tendency to rationalise it from our point of view. The same is arguably true of mechanistic models for behaviour and the brain. We form things into neat conceptual models and try to explain behaviour and related neural structures in ways which make sense *to us*. Thinking about neural computation in this way is useful, but the resultant conceptual models can be overly accommodating and are susceptible to engineering bias. When we think as engineers, we tend to think about how systems ought to behave according to our criteria, rather than how they actually behave. We then miss the assumptions embedded in our intuition and overlook system interactions which may interfere with our neat model. It is therefore important that we provide functional implementations of conceptual models.

Functional implementations highlight the assumptions hidden by our engineering bias, this is particularly powerful where hidden assumptions are practically testable. Where assumptions are plausible they become predictions and where they fail, they prompt a re-examination of our conceptual model. In turn, testing implementations on robots forces us to consider the pathway from sensory input to model input, and from model output to motor output. Testing implementations on robots demonstrates that our model can function outside a simulation, or points out complications not captured by the simulation.

In all, testing fully-specified functional models in the real world forces us to confront engineering bias in our conceptual model. This is a strong test of our conceptual model, with the added benefit of exposing the problems the biological system has solved which we have yet to consider.

Bibliography

Frank Cole Babbitt. *Moralia, Volume V: Isis and Osiris. The E at Delphi. The Oracles at Delphi No Longer Given in Verse. The Obsolescence of Oracles.* Loeb Classical Library 306, Harvard University Press, Cambridge, MA, 1936. Translation by Frank Cole Babbitt.

Emily Baird, Marcus J. Byrne, Clarke H. Scholtz, Eric J. Warrant, and Marie Dacke. Bearing selection in ball-rolling dung beetles: Is it constant? *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, 196(11):801–806, 2010. ISSN 03407594. doi: 10.1007/s00359-010-0559-8.

Emily Baird, Marcus J Byrne, Jochen Smolka, Eric J Warrant, and Marie Dacke. The dung beetle dance: an orientation behaviour? *PLoS One*, 7(1):e30211, 2012.

Edward Batschelet. *Circular statistics in biology.* ACADEMIC PRESS, 111 FIFTH AVE., NEW YORK, NY 10003, 1981, 388, 1981.

M Jerome Beetz, Christian Kraus, and Basil El Jundi. Neural representation of goal direction in the monarch butterfly brain. *Nature Communications*, 14(1):5859, 2023.

George Boas. *The Hieroglyphics of Horapollo.* Princeton University Press, 1993. ISBN 978-0691000923. Translation by George Boas.

Jeffrey S Bowers and Colin J Davis. Bayesian just-so stories in psychology and neuroscience. *Psychological bulletin*, 138(3):389, 2012a.

Jeffrey S Bowers and Colin J Davis. Is that what bayesians believe? reply to griffiths, chater, norris, and pouget (2012). 2012b.

Michael L Brines and James L Gould. Bees have rules. *Science*, 206(4418): 571–573, 1979.

Paul Bromiley. Products and convolutions of gaussian probability density functions. *Tina-Vision Memo*, 2003. URL

<https://amitbny.github.io/akb.github.io/studymat/Sem4/convolution.pdf>. Last accessed on 24/07/2023.

- Marcus Byrne, Marie Dacke, Peter Nordström, Clarke Scholtz, and Eric Warrant. Visual cues used by ball-rolling dung beetles for orientation. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, 189(6):411–418, June 2003. ISSN 03407594. doi: 10.1007/s00359-003-0415-1.
- Ken Cheng, Sara J Shettleworth, Janellen Huttenlocher, and John J Rieser. Bayesian integration of spatial information. *Psychological bulletin*, 133(4):625, 2007.
- Allen Cheung, Shaowu Zhang, Christian Stricker, and Mandyam V Srinivasan. Animal navigation: the difficulty of moving in a straight line. *Biological cybernetics*, 97:47–61, 2007.
- William G Cochran. Problems arising in the analysis of a series of similar experiments. *Supplement to the Journal of the Royal Statistical Society*, 4(1):102–118, 1937.
- Matthew Collett. How navigational guidance systems are combined in a desert ant. *Current Biology*, 22(10):927–932, 2012.
- Matteo Colombo and Peggy Seriès. Bayes in the brain—on bayesian modelling in neuroscience. *The British journal for the philosophy of science*, 2012.
- Alex J Cope, Chelsea Sabo, Eleni Vasilaki, Andrew B Barron, and James AR Marshall. A computational model of the integration of landmarks and motion in the insect central complex. *PloS one*, 12(2): e0172325, 2017.
- M. Dacke, M. J. Byrne, E. Baird, C. H. Scholtz, and E. J. Warrant. How dim is dim? Precision of the celestial compass in moonlight and sunlight. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1565):697–702, March 2011. ISSN 14712970. doi: 10.1098/rstb.2010.0191.
- M. Dacke, Basil el Jundi, Jochen Smolka, Marcus Byrne, and Emily Baird. The role of the sun in the celestial compass of dung beetles. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1636), February 2014. ISSN 14712970. doi: 10.1098/rstb.2013.0036.
- Marie Dacke, P Nordström, C Scholtz, and Eric Warrant. A specialized dorsal rim area for polarized light detection in the compound eye of

- the scarab beetle *pachysoma striatum*. *Journal of Comparative Physiology A*, 188:211–216, 2002.
- Marie Dacke, Dan Eric Nilsson, Clarke H. Scholtz, Marcus Byrne, and Eric J. Warrant. Animal behaviour: insect orientation to polarized moonlight. *Nature*, 424(6944):33, 2003a. ISSN 14764687.
- Marie Dacke, Peter Nordström, and Clarke H. Scholtz. Twilight orientation to polarised light in the crepuscular dung beetle *Scarabaeus zambesianus*. *Journal of Experimental Biology*, 206(9):1535–1543, May 2003b. ISSN 00220949. doi: 10.1242/jeb.00289.
- Marie Dacke, Marcus J. Byrne, Clarke H. Scholtz, and Eric J. Warrant. Lunar orientation in a beetle. *Proceedings of the Royal Society B: Biological Sciences*, 271(1537):361–365, 2004. ISSN 14712970. doi: 10.1098/rspb.2003.2594.
- Marie Dacke, Marcus Byrne, Jochen Smolka, Eric Warrant, and Emily Baird. Dung beetles ignore landmarks for straight-line orientation. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, 199(1):17–23, 2012. ISSN 03407594. doi: 10.1007/s00359-012-0764-8.
- Marie Dacke, Emily Baird, Marcus Byrne, Clarke H. Scholtz, and Eric J. Warrant. Dung beetles use the milky way for orientation. *Current Biology*, 23(4):298–300, February 2013. ISSN 09609822. doi: 10.1016/j.cub.2012.12.034.
- Marie Dacke, Adrian T.A. Bell, James J. Foster, Emily J. Baird, Martin F. Strube-Bloss, Marcus J. Byrne, and Basil El Jundi. Multimodal cue integration in the dung beetle compass. *Proceedings of the National Academy of Sciences of the United States of America*, 116(28):14248–14253, 2019. ISSN 10916490. doi: 10.1073/pnas.1904308116.
- Marie Dacke, Basil El Jundi, Yakir Gagnon, Ayse Yilmaz, Marcus Byrne, and Emily Baird. A dung beetle that path integrates without the use of landmarks. *Animal cognition*, 23:1161–1175, 2020.
- Marie Dacke, Emily Baird, Basil El Jundi, Eric J. Warrant, and Marcus Byrne. How Dung Beetles Steer Straight. *Annual Review of Entomology*, 66(1):243–256, January 2021. ISSN 0066-4170, 1545-4487. doi: 10.1146/annurev-ento-042020-102149.
- Chuntao Dan, Brad K Hulse, Vivek Jayaraman, and Ann M Hermundstad. Flexible control of behavioral variability mediated by an internal representation of head direction. *bioRxiv*, pages 2021–08, 2021.

- Karl Deisseroth. Optogenetics. *Nature methods*, 8(1):26–29, 2011.
- Basil el Jundi, Jochen Smolka, Emily Baird, Marcus J. Byrne, and Marie Dacke. Diurnal dung beetles use the intensity gradient and the polarization pattern of the sky for orientation. *Journal of Experimental Biology*, 217(13):2422–2429, July 2014. ISSN 00220949. doi: 10.1242/jeb.101154.
- Basil el Jundi, James J. Foster, Marcus J. Byrne, Emily Baird, and Marie Dacke. Spectral information as an orientation cue in dung beetles. *Biology Letters*, 11(11):2–5, 2015a. ISSN 1744957X. doi: 10.1098/rsbl.2015.0656.
- Basil el Jundi, Eric J. Warrant, Marcus J. Byrne, Lana Khaldy, Emily Baird, Jochen Smolka, and Marie Dacke. Neural coding underlying the cue preference for celestial orientation. *Proceedings of the National Academy of Sciences of the United States of America*, 112(36):11395–11400, 2015b. ISSN 10916490. doi: 10.1073/pnas.1501272112.
- Basil el Jundi, Eric J. Warrant, Keram Pfeiffer, and Marie Dacke. Neuroarchitecture of the dung beetle central complex. *Journal of Comparative Neurology*, 526(16):2612–2630, 2018. ISSN 10969861. doi: 10.1002/cne.24520.
- Basil el Jundi, James J J. Foster, Lana Khaldy, Marcus J J. Byrne, Marie Dacke, and Emily Baird. A Snapshot-Based Mechanism for Celestial Orientation. *Current Biology*, 26(11):1456–1462, 2016. ISSN 09609822. doi: 10.1016/j.cub.2016.03.030.
- Marc O Ernst and Martin S Banks. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870):429–433, 2002.
- Marc O Ernst and Heinrich H Bülthoff. Merging the senses into a robust percept. *Trends in cognitive sciences*, 8(4):162–169, 2004.
- Yvette E Fisher. Flexible navigational computations in the drosophila central complex. *Current opinion in neurobiology*, 73:102514, 2022.
- Yvette E Fisher, Jenny Lu, Isabel D’Alessandro, and Rachel I Wilson. Sensorimotor experience remaps visual input to a heading-direction network. *Nature*, 576(7785):121–125, 2019.
- Yvette E Fisher, Michael Marquis, Isabel D’Alessandro, and Rachel I Wilson. Dopamine promotes head direction plasticity during orienting movements. *Nature*, pages 1–7, 2022.

- James J. Foster, Basil el Jundi, Jochen Smolka, Lana Khaldy, Dan Eric Nilsson, Marcus J. Byrne, and Marie Dacke. Stellar performance: Mechanisms underlying milky way orientation in dung beetles. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1717), April 2017. ISSN 14712970. doi: 10.1098/rstb.2016.0079.
- James J. Foster, John D. Kirwan, Basil el Jundi, Jochen Smolka, Lana Khaldy, Emily Baird, Marcus J. Byrne, Dan Eric Nilsson, Sönke Johnsen, and Marie Dacke. Orienting to polarized light at night – matching lunar skylight to performance in a nocturnal beetle. *Journal of Experimental Biology*, 222(2), January 2019. ISSN 00220949. doi: 10.1242/jeb.188532.
- James J. Foster, Claudia Tocco, Jochen Smolka, Lana Khaldy, Emily Baird, Marcus J. Byrne, Dan-Eric Nilsson, and Marie Dacke. Light pollution forces a change in dung beetle orientation behavior. *Current Biology*, 31(17):3935–3942.e3, September 2021. ISSN 09609822. doi: 10.1016/j.cub.2021.06.038.
- Ysabel Milton Giraldo, Katherine J Leitch, Ivo G Ros, Timothy L Warren, Peter T Weir, and Michael H Dickinson. Sun navigation requires compass neurons in drosophila. *Current Biology*, 28(17):2845–2852, 2018.
- Evripidis Gkaniias, Benjamin Risse, Michael Mangan, and Barbara Webb. From skylight input to behavioural output: a computational model of the insect polarised light compass. *PLoS computational biology*, 15(7): e1007123, 2019.
- Evripidis Gkaniias, Robert Mitchell, Jan Stankiewicz, Sadeque R. Khan, Srinjoy Mitra, and Barbara Webb. Celestial compass design mimics the fan-like polarisation filter array of insect eyes. *Nature Communications Engineering*, in press.
- Roman Goulard, Cornelia Buehlmann, Jeremy E Niven, Paul Graham, and Barbara Webb. A unified mechanism for innate and learned visual landmark guidance in the insect central complex. *PLoS computational biology*, 17(9):e1009383, 2021.
- Thomas L Griffiths, Nick Chater, Dennis Norris, and Alexandre Pouget. How the bayesians got their beliefs (and what those beliefs actually are): comment on bowers and davis (2012). 2012.
- Ilkka Hanski and Yves Cambefort, editors. *Dung beetle ecology*. Princeton University Press, Princeton, N.J, 1991. ISBN 978-0-691-08739-9.

- Ben J Hardcastle, Jaison J Omoto, Pratyush Kandimalla, Bao-Chau M Nguyen, Mehmet F Keleş, Natalie K Boyd, Volker Hartenstein, and Mark A Frye. A visual pathway for skylight polarization processing in *Drosophila*. *Elife*, 10:e63225, 2021.
- Stanley Heinze, Basil El Jundi, Bente G Berg, Uwe Homberg, Randolph Menzel, Keram Pfeiffer, Ronja Hensgen, Frederick Zittrell, Marie Dacke, Eric Warrant, et al. Insect brain database. URL <https://insectbraindb.org/app/>. Last accessed 22/10/23.
- Stanley Heinze, Basil El Jundi, Bente G Berg, Uwe Homberg, Randolph Menzel, Keram Pfeiffer, Ronja Hensgen, Frederick Zittrell, Marie Dacke, Eric Warrant, et al. A unified platform to manage, share, and archive morphological and functional data in insect neuroscience. *ELife*, 10:e65376, 2021.
- Thierry Hoinville and Rüdiger Wehner. Optimal multiguide integration in insect navigation. *Proceedings of the National Academy of Sciences*, 115(11):2824–2829, 2018.
- Anna Honkanen, Andrea Adden, Josiane da Silva Freitas, and Stanley Heinze. The insect central complex and the neural basis of navigational strategies. *Journal of Experimental Biology*, 222(Suppl_1):jeb188854, 2019.
- Brad K Hulse, Hannah Haberkern, Romain Franconville, Daniel B Turner-Evans, Shin-ya Takemura, Tanya Wolff, Marcella Noorman, Marisa Dreher, Chuntao Dan, Ruchi Parekh, et al. A connectome of the *Drosophila* central complex reveals network motifs suitable for flexible navigation and context-dependent action selection. *eLife*, 10, 2021.
- Esa Ville Immonen, Marie Dacke, Stanley Heinze, and Basil el Jundi. Anatomical organization of the brain of a diurnal and a nocturnal dung beetle. *Journal of Comparative Neurology*, 525(8):1879–1908, June 2017. ISSN 10969861. doi: 10.1002/cne.24169.
- Sönke Johnsen. *The optics of life: a biologist's guide to light in nature*. Princeton University Press, 2012.
- Pete R Jones. A tutorial on cue combination and signal detection theory: Using changes in sensitivity to evaluate how observers integrate sensory information. *Journal of Mathematical Psychology*, 73:117–139, 2016.
- Lana Khaldy, Orit Peleg, Claudia Tocco, L. Mahadevan, Marcus Byrne, and Marie Dacke. The effect of step size on straight-line orientation. *Journal of The Royal Society Interface*, 16(157):20190181,

- August 2019a. ISSN 1742-5689. doi: 10.1098/rsif.2019.0181.
URL <https://royalsocietypublishing.org/doi/10.1098/rsif.2019.0181>.
- Lana Khaldy, Claudia Tocco, Marcus Byrne, Emily Baird, and Marie Dacke. Straight-line orientation in the woodland-living beetle *Sisyphus fasciculatus*. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, 2019b. ISSN 14321351. doi: 10.1007/s00359-019-01331-7.
- Lana Khaldy, Claudia Tocco, Marcus Byrne, and Marie Dacke. Compass cue integration and its relation to the visual ecology of three tribes of ball-rolling dung beetles. *Insects*, 12(6):526, 2021.
- Lana Khaldy, James J. Foster, Ayse Yilmaz, Gregor Belušič, Yakir Gagnon, Claudia Tocco, Marcus J. Byrne, and Marie Dacke. The interplay of directional information provided by unpolarised and polarised light in the heading direction network of the diurnal dung beetle *Kheper lamarcki*. *Journal of Experimental Biology*, 225(3):jeb243734, February 2022. ISSN 0022-0949, 1477-9145. doi: 10.1242/jeb.243734.
- Sung Soo Kim, Ann M Hermundstad, Sandro Romani, LF Abbott, and Vivek Jayaraman. Generation of stable heading representations in diverse visual scenes. *Nature*, 576(7785):126–131, 2019.
- Rebecca Knight, Caitlin E Piette, Hector Page, Daniel Walters, Elizabeth Marozzi, Marko Nardini, Simon Stringer, and Kathryn J Jeffery. Weighted cue integration in the rodent head direction system. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1635): 20120512, 2014.
- David C Knill and Alexandre Pouget. The bayesian brain: the role of uncertainty in neural coding and computation. *TRENDS in Neurosciences*, 27(12):712–719, 2004.
- Konrad Kording. Decision theory: what" should" the nervous system do? *Science*, 318(5850):606–610, 2007.
- Anna Kutschireiter, Melanie A Basnak, Rachel I Wilson, and Jan Dru-gowitsch. Bayesian inference in ring attractor networks. *Proceedings of the National Academy of Sciences*, 120(9):e2210622120, 2023.
- Thomas Labhart. Polarization-opponent interneurons in the insect visual system. *Nature*, 331(6155):435–437, 1988.

- Thomas Labhart and Eric P Meyer. Detectors for polarized skylight in insects: a survey of ommatidial specializations in the dorsal rim area of the compound eye. *Microscopy research and technique*, 47(6):368–379, 1999.
- Dimitrios Lambrinos, Ralf Möller, Thomas Labhart, Rolf Pfeifer, and Rüdiger Wehner. A mobile robot employing insect strategies for navigation. *Robotics and Autonomous systems*, 30(1-2):39–64, 2000.
- Michael F Land and Dan-Eric Nilsson. *Animal eyes*. OUP Oxford, 2012.
- Michael S Landy, Martin S Banks, and David C Knill. Ideal-observer models of cue integration. *Sensory cue integration*, pages 5–29, 2011.
- Richard D Lange, Sabyasachi Shivkumar, Ankani Chattoraj, and Ralf M Haefner. Bayesian encoding and decoding as distinct perspectives on neural coding. *Nature Neuroscience*, pages 1–10, 2023.
- Florent Le Moël, Thomas Stone, Mathieu Lihoreau, Antoine Wystrach, and Barbara Webb. The central complex as a potential substrate for vector based navigation. *Frontiers in psychology*, 10:690, 2019.
- Eric LG Legge, Antoine Wystrach, Marcia L Spetch, and Ken Cheng. Combining sky and earth: desert ants (*Melophorus bagoti*) show weighted integration of celestial and terrestrial cues. *Journal of Experimental Biology*, 217(23):4159–4166, 2014.
- David D Lent, Paul Graham, and Thomas S Collett. Image-matching during ant navigation occurs through saccade-like body turns controlled by learned visual features. *Proceedings of the National Academy of Sciences*, 107(37):16348–16353, 2010.
- Eric G. Matthews. Observations on the Ball-Rolling Behavior of *Canthon pilularius* (L.) (Coleoptera, Scarabaeidae). *Psyche: A Journal of Entomology*, 70(2):75–93, 1963. ISSN 0033-2615, 1687-7438. doi: 10.1155/1963/23630.
- Robert Mitchell. Main beetlebot software repository. URL https://github.com/refmitchell/beetlebot_software. Made public 23/08/2023.
- Robert Mitchell, Shahrzad Shaverdian, Marie Dacke, and Barbara Webb. A model of cue integration as vector summation in the insect brain. *Proceedings of the Royal Society B*, 290(2001):20230767, 2023.
- Martin Müller and Rüdiger Wehner. Wind and sky as compass cues in desert ant navigation. *Naturwissenschaften*, 94:589–594, 2007.

- Richard F Murray and Yaniv Morgenstern. Cue combination on the circle and the sphere. *Journal of vision*, 10(11):15–15, 2010.
- Peter Mussells Pires, Lingwei Zhang, Victoria Parache, LF Abbott, and Gaby Maimon. Converting an allocentric goal into an egocentric steering signal. *Nature*, pages 1–11, 2024.
- Ajay Narendra. Homing strategies of the australian desert ant *melophorus bagoti* ii. interaction of the path integrator with visual cue information. *Journal of Experimental Biology*, 210(10):1804–1812, 2007.
- Tu Anh Thi Nguyen, M Jerome Beetz, Christine Merlin, and Basil El Jundi. Sun compass neurons are tuned to migratory orientation in monarch butterflies. *Proceedings of the Royal Society B*, 288(1945): 20202988, 2021.
- Tatsuo S Okubo, Paola Patella, Isabel D’Alessandro, and Rachel I Wilson. A neural network for wind-guided compass navigation. *Neuron*, 107(5):924–940, 2020.
- Ipek Oruç, Laurence T Maloney, and Michael S Landy. Weighted linear cue combination with possibly correlated error. *Vision research*, 43(23): 2451–2468, 2003.
- Robert P O’Shea, Gary P Misson, and Shelby E Temple. Seeing polarization of light with the naked eye. *Current Biology*, 31(4):R178–R179, 2021.
- Ioannis Pisokas, Stanley Heinze, and Barbara Webb. The head direction circuit of two insect species. *Elife*, 9:e53985, 2020.
- Ioannis Pisokas, Wolfgang Rössler, Barbara Webb, Jochen Zeil, and Ajay Narendra. Anesthesia disrupts distance, but not direction, of path integration memory. *Current Biology*, 32(2):445–452, 2022.
- Morgan Quigley, Brian Gerkey, Ken Conley, Josh Faust, Tully Foote, Jeremy Leibs, Eric Berger, Rob Wheeler, and Andrew Ng. Ros: an open-source robot operating system. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) Workshop on Open Source Robotics*, Kobe, Japan, 2009. URL <https://wiki.ros.org>. Last accessed 13/10/2023.
- Dobromir Rahnev and Rachel N Denison. Suboptimality in perceptual decision making. *Behavioral and brain sciences*, 41:e223, 2018.
- Steven M Reppert and Jacobus C de Roode. Demystifying monarch butterfly migration. *Current Biology*, 28(17):R1009–R1022, 2018.

- ROBOTIS. Turtlebot3 e-manual. URL <https://emanual.robotis.com/docs/en/platform/turtlebot3/overview/>. Last accessed 26/07/2023.
- Pedro Rosas and Felix A Wichmann. Cue combination: Beyond optimality. *Sensory cue integration*, pages 144–52, 2011.
- Emilio Salinas and LF Abbott. Vector reconstruction from firing rates. *Journal of computational neuroscience*, 1(1-2):89–107, 1994.
- Marcel Ethan Sayre, Rachel Templin, Johanna Chavez, Julian Kempenaers, and Stanley Heinze. A projectome of the bumblebee central complex. *Elife*, 10:e68911, 2021.
- Peter Scarfe. Experimentally disambiguating models of sensory cue integration. *Journal of Vision*, 22(1):5–5, 2022.
- Louis K Scheffer, C Shan Xu, Michal Januszewski, Zhiyuan Lu, Shin-ya Takemura, Kenneth J Hayworth, Gary B Huang, Kazunori Shinomiya, Jeremy Maitlin-Shepard, Stuart Berg, et al. A connectome and analysis of the adult drosophila central brain. *Elife*, 9:e57443, 2020.
- Johannes D Seelig and Vivek Jayaraman. Feature detection and orientation tuning in the drosophila central complex. *Nature*, 503(7475):262–266, 2013.
- Johannes D Seelig and Vivek Jayaraman. Neural dynamics for landmark orientation and angular path integration. *Nature*, 521(7551):186–191, 2015.
- Johannes D Seelig, M Eugenia Chiappe, Gus K Lott, Anirban Dutta, Jason E Osborne, Michael B Reiser, and Vivek Jayaraman. Two-photon calcium imaging from head-fixed drosophila during optomotor walking behavior. *Nature methods*, 7(7):535–540, 2010.
- Doron J Shahar. Minimizing the variance of a weighted average. *Open Journal of Statistics*, 7(2):216–224, 2017.
- Shahrzad Shaverdian, Elin Dirlik, Robert Mitchell, Claudia Tocco, Barbara Webb, and Marie Dacke. Weighted cue integration for straight-line orientation. *iScience*, 25(10):105207, October 2022. ISSN 2589-0042. doi: 10.1016/J.ISCI.2022.105207.
- Lori A Sjolund, Jonathan W Kelly, and Timothy P McNamara. Optimal combination of environmental cues and path integration during navigation. *Memory & Cognition*, 46:89–99, 2018.

- William Skaggs, James Knierim, Hemant Kudrimoti, and Bruce McNaughton. A model of the neural basis of the rat's sense of direction. *Advances in neural information processing systems*, 7, 1994.
- Jochen Smolka, Emily Baird, Basil el Jundi, Therese Reber, Marcus J. Byrne, and Marie Dacke. Night sky orientation with diurnal and nocturnal eyes: Dim-light adaptations are critical when the moon is out of sight. *Animal Behaviour*, 111:127–146, January 2016. ISSN 00033472. doi: 10.1016/j.anbehav.2015.10.005.
- Thomas Stone, Barbara Webb, Andrea Adden, Nicolai Ben Weddig, Anna Honkanen, Rachel Templin, William Wcislo, Luca Scimeca, Eric Warrant, and Stanley Heinze. An anatomically constrained model for path integration in the bee brain. *Current Biology*, 27(20):3069–3085, 2017.
- Xuelong Sun, Michael Mangan, and Shigang Yue. An analysis of a ring attractor model for cue integration. In *Biomimetic and Biohybrid Systems: 7th International Conference, Living Machines 2018, Paris, France, July 17–20, 2018, Proceedings 7*, pages 459–470. Springer, 2018.
- Marie P Suver, Andrew MM Matheson, Sinekdha Sarkar, Matthew Damiata, David Schoppik, and Katherine I Nagel. Encoding of wind direction by central neurons in drosophila. *Neuron*, 102(4):828–842, 2019.
- Claudia Tocco, Marie Dacke, and Marcus Byrne. Eye and wing structure closely reflects the visual ecology of dung beetles. *Journal of Comparative Physiology A*, 205:211–221, 2019.
- Claudia Tocco, Marie Dacke, and Marcus Byrne. The finely defined shift work schedule of dung beetles and their eye morphology. *Ecology and Evolution*, 11(22):15947–15960, November 2021. ISSN 2045-7758, 2045-7758. doi: 10.1002/ece3.8264.
- Julia Trommershauser, Konrad Kording, and Michael S Landy. *Sensory cue integration*. Oxford University Press, 2011.
- Daniel Turner-Evans, Stephanie Wegener, Herve Rouault, Romain Franconville, Tanya Wolff, Johannes D Seelig, Shaul Druckmann, and Vivek Jayaraman. Angular velocity integration in a fly heading circuit. *Elife*, 6:e23496, 2017.
- Alberto Ugolini. Relationship between compass systems of orientation in equatorial sandhoppers. *Animal Behaviour*, 62(2):193–199, 2001.

- Karl von Frisch. *The dance language and orientation of bees*. Harvard University Press, 1967. Originally published in German in 1965, first published in English in 1967. Translation by Leigh E. Chadwick.
- Herman von Helmholtz. *Helmholtz's Treatise on Physiological Optics. Volume III: The Perceptions of Vision*. The Optical Society of America, 1925. First published in German in 1867. Translation by James P. C. Southall, translated from the Third German Edition.
- Eric Warrant, Barrie Frost, Ken Green, Henrik Mouritsen, David Dreyer, Andrea Adden, Kristina Brauburger, and Stanley Heinze. The Australian bogong moth *agrotis infusa*: a long-distance nocturnal navigator. *Frontiers in behavioral neuroscience*, 10:77, 2016.
- Barbara Webb. What does robotics offer animal behaviour? *Animal behaviour*, 60(5):545–558, 2000.
- Barbara Webb and Thomas Consi. *Biorobotics*. Mit Press, 2001.
- Rüdiger Wehner and Martin Müller. The significance of direct sunlight and polarized skylight in the ant's celestial system of navigation. *Proceedings of the National Academy of Sciences*, 103(33):12575–12579, 2006.
- Rüdiger Wehner and Sibylle Wehner. Insect navigation: use of maps or ariadne's thread? *Ethology Ecology & Evolution*, 2(1):27–48, 1990.
- Elena A Westeinde, Emily Kellogg, Paul M Dawson, Jenny Lu, Lydia Hamburg, Benjamin Midler, Shaul Druckmann, and Rachel I Wilson. Transforming a head direction signal into a goal-oriented steering command. *Nature*, pages 1–8, 2024.
- Rachel I Wilson. Neural networks for navigation: From connections to computations. *Annual Review of Neuroscience*, 46:403–423, 2023.
- Tanya Wolff and Gerald M Rubin. Neuroarchitecture of the drosophila central complex: A catalog of nodulus and asymmetrical body neurons and a revision of the protocerebral bridge catalog. *Journal of Comparative Neurology*, 526(16):2585–2611, 2018.
- Antoine Wystrach, Michael Mangan, and Barbara Webb. Optimal cue integration in ants. *Proceedings of the Royal Society B: Biological Sciences*, 282(1816):20151484, 2015.
- Ayse Yilmaz, Basil El Jundi, Gregor Belušič, Marcus Byrne, Emily Baird, and Marie Dacke. Mechanisms of spectral orientation in a diurnal dung beetle. *Philosophical Transactions of the Royal Society B*, 377(1862), October 2022a. ISSN 14712970. doi: 10.1098/RSTB.2021.0287.

Ayse Yilmaz, Yakir Gagnon, Marcus Byrne, Emily Baird, and Marie Dacke. Cold-induced anesthesia impairs path integration memory in dung beetles. *Current Biology*, 32(2):438–444, 2022b.

Appendix A

Polarised light detection

Rather than try to come up with yet another way to explain the detection of polarised light, I will quote [Johnsen \(2012\)](#) directly as his is the most straightforward explanation I have come across. Recall that light which is polarised oscillates (primarily) along one axis in the plane perpendicular to the direction of travel (Figure 2.3). Assume we have a linear polariser and a light meter.

“First, we place the light meter in the path of the beam (facing the source) and measure the irradiance (at a given wavelength), which we’ll call I . Nothing mysterious here, we are just measuring how bright the beam is. Then, leaving the light meter where it is, we put the linear polarizer in the beam with its transmission axis oriented horizontally and take a measurement. Then turn the polarizer 90° (so that it now maximally transmits vertically polarized light) and take another measurement. Subtract the third measurement from the second and call this number Q . If the light is unpolarized, it doesn’t matter how the polarizer is oriented and Q equals zero. If the light is completely horizontally polarized, Q will equal about $I - 0 = I$ (it won’t exactly equal I because no real polarizer is perfect). So it is tempting to conclude that Q/I tells us how polarized the light is. . . ” ([Johnsen, 2012](#), Ch. 8)

The photoreceptive microvilliar strands of insect eyes are naturally sensitive to light which is polarised along the length of the microvilli ([Land and Nilsson, 2012](#)). Thus, the perpendicular arrangement of microvilli in DRA rhabdoms (described in Chapter 2) permits downstream neurons to perform the measurement given by [Johnsen \(2012\)](#). This is known as polarisation opponent processing and was first demonstrated in crickets by [Labhart \(1988\)](#).

It should be noted that the polarisation opponent response, Q/I , does not uniquely specify how polarised the light is or the angle at which it is polarised (Johnsen, 2012). However, Q/I does tell us the degree to which the incident light is polarised along the axis of the sensor. If Q/I is -1, the light is maximally polarised along the axis perpendicular to that of our sensor; similarly, if Q/I is 1, the light is maximally polarised along the axis parallel to the sensor. For this sensor to extract the angle and degree of polarisation, it would need to rotate through 360° on the spot. Alternatively, multiple polarisation opponent units can be used to compute the angle and degree without rotating the sensor.

One such method involves having multiple units which all point directly at the zenith, each of which is set to have a different axis. For example, Lambrinos et al. (2000) used three such units with orientations of 0° , 60° , and 120° , and by combining the outputs they could determine the zenithal angle of polarisation, and therefore the orientation of a robot within a period of 180° . Polarised light is often discussed as an inherently ambiguous cue and this is reflected in behavioural results which show animals with axial orientation (e.g. (el Jundi et al., 2015b; Khaldy et al., 2021)).

An alternative model proposed by Gkanias et al. (2019) exploits the 3D geometry of the dorsal rim and its relation to the pattern of polarisation across the sky. In brief, the rhabdoms in the dorsal rim all point at slightly different parts of the sky and can thus expect a slightly different angle of polarisation. By setting the major axis of each rhabdom in approximate alignment with the expected angle of polarisation in the target patch of sky, Q/I becomes a measure of how well the sensor and the incident polarised light are aligned. $Q/I = 1$ if they are perfectly aligned, and $Q/I = -1$ if they are perpendicular. Due to the concentric pattern of polarised light, the polarisation opponent sensors pointing to the anti-solar part of the sky will always be best aligned with the polarised skylight, and an array of polarisation opponent sensors will have a stronger response in the anti-solar region.

By measuring the polarisation opponent response across the sky and, a visual system can extract the position of the sun from the pattern of polarisation and thus its orientation without the aforementioned 180° ambiguity. This principle has been demonstrated in simulation (Gkanias et al., 2019) and *in silico* (Gkanias et al., in press) (Appendix F). In addition, Gkanias et al. (2019) also demonstrated that axial orientation responses observed in behavioural experiments (e.g. Khaldy et al. (2019b)) could be an a result of the linear polarisers used in the experiment, rather than a limitation of the visual system. (While we did no formal data collection, I also tested the robot platform from Gkanias et al. (in press) under one of the polarising filters typically used in such experiments and found the

same thing.)

The fan-like arrangement of the DRAs in *Pachysoma striatum* (Dacke et al., 2002) and *Scarabaeus zambesianus* (Dacke et al., 2003b) would seem to be well-suited to computing polarised light orientation in this fashion.

Appendix B

Derivation of the linear rule for ‘optimal’ cue integration

Here I will give the full derivation for Equation 3.13 starting from Bayes’ theorem. Insight was drawn from (Murray and Morgenstern, 2010) and this was checked against (Bromiley, 2003).

$$P(X|Y_1, Y_2) = \frac{P(Y_1, Y_2|X)P(X)}{P(Y)} \quad (\text{B.1})$$

Ignoring the normalising constant $P(Y)$

$$P(X|Y_1, Y_2) \propto P(Y_1, Y_2|X)P(X) \quad (\text{B.2})$$

Assuming we have a uniform prior distribution over values of X

$$P(X|Y_1, Y_2) \propto P(Y_1, Y_2|X) \quad (\text{B.3})$$

Assuming Y_1 and Y_2 are independent

$$P(X|Y_1, Y_2) \propto P(Y_1|X)P(Y_2|X) \quad (\text{B.4})$$

In the context of cue integration, the likelihoods $P(Y_1|X)$ and $P(Y_2|X)$ are interpreted as ‘cue noise distributions’ which describe how an individual perceives a cue (Y_1 or Y_2) to vary with respect to the true state of the world (X).

We could substitute any noise distribution in, but the linear rule assumes Gaussian noise. The Gaussian probability density function is given by

$$f(y; \mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{\mu-y}{\sigma}\right)^2} \quad (\text{B.5})$$

The shape of the Gaussian is given by the exponential term

$$\left(\frac{\mu - y}{\sigma}\right)^2 = \frac{y^2 - 2\mu y + \mu^2}{\sigma^2} \quad (\text{B.6})$$

This form will be relevant later. Substituting the Gaussian p.d.f. (Equation B.5) into Equation B.4 we get

$$P(X|Y_1, Y_2) \propto f(y_1; x, \sigma_1)f(y_2; x, \sigma_2) \quad (\text{B.7})$$

Here we are assuming that the Gaussian noise distributions are centred on the true (unknown) value of the stimulus x . For convenience in notation, we use the fact that $f(y; x, \sigma) = f(x; y, \sigma)$ to rewrite Equation B.7 as¹

$$P(X|Y_1, Y_2) \propto f(x; y_1, \sigma_1)f(x; y_2, \sigma_2) \quad (\text{B.8})$$

To standardise notation we swap the ys for μs (as these are now the means of our noise distributions)

$$P(X|Y_1, Y_2) \propto f(x; \mu_1, \sigma_1)f(x; \mu_2, \sigma_2) \quad (\text{B.9})$$

Then

$$P(X|Y_1, Y_2) \propto \frac{1}{\sigma_1\sqrt{2\pi}}e^{\frac{1}{2}\left(\frac{\mu_1-x}{\sigma_1}\right)^2} \cdot \frac{1}{\sigma_2\sqrt{2\pi}}e^{\frac{1}{2}\left(\frac{\mu_2-x}{\sigma_2}\right)^2} \quad (\text{B.10})$$

$$\propto \frac{1}{\sigma_1\sigma_2 2\pi} e^{\left[\left(\frac{1}{2}\left(\frac{\mu_1-x}{\sigma_1}\right)^2\right) + \left(\frac{1}{2}\left(\frac{\mu_2-x}{\sigma_2}\right)^2\right)\right]} \quad (\text{B.11})$$

$$\propto \frac{1}{\sigma_1\sigma_2 2\pi} e^{\frac{1}{2}\left[\left(\frac{\mu_1-x}{\sigma_1}\right)^2 + \left(\frac{\mu_2-x}{\sigma_2}\right)^2\right]} \quad (\text{B.12})$$

The exponential term defines the shape of this function, so we will focus on this part. Let

$$a = \left(\frac{\mu_1 - x}{\sigma_1}\right)^2 + \left(\frac{\mu_2 - x}{\sigma_2}\right)^2$$

Then

¹The result is identical either way, this swap is entirely for convenience in notation.

$$a = \left(\frac{\mu_1 - x}{\sigma_1} \right)^2 + \left(\frac{\mu_2 - x}{\sigma_2} \right)^2 \quad (\text{B.13})$$

$$= \frac{\mu_1^2 - 2\mu_1x + x^2}{\sigma_1^2} + \frac{\mu_2^2 - 2\mu_2x + x^2}{\sigma_2^2} \quad (\text{B.14})$$

$$= \frac{\sigma_2^2(\mu_1^2 - 2\mu_1x + x^2)}{\sigma_2^2\sigma_1^2} + \frac{\sigma_1^2(\mu_2^2 - 2\mu_2x + x^2)}{\sigma_2^2\sigma_1^2} \quad (\text{B.15})$$

$$= \frac{\sigma_2^2\mu_1^2 - 2\mu_1\sigma_2^2x + x^2\sigma_2^2 + \sigma_1^2\mu_2^2 - 2\mu_2\sigma_1^2x + x^2\sigma_1^2}{\sigma_1^2\sigma_2^2} \quad (\text{B.16})$$

Rearranging this expression to collect common terms

$$a = \frac{x^2\sigma_2^2 + x^2\sigma_1^2 - 2\mu_1\sigma_2^2x - 2\mu_2\sigma_1^2x + \sigma_2^2\mu_1^2 + \sigma_1^2\mu_2^2}{\sigma_1^2\sigma_2^2} \quad (\text{B.17})$$

$$= \frac{(\sigma_1^2 + \sigma_2^2)x^2 - 2(\mu_1\sigma_2^2 + \mu_2\sigma_1^2)x + \sigma_1^2\mu_2^2 + \sigma_2^2\mu_1^2}{\sigma_1^2\sigma_2^2} \quad (\text{B.18})$$

Which is quadratic in x as in Equation B.6, thus the shape of this joint distribution is (at least proportionally) Gaussian. To determine the mean and variance, we divide throughout by $(\sigma_1^2 + \sigma_2^2)$ such that the x^2 term has a coefficient of 1 (as in Equation B.6).

$$a = \frac{x^2 - 2 \left(\frac{\mu_1\sigma_2^2 + \mu_2\sigma_1^2}{\sigma_1^2 + \sigma_2^2} \right) x + \left(\frac{\sigma_1^2\mu_2^2 + \sigma_2^2\mu_1^2}{\sigma_1^2 + \sigma_2^2} \right)}{\left(\frac{\sigma_1^2\sigma_2^2}{\sigma_1^2 + \sigma_2^2} \right)} \quad (\text{B.19})$$

The mean is given by the coefficient of the x term and the variance is given by the denominator (again, as in Equation B.6).

Thus, the pointwise product of two Gaussian distributions is proportionally² Gaussian with

$$\mu_c = \frac{\mu_1\sigma_2^2 + \mu_2\sigma_1^2}{\sigma_1^2 + \sigma_2^2} \quad (\text{B.20})$$

and

$$\sigma_c^2 = \frac{\sigma_1^2\sigma_2^2}{\sigma_1^2 + \sigma_2^2} \quad (\text{B.21})$$

²The ‘proportional’ qualifier is due to the fact that the coefficient of the x^0 term is not the square of the coefficient of the x term (as in Equation B.6). However, because the x^0 term does not depend on x , we can simply choose an additive constant to make this equality work (Bromiley, 2003).

To express in terms of reliabilities (Equation 3.13),

$$\mu_c = \frac{\mu_1 \sigma_2^2 + \mu_2 \sigma_1^2}{\sigma_1^2 + \sigma_2^2} \quad (\text{B.22})$$

$$= \frac{\mu_1 + \mu_2 \frac{\sigma_1^2}{\sigma_2^2}}{\frac{\sigma_1^2}{\sigma_2^2} + 1} \quad (\text{divide by } \sigma_2^2) \quad (\text{B.23})$$

$$= \frac{\frac{\mu_1}{\sigma_1^2} + \frac{\mu_2}{\sigma_2^2}}{\frac{1}{\sigma_2^2} + \frac{1}{\sigma_1^2}} \quad (\text{divide by } \sigma_1^2) \quad (\text{B.24})$$

$$\mu_c = \frac{\frac{1}{\sigma_1^2} \mu_1 + \frac{1}{\sigma_2^2} \mu_2}{\frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2}} \quad (\text{reformat to match Equation 3.13}) \quad (\text{B.25})$$

Similarly, for the variance

$$\sigma_c^2 = \frac{\sigma_1^2 \sigma_2^2}{\sigma_1^2 + \sigma_2^2} \quad (\text{B.26})$$

$$\frac{1}{\sigma_c^2} = \frac{\sigma_1^2 + \sigma_2^2}{\sigma_1^2 \sigma_2^2} \quad (\text{flip expression}) \quad (\text{B.27})$$

$$= \frac{\frac{\sigma_1^2}{\sigma_2^2} + 1}{\sigma_1^2} \quad (\text{divide by } \sigma_2^2) \quad (\text{B.28})$$

$$= \frac{\frac{1}{\sigma_2^2} + \frac{1}{\sigma_1^2}}{1} \quad (\text{divide by } \sigma_1^2) \quad (\text{B.29})$$

$$\frac{1}{\sigma_c^2} = \frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2} \quad (\text{reformat to match Equation 3.13}) \quad (\text{B.30})$$

Appendix C

Supplementary information for Paper 1

iScience, Volume 25

Supplemental information

**Weighted cue integration
for straight-line orientation**

Shahrzad Shaverdian, Elin Dirlik, Robert Mitchell, Claudia Tocco, Barbara Webb, and Marie Dacke

Supplementary material




Figure S.1: **The behavioural setup, related to Figures 1 and 2** - The setup used for behavioural experiments, including a circular sand painted arena, two arches lined with green LEDs and three wind generators (each denoted with a white asterisk) positioned at a 0°, 60° and 120° angle relative the green LED used as a sun cue. For detailed information see STAR Methods, Experimental setup.

Table S.1: **The number of beetles tested for tactic behaviour, related to Figure 1** - The number of beetles tested for tactic behaviour at different elevations of an ersatz sun and at different speeds of a wind cue when the ersatz sun was positioned in zenith. Taxis was observed at a 75° solar elevation, as well as at wind speeds 0.5, 0.8, 3.0 and 4.0m/s. The 90° solar elevation has been omitted as there can be no tactic behaviour with respect to a cue in zenith.

Elevation	n	Wind speed	n
5°	18	0.5m/s	10
20°	17	0.8m/s	9
45°	16	1.0m/s	7
60°	17	1.5m/s	14
75°	13	1.9m/s	14
80°	10	2.5m/s	14
82°	15	3.0m/s	13
84°	9	4.0m/s	14
86°	6		
88°	4		




Figure S.2: The proportion of excluded beetles at cue conflict experiments, related to Figure 2
 - To determine the significance threshold for our exclusion criterion, a comparison of the proportion of excluded beetles was carried out for Rayleigh tests using $p = 0.05$ and $p = 0.1$. We found that fewer beetles were excluded when using $p = 0.1$. The behavioural results, i.e. changes in heading, were not different between the two significance levels ($p > 0.5$, Mardia-Watson-Wheeler tests, $14 \leq n \leq 30$). This held true for all conditions where the Mardia-Watson-Wheeler test was applicable; the number of beetles tested at an 86° solar elevation and a 1.25m/s wind speed was less than 10, which is the minimum number required for the Mardia-Watson-Wheeler test. As the behavioural results were not different between the two significance levels we decided to move forward with the threshold of $p = 0.1$ to maximise the amount of data available for our modelling.

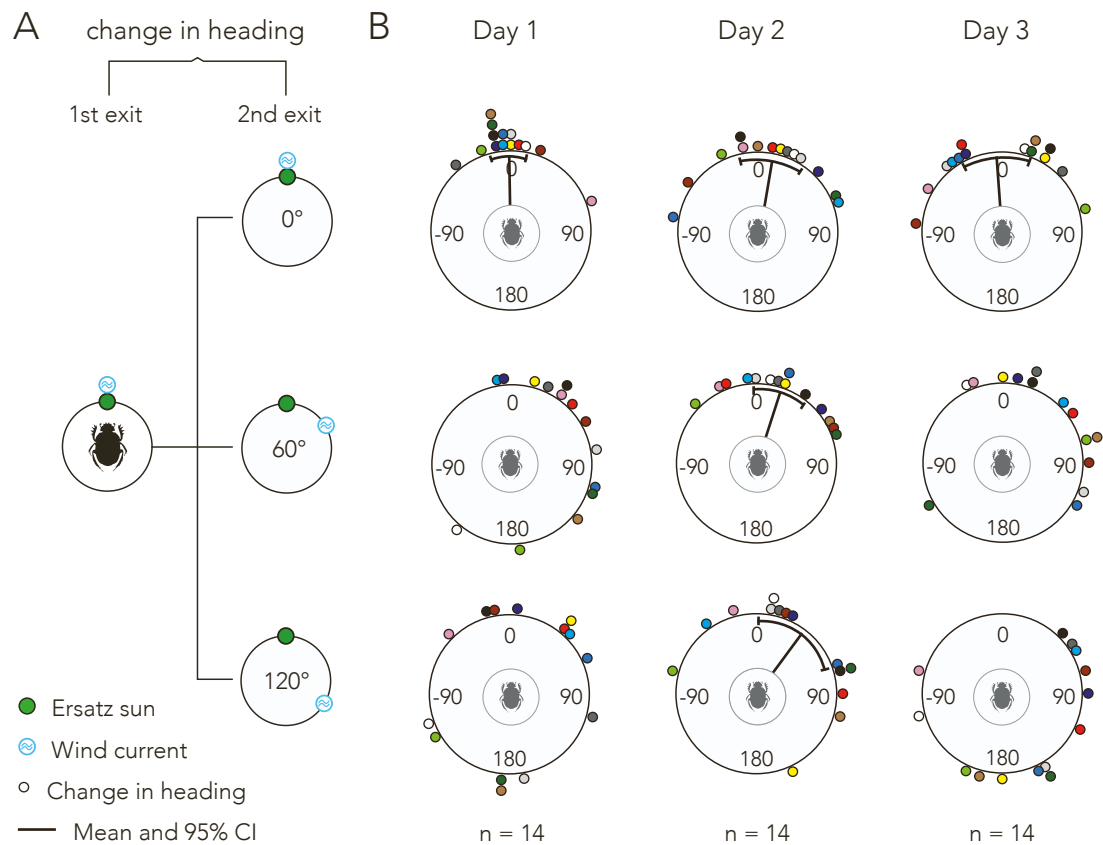


Figure S.3: Changes in heading of the same individuals across three days, related to Figure 2 - The solar elevation and wind speed were set to 60° and 2.5m/s respectively. **A.** Schematic procedure of the cue conflict experiment. Change in heading was calculated between two consecutive exits; initial condition (1st exit) to conflict condition (2nd exit). **B.** Each coloured data point illustrates the change in heading of the same individual across conflicts and days. The line extending from the centre of each circular graph indicates the mean vector and ends in a 95% confidence interval. Note that the data in the third row (120° conflict) is the same as in Figure 2 of this article.

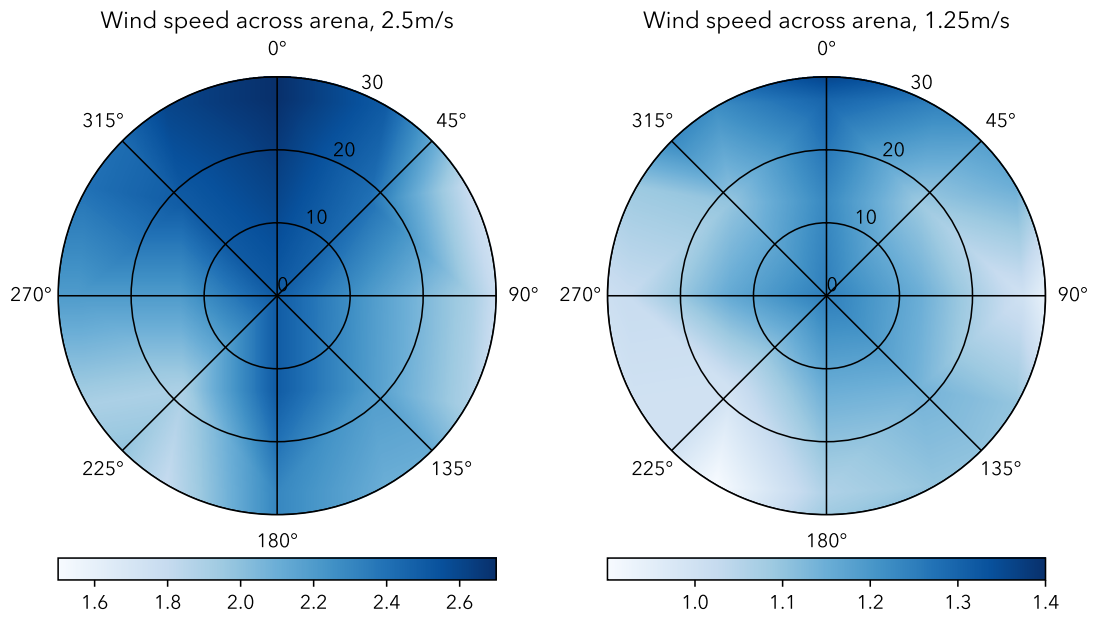


Figure S.4: **Distribution of wind speed, related to Figure 2** - The distribution of wind speeds across our circular arena when the wind speed in the centre is set to 2.5m/s and 1.25m/s. The wind generator is positioned at 0°.




Figure S.5: **The effect of the weight adjustment function g , related to STAR Methods (Equation 16)** - The effect of the adjustment function g when applied to the normalised weight W_{Wind} . After adjustment, the cue with greater weight will dominate the integration. Weights are defined w.r.t. kappas for the cue distributions, see STAR Methods, Integration models.




Figure S.6: **Illustration of the difference between WAM and WVS outputs for different conflicts and weights, related to STAR Methods (Equations 10 and 13)** - Adapted from (Murray and Morgenstern, 2010). The contrasting outputs of the linear weighted arithmetic mean and the circular weighted vector sum. The difference between the two strategies becomes more striking at large conflicts, especially when observing noisy populations (see Figure 3). Experimental conflicts of 60° and 120° shown with dashed lines.




Figure S.7: **The κ approximation from Mardia and Jupp (2009) (Red, dot-dash) compared to the true MLE for κ given by solutions to $R = \frac{I_1(\kappa)}{I_0(\kappa)}$ (Blue, solid), related to STAR Methods (Equation 4)** - We opted for the approximation as it was slightly faster to compute and (as can be seen) sufficiently accurate. The approximation is due to Mardia and Jupp (2009); note that for large R we specifically use Equation 5.3.8 from (Mardia and Jupp, 2009) as it provides a better approximation than their final expression in Equation 5.3.9 (tested in the same manner, by comparing to the true solutions).




Figure S.8: **NVS parameter analysis; selection of parameter a in the weight adjustment function, related to STAR Methods (Equation 16)** - In order to choose a slope parameter a we performed the same likelihood analysis on NVS across a range of parameters. For convenience we defined $a = 0$ such that $g(x; 0) = x$; thus, varying a gives an effective spectrum from 'optimal' weights to winner-take-all. As can be seen, our MLE lies somewhere in the middle of this spectrum, indicating that while one cue will tend to dominate, both are always considered. This analysis indicated the maximally likely slope $\hat{a} = 53$. The flat nature of the peak and the randomness inherent in the population generation process means that \hat{a} would be likely to change if the analysis were run again, however it would remain in the same general area.




Figure S.9: **BVS parameter analysis; selection of σ^2 for Gaussian bias distributions, related to STAR Methods (Equations 17 and 18)** - In order to choose a representative bias variance σ^2 , we performed a likelihood analysis on BVS across a range of parameters. σ^2 is the variance of the Gaussian distribution from which the biases are drawn. This analysis indicated a maximally likely variance $\hat{\sigma}^2 = 0.000303$. Again this is subject to change if the full analysis were to be run again, however the narrow peak indicates this value would be more stable than the slope parameter.

Appendix D

Supplementary information for Paper 2

This file contains the electronic supplementary information for ‘A model of cue integration as vector summation in the insect brain’ (DOI: 10.1098/rspb.2023.0767). Any references to the main text will be made explicit. All other references (equations, figures, sections) are internal to this document. Citation numbering is internal to this document (there is a reference list at the end); most cited works are also cited in the main text.

1 SI: Additional modelling details

1.1 The model in the context of canonical CX neuropils




Figure S1: Example model connections illustrated in the context of canonical CX neuropils. The language of this figure matches that of Figure 2 in the main text; connection colour indicates source, a marker (dot or arrowhead) indicates a receiver. Black dotted lines indicate information flow which was not modelled but for which there is evidence. For example, angular velocity input is thought to arrive in the P-ENs via the noduli [14].

1.2 Model specification

Firing rate model We use the same basic rate encoding as Stone et al. [21]. Neural firing rate (r) is a sigmoid of the input I :

$$r = \frac{1}{(1 + e^{-(aI-b)})} \quad (1)$$

\bar{r} is used for vectors which describe the firing rates over a whole population. Parameters for different neuron types are given in Table S1. Rate parameter selection and tuning was performed by hand.

R neurons Each R neuron has a receptive field centred on its preferred direction: $\theta_R \in \{0, \frac{\pi}{4}, \frac{\pi}{2}, \frac{3\pi}{4}, \pi, \frac{5\pi}{4}, \frac{3\pi}{2}, \frac{7\pi}{4}\}$. The R neuron input is:

Neuron connectivity (R neurons not shown)

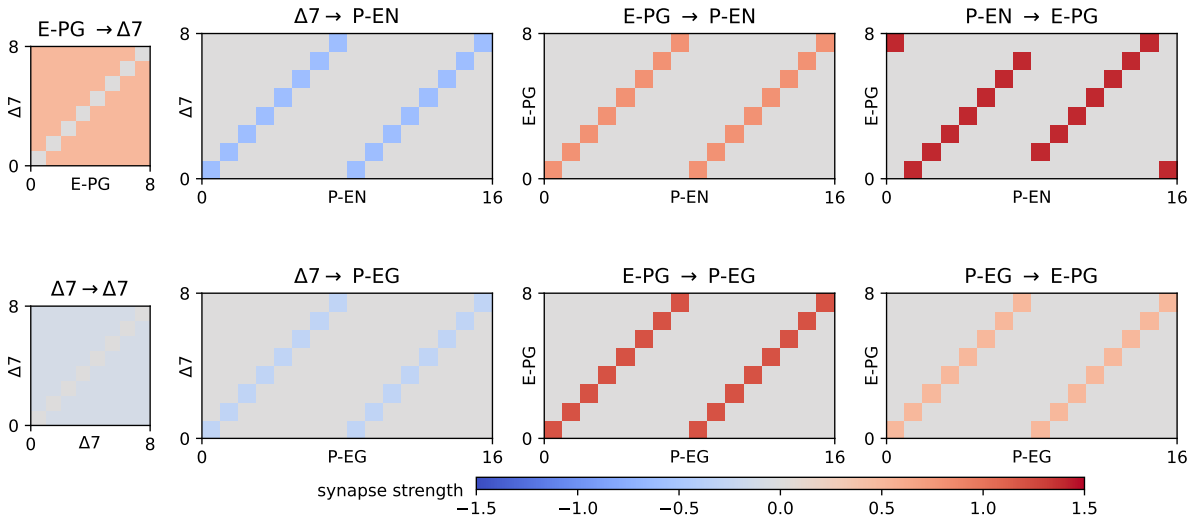


Figure S2: Adjacency matrices documenting neuron interactions in the model. Inhibitory connections are blue, excitatory are red. An absence of connection is grey. $R \rightarrow E$ -PG connections are omitted. These connections represent a functional hybrid of the fruit-fly and locust circuits compared by Pisokas et al. [19]. Also see Goulard et al. [12].

Table S1: The sigmoid tuning parameters used for each neuron type, corresponding to a and b in Equation 1. These parameters were manually tuned to achieve the desired network characteristics.

Neuron type	a	b
R	4	1.8
E-PG	4	1.8
$\Delta 7$	3	3
P-EN	4	3
P-EG	4	5

$$I_{R_{l,i}}(t) = \frac{\cos(\theta_{R_{l,i}} - \theta_C)}{2} + 1 \quad (2)$$

(adapted from Stone et al. [21]), where $R_{l,i}$ is the i th R neuron of group l , θ_C is the azimuthal angle of the input cue, and $\theta_{R_{l,i}}$ is the preferred direction of the R neuron.

R sinusoidal amplitude Sinusoidal amplitudes for R_1 and R_2 are governed by parameters w_1 and w_2 respectively. These amplitude modifiers are applied after the R rates are computed (i.e. we scale the R responses):

$$\bar{r}_l := w_l \cdot \bar{r}_l \quad (3)$$

It is assumed that $w_1 + w_2 = 1$ so that input to the E-PGs remains within an acceptable range, regardless of which amplitude is higher. While we include this normalisation for practical reasons, data from Hulse et al. [14] suggests that different groups of R neurons inhibit each other, which could implement a kind of hierarchical normalisation. Hulse et al. [14] report that this inhibition is not equal between groups which would suggest that one cue’s weight is moderated by another, rather than the cues being normalised against each other. We model the latter so

as not to include assumptions about “preferred” cue types. Beyond the constraint that w_1 and w_2 sum to 1, we can set these weights arbitrarily to explore the effects of anchoring the weights to different properties of the cues.

P-EN neurons We use sixteen P-EN neurons, with input given by:

$$I_{\text{P-EN, Left}}(t) = (1 - h(\mathbf{av})) + W_{\text{E-PG} \rightarrow \text{P-EN}} \cdot \bar{r}_{\text{E-PG}}(t - 1) + W_{\Delta 7 \rightarrow \text{P-EN}} \cdot \bar{r}_{\Delta 7} \quad (4)$$

$$I_{\text{P-EN, Right}}(t) = h(\mathbf{av}) + W_{\text{E-PG} \rightarrow \text{P-EN}} \cdot \bar{r}_{\text{E-PG}}(t - 1) + W_{\Delta 7 \rightarrow \text{P-EN}} \cdot \bar{r}_{\Delta 7} \quad (5)$$

With \mathbf{av} being a signed floating-point number which represents angular velocity in $^\circ/\Delta t$. The function h is given by:

$$h(x) = k \cdot x + 0.5 \quad (6)$$

With $k = \frac{0.5}{24}$ and $h(x)$ clipped such that $0 \leq h(x) \leq 1$. The terms $h(\mathbf{av})$ and $(1 - h(\mathbf{av}))$ above generate the differential speed encoding across the P-EN population while keeping overall activity constant (if one side’s average activity increases, the other’s decreases). The value of k was chosen so as not to saturate or starve any other network components when varying angular velocity. This means the maximum angular velocity which can be detected by self-motion alone is $24^\circ/\Delta t$. This maximum is arbitrary, one could re-tune the network to achieve a better range of responses. In practice, our simulated beetle sometimes exceeds this angular velocity, but the inaccuracy this introduces is effectively corrected by the external cue input; it is only critical to stay within this range during learning.

E-PG neurons E-PG input is given by:

$$I_{\text{E-PG}}(t) = S_{\text{R}} \cdot (W_{\text{R}_0 \rightarrow \text{E-PG}} \cdot \bar{r}_{\text{R}_0}(t) + W_{\text{R}_1 \rightarrow \text{E-PG}} \cdot \bar{r}_{\text{R}_1}(t)) + W_{\text{P-EN} \rightarrow \text{E-PG}} \cdot \bar{r}_{\text{P-EN}}(t) + W_{\text{P-EG} \rightarrow \text{E-PG}} \cdot \bar{r}_{\text{P-EG}}(t - 1) \quad (7)$$

Where S_{R} provides the final, context-dependent scaling of R input to E-PGs. To scale the R input relative to input from other neuron types S_{R} is given explicitly here. In all other cases, sign and magnitude of input is encoded in the adjacency matrix (Figure S2).

$$S_{\text{R}} = \begin{cases} -1.4 \text{ during normal behaviour,} \\ -0.28 \text{ during learning.} \end{cases}$$

In all of our results, this mechanism only serves to facilitate learning (i.e. the scaling factor is only changed during the learning process). Cope et al. [5] similarly adjusted the inputs to their heading neurons during the learning process (where heading was guided only by self-motion inputs). Dan et al. [7] also modelled plasticity between visual ring neurons but they assume that the E-PG bump moves in lock-step with agent orientation without modelling a mechanism to drive the bump while R input to E-PGs is fluctuating due to changing weights. Here we aimed to keep all neural input to the E-PGs active at all times though the balance can shift contextually.

$\Delta 7$ neurons We use eight $\Delta 7$ neurons with input given by:

$$I_{\Delta 7}(t) = W_{\text{E-PG} \rightarrow \Delta 7} \cdot \bar{r}_{\text{E-PG}}(t) + W_{\Delta 7 \rightarrow \Delta 7} \cdot \bar{r}_{\Delta 7}(t - 1) \quad (8)$$

Following Kim et al. [16] and Pisokas et al. [19], we used uniform inhibition - each neuron inhibits all others by the same amount, regardless of position. Kim et al. [16] note that this form of inhibition best matches the observed dynamics in the fruit fly head-direction circuit. In a comparative study, Pisokas et al. [19] observed that the structure of these connections can vary between species. In fruit-flies, each $\Delta 7$ neuron is thought to inhibit all others uniformly; in locusts the inhibition was more concentrated, capturing only those neurons directly opposite (with a couple of neighbours). Pisokas et al. [19] found the different inhibition patterns were functionally identical.

P-EG neurons We use sixteen P-EG neurons with input given by:

$$I_{\text{P-EG}}(t) = W_{\text{E-PG} \rightarrow \text{P-EG}} \cdot \bar{r}_{\text{E-PG}}(t) + W_{\Delta 7 \rightarrow \text{P-EG}} \cdot \bar{r}_{\Delta 7}(t) \quad (9)$$

Learning rotations (‘dances’) Without placing structural constraints on the network, angular self-motion information is required to learn useful mappings (Figure 4C and D, main text). As our inspiration is the ball-rolling dung beetle, we chose to mimic the dung beetle dance, a short stereotypical rotation which occurs at the beginning of the beetle’s orientation behaviour [2]. In addition, this dance is thought to play a critical role in the formation of the dung beetle orientation snapshot [8]. We therefore allow the agent to rotate through 360° in a single direction during learning (in comparison to random walk learning, the learning rotation is extremely time efficient).

It should be noted that Dung beetles do not always perform a full 360° rotation when dancing [2]. The dance may be the last part of the learning process (as there is a lot of movement on-and-around the dung pile). There may also be some predictive capacity; a partial rotation may be enough to construct the rest of the mapping. Even without a predictive capacity, partial rotations may not be useless.

1.3 Model initialisation

In our model, the activity bump in the E-PGs emerges due to input from the R neurons, then propagates through the rest of the network and becomes self-sustaining. To produce this initial bump in the E-PG layer from R input, the R-to-E-PG synapses require some sensible default configuration, described below. This is used in combination with a standard initialisation procedure in which the model rotates on the spot in the presence of cues. Note that after this procedure (once a bump is established) the R-to-E-PG synapses can be reset (e.g. to have all equal weights) so that the default values do not affect what is subsequently learnt.

1.4 Default connections

Each R and E-PG neuron has an associated angle. For R neurons, this angle describes the centre of their receptive field. For E-PG neurons, the angle describes their topographical position within the ring. We set the synaptic strength in inverse proportion to the distance between these R and E-PG angles. To make this general to any R or E-PG population size, the raw synaptic strength is given by

$$w_{R_{l,i} \rightarrow \text{E-PG}_j} = \pi - \cos^{-1} \left(\frac{\mathbf{r}_{l,i} \cdot \mathbf{e}_j}{|\mathbf{r}_{l,i}| |\mathbf{e}_j|} \right) \quad (10)$$

where $\mathbf{r}_{l,i}$ and \mathbf{e}_j are unit vectors with the angle associated to $R_{l,i}$ and E-PG_j respectively. The raw synaptic strengths are then normalised such that all of the connections onto each E-PG sum to 1.

1.5 Behavioural simulations

1.5.1 Procedure

We assume each agent has a (random) desired heading which it tries to maintain [1]. For each step, we compute the signed error (θ_{error}) between the current (θ_c) and desired (θ_d) E-PG headings as follows:

$$\bar{\mathbf{d}} = [\cos(\theta_d), \sin(\theta_d)] \quad \bar{\mathbf{c}} = [\cos(\theta_c), \sin(\theta_c)] \quad (11)$$

$$\theta_{\text{error}} = \text{atan2}(\bar{\mathbf{d}}^\perp \cdot \bar{\mathbf{c}}, \bar{\mathbf{d}} \cdot \bar{\mathbf{c}}) \quad (12)$$

where $\bar{\mathbf{d}}^\perp$ is the unit vector orthogonal to $\bar{\mathbf{d}}$. As $\bar{\mathbf{c}}$ and $\bar{\mathbf{d}}$ are unit vectors, the dot products give the projection of $\bar{\mathbf{c}}$ onto $\bar{\mathbf{d}}$ and $\bar{\mathbf{d}}^\perp$ (i.e. the extent to which $\bar{\mathbf{c}}$ aligns with each of them). This means that $\bar{\mathbf{d}}$ and $\bar{\mathbf{d}}^\perp$ can be thought of as cartesian axes (x and y respectively) and $\bar{\mathbf{c}}$ a vector on those axes. Equation 12 then tells us the polar angle between $\bar{\mathbf{c}}$ and the x -axis (given by $\bar{\mathbf{d}}$). This angle will be positive for an anti-clockwise rotation and negative for a clockwise rotation.

Once we have θ_{error} , the intended turn (T_{intended}) is computed as:

$$T_{\text{intended}} = \begin{cases} -\theta_{\text{error}} & \text{if } |\theta_{\text{error}}| < T_{\text{max}} \\ -T_{\text{max}} & \text{if } |\theta_{\text{error}}| \geq T_{\text{max}} \text{ and } \theta_{\text{error}} > 0 \\ T_{\text{max}} & \text{if } |\theta_{\text{error}}| \geq T_{\text{max}} \text{ and } \theta_{\text{error}} < 0 \end{cases} \quad (13)$$

where $T_{\text{max}} = 24^\circ$ is the maximum turn which can occur in a given timestep and $|\cdot|$ denotes the absolute value. To account for motor error, we add some noise to T_{intended} :

$$T_{\text{actual}} = T_{\text{intended}} + \epsilon_{\text{motor}} \quad (14)$$

where $\epsilon_{\text{motor}} \sim \text{Uniform}(-30, 30)$ generates a random integer angle between -30° and 30° inclusive - bounds chosen so as to loosely approximate the per-step error of *K. lamarcki* [15]. The bounded uniform approximation is chosen as Khaldy et al. [15] report an average standard deviation of motor error but do not specify whether this is circular standard deviation or angular deviation [13] (or an ordinary standard deviation calculation applied to circular data). The approximation appears to be satisfactory but improving this estimate is a goal for future work.

The agent's orientation is updated on each step:

$$\theta_{\text{orientation}} := \theta_{\text{orientation}} + T_{\text{actual}} \quad (15)$$

The cue positions are updated with respect to the agent's orientation:

$$\theta_{C1} = \theta_{\text{orientation}} + \epsilon_{C1} \quad (16)$$

$$\theta_{C2} = \theta_{\text{orientation}} + \epsilon_{C2} \quad (17)$$

where $\epsilon_{C1} \sim \text{von Mises}(\mu_1 = 0, \kappa_1)$ and $\epsilon_{C2} \sim \text{von Mises}(\mu_2 = 0, \kappa_2)$. Note that the means μ_i are always zero (noise is always centred on the true cue position) but concentrations κ_i vary depending on the experiment.

Thus far we have only changed the simulated position of the agent, but we have not updated the model. Model state is updated using a function `update_state(c1, c2, av, w1, w2)`. The parameters `c1` and `c2` are the cue positions, `av` is the angular velocity experienced, and `w1/w2` are the *weights* associated with each cue. Specifically, `w1` and `w2` govern the amplitudes of

the sinusoids encoded by their respective R populations (see Section 2.1 of the main text and Section 1 - R sinusoidal amplitude for how weights are applied). $c1$, $c2$, and av are simply:

$$\begin{aligned} c1 &= \theta_{C1} \\ c2 &= \theta_{C2} \\ av &= T_{intended} \end{aligned} \tag{18}$$

Note that noise is added to the cues around the ‘real-world’ but the turn (av) does not include noise; i.e. the introduced self-motion error is not sensed directly. In other words, av is an efference copy of the intended turn, and any error will be inferred from cue position. Whether P-EN activity is representative of proprioceptive feedback or motor efference is not currently known [22, 9]. $w1$ and $w2$ vary depending on the experiment.

Finally, we update the agent’s cartesian position, $\bar{\mathbf{p}}$:

$$\bar{\mathbf{p}} := \bar{\mathbf{p}} + [s \cdot \cos(\theta_{orientation}), s \cdot \sin(\theta_{orientation})] \tag{19}$$

where s , the step-size is given by:

$$s = \begin{cases} 2.89cm, & \text{if } |T_{intended}| < T_{max}, \\ 0cm, & \text{otherwise.} \end{cases} \tag{20}$$

For large errors, the agent would correct on-the-spot to minimise the effect of the (simple) proportional control mechanism on the exit angles. The step-size of $2.89cm$ is given by Khaldy et al. [15] for *K. lamarcki*. Note that our behavioural simulation does not take account of the parallax effect inherent in our behavioural setup.

1.5.2 Statistics

Our statistics are based on the angle at which the agent exits the arena. Exit angles are rounded to the nearest five degrees for the sake of plotting and comparison to behavioural data. All circular statistics are computed using custom software in Python to allow automation of the simulation, stats computation, and subsequent plot production.

Mean vectors are computed using standard methods from Batschelet [3]. Rayleigh tests are implemented again using standard statistics on the mean vector from Batschelet [3] with significance determined using the methods of Wilkie [24]. The V test is implemented using the methods of Batschelet [3] and additional information from Berens [4] (and by extension [13]).

The Wilcoxon signed rank test given for the precision experiments in Figure 7B (main text) is computed using the SciPy statistics module [23].

1.6 Note on R influence over E-PG neurons

The combined normalisation of cue weights and synaptic strengths leads to a counterintuitive effect where a single cue exerts greater influence over the E-PGs than multiple cues. Assume we have two cues with $w1 = w2 = 0.5$. During learning the total synaptic strength given to each cue (onto each E-PG) will also be 0.5. This means that each cue gets an influence of $0.5 \times 0.5 = 0.25$. The total influence from the Rs is summed as 0.5 and this determines how the R neurons in combination interact with angular velocity and past orientation information (P-ENs and P-EGs). Now assume a single cue with $w1 = 1$. During learning, the total synaptic strength onto each E-PG will also be 1. Total influence given to the cue is then $1 \times 1 = 1$ giving twice the overall R influence compared to the dual cue scenario above. This does not predictably influence behaviour.

This effect could be reflected in reality or, more likely, E-PGs could adapt to keep total R influence constant. It could be useful to know if or how E-PG activity changes in the presence of multiple cues.

2 SI: Additional results

2.1 Simulated behaviour - Modality transfer [6]




Figure S3: Simulated results from a modality transfer experiment which replicates a previous behavioural experiment by Dacke et al. [6]. The behavioural data is included here (with permission) in the right-hand plot (replicating Figure 4A from [6]).

Dacke et al. [6] demonstrated that the ball-roller *K. lamarcki* was able to learn to orientate to a wind cue placed at 90° to an existing sun cue (a green LED); when the ersatz sun was removed, the beetles kept their bearing. Dacke et al. [6] describe this as a modality transfer, with the beetles transferring information between a sun and wind compass. Here we explain this behaviour as cue synchronisation; any new cue positions are learned with respect to existing $R \rightarrow E$ -PG mappings. These are behaviourally equivalent but we interpret this system as a centralised compass where inputs are synchronised rather than multiple independent compass systems.

To simulate this experiment we assume that the cues are of equal reliability ($\kappa_1 = \kappa_2$) and equal weight ($w_1 = w_2$); this is a simplifying assumption. The experiment is then simulated as described by Dacke et al. [6]. $n = 40$ beetles exit the arena four times each. On the first roll they have a single cue, on the second and third they have a second cue added at 90° to the initial cue, and on the final roll the initial cue is removed. If cues are synchronised, the changes in bearing between the first and fourth roll should be clustered around zero (determined using a V test, as in Dacke et al. [6]).

As in the beetles, our model was able to integrate the new cue alongside an existing one; when the first was removed, the population largely retained their initial bearing (change in bearing clustered around 0).

2.2 Simulated behaviour - Cue conflict

Here we simulated the cue conflict experiments performed by Shaverdian et al. [20]. Beetles were presented with an ersatz sun (LED) cue and wind cue. The wind speed and light elevation were changed in order to manipulate the reliability of the cues. Beetles performed eight rolls each. Three with cues aligned, one with cues in conflict, one with cues aligned, one with cues in

conflict, then two final rolls as controls. Each beetle experienced a conflict of 60° or 120° , but the order in which these conflicts were experienced was randomised. Light elevations of 45° , 60° , 75° , and 86° were tested with a wind speed of $2.5m/s$, and elevations of 60° , 75° , and 86° were further tested with a wind speed of $1.25m/s$. As noted in the main text, it is currently not clear how cues are weighted by the beetles. We therefore tried weighting cues according to the adjusted reliability used by Shaverdian et al. [20] and compared against a rough contrast model to see if they produced qualitatively different outcomes.

We devised a simple weight model based on a hypothetical measure of cue contrast. If we draw a unit vector from the centre of the arena to a simulated sun, then project this into the ground plane, the magnitude of this projection gives us a measure of the contrast across the sky. If the light elevation is given by ϕ then the contrast would simply be:

$$s_{\text{light}} = \cos(\phi) \quad (21)$$

For wind we do not have such a neat formulation. We made the assumption that contrast was linear with wind speed. We constrain the model for wind speed so that the relative contrast of the wind at $2.5m/s$ is the same as the light at an elevation of 60° . This constraint is based on the observations of Shaverdian et al. [20]; they note that this is the point at which cues appear to be equally weighted. We chose:

$$s_{\text{wind}} = a(\text{speed}) + \left(\frac{1}{2} - a(2.5)\right) \quad (22)$$

with $a = 0.3$. The wind and light contrast models are shown in Figure S4. The effect of constructing cue weights based on contrast and adjusted reliability can be seen in the top and bottom rows.

To compare the models, we simulated the cue conflict experiment from [20] using both weighting models. The resulting conflict plots are shown in Figures S5 and S7 for the contrast model, and Figures S6 and S8 for the adjusted reliability model [20]. For cue noise, we use the same kappa estimation as Shaverdian et al. [20]; while it is not clear that reliability sets cue weight, single-cue performance could still estimate reliability.

Our results are mixed. Using the contrast model in Figure S4, our simulated beetles do turn (or not turn) under the correct combinations of light elevation and wind speed. The turns are not cleanly clustered around the conflicting cue locations. The means are not always centred on the ‘preferred’ cue locations, however, this is to be expected with our more comprehensive simulation. Interestingly, despite the clear difference in the weight models (compare weight maps in Figure S4), simulations which use the weighting model from [20] ends up producing very similar data. Shaverdian et al. [20] noted that with increased conflict, there was a systematic change in mean vector length; mean vectors shrank as conflict increased. We do not see any such effect in our simulations (via the proxy of circular standard deviation, s_0).

Shaverdian et al. [20] also saw variation in behaviour at large conflicts where cues were thought to be equal in weight. In our case, such variation did not depend on the balance of weights but rather depended on the amount of noise present in the cues. As an example from our simulations, 86° elevation in combination with $1.25m/s$ wind speed tends to be a very noisy condition despite the fact that cues are not equal in weight. It was not uncommon to find cases (different random initialisations) where the confidence limits could not be computed because the mean vector length was insufficient.

Despite implementing the same exclusion criterion (removing disoriented agents), we found that we generally retained more agents than Shaverdian et al. [20]. This likely means that our noise estimates (both for cue reliability and motor error) are optimistic, as our agents remained more directed than the real beetles. As noted in the main text, our simulations do not include parallax which may also contribute to our higher agent retention.

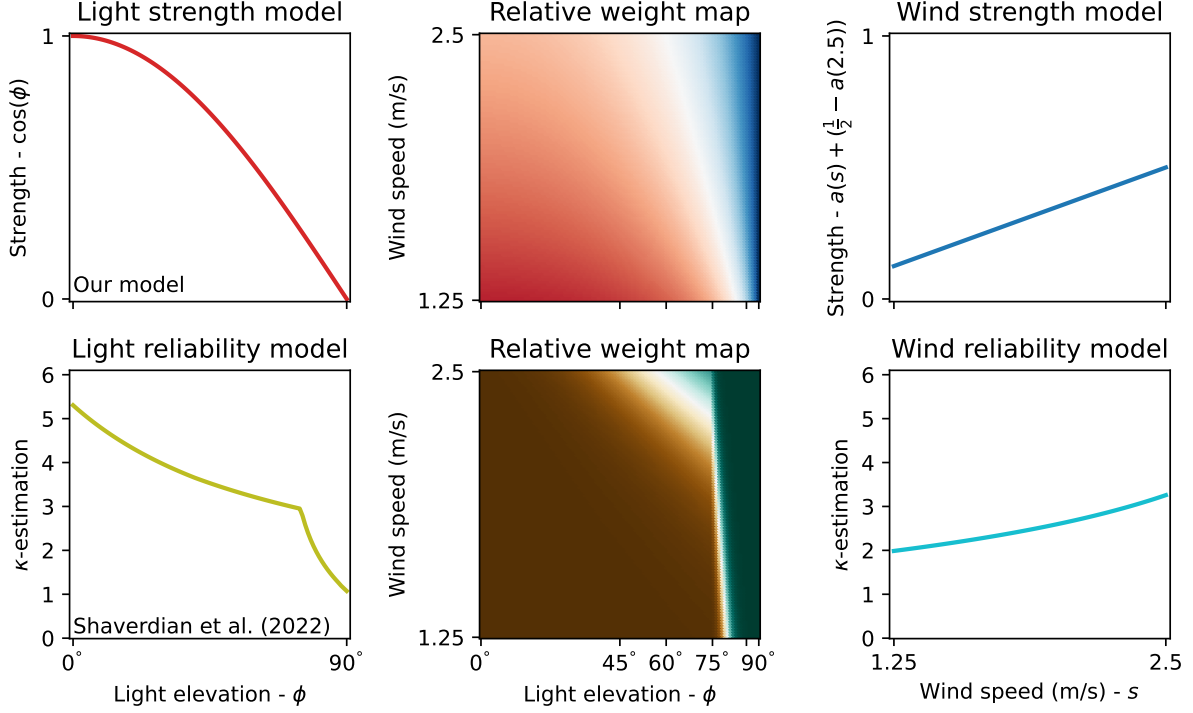


Figure S4: **Top row:** Light and wind contrast models (left and right respectively) and the relative weight afforded each cue if cues are weighted by contrast (centre). On the relative weight map, red indicates a preference for light, blue indicates a preference for wind. **Bottom row:** The light and wind reliability models from Shaverdian et al. [20] (left and right respectively). Note that these lines are generated using their full κ -estimation process. In the centre plot, we show the weights given to each cue if one uses the sigmoid adjusted relative reliability as used by Shaverdian et al. [20]. In both relative weight maps, the unmarked tick on the x-axis is for an elevation of 86° , omitted for readability.

In summary, our simulations suggest that we really do not have a clear picture of cue reliability, contrast, weight, or influence in the beetles. Further, our simulations suggest that two different weighting strategies can produce very similar outcomes which further illustrates the difficulty in determining a cue conflict strategy using behaviour alone.

2.3 Insensitivity to R neuron population size

In circuit construction we chose to use eight R neurons per group. There are three primary reasons: (1) the default R-to-E-PG mapping was easy to construct and visualise, (2) precise R neuron counts (on a group level) are not yet available, and (3) our encoding should be invariant to population size. Here we verify that the choice of eight R neurons was not critical to the proposed function. The learning rule does not need to be changed depending on the number of neurons. There are two scenarios we explored: equal or unequal group size.

2.3.1 R groups of the same size

We chose R groups of size $n = 3, 6, 12$, and 24 . This gives a reasonable range of group sizes and should also be sufficient to reveal any variability. These group sizes are all smaller than we might expect to find in the brain [18, 14], however, as we will show, increasing in group size seems unlikely to cause problems. The results here shown in Figures S9 and S10. Again, weights are applied by manipulating the amplitude of the R neuron signal. In all cases, cues

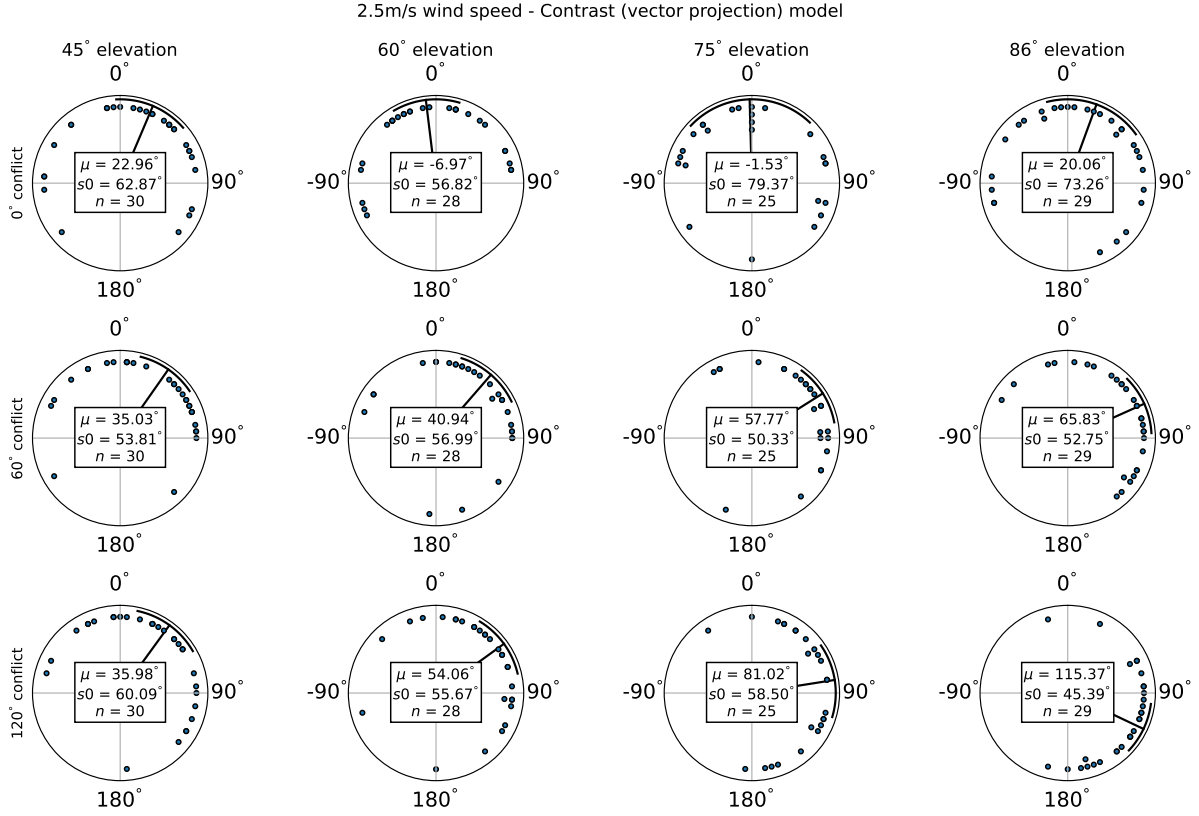


Figure S5: Cue conflict results for 2.5m/s wind speed. Different conflicts are shown in each row and different elevations are shown in each column. Rayleigh statistical significance ($p < 0.05$) is achieved in all cases. μ is the angle of the mean vector, $s0$ is the circular standard deviation, and n is the number of agents included in the analysis. Some agents are excluded from the analysis as they are not orientated when cues are aligned; we use the same exclusion criterion as Shaverdian et al. [20]. Arcs indicate the 95% confidence interval for the mean.

are aligned, of equal weight, equally useful, and noiseless.

Learned mappings are presented in the left-hand columns of Figures S9 and S10. In all cases, the adjacency matrices form diagonal patterns (though all are horizontally shifted). The connections appear stronger with fewer neurons which is an artifact of the normalisation process. In all four cases, the connections in each row must sum to 1; the more columns you have, the more spread-out the synaptic strength will be. This does not seem to cause any issues with the cue integration function of the network.

Generally, the angular vector sum is preserved (compare right-hand plots with Figure 4A, main text). Distortion is seen for the minimum example of $n = 3$ R neurons (Figure S9a), with the most obvious change being the behaviour for equal weights (purple line) and large conflicts. Minor distortion continues to be present in all cases but the results better approximate vector summation with more R neurons. This distortion is similar to that seen in the base case of $n = 8$ neurons (Figure 4D, main text); given that the distortion is similar in all cases it is likely that the distortion is caused by the learning process rather than the R population size variation. The angular output of the network for varying cue conflict does not appear to be seriously affected by the number of R neurons which provide input.

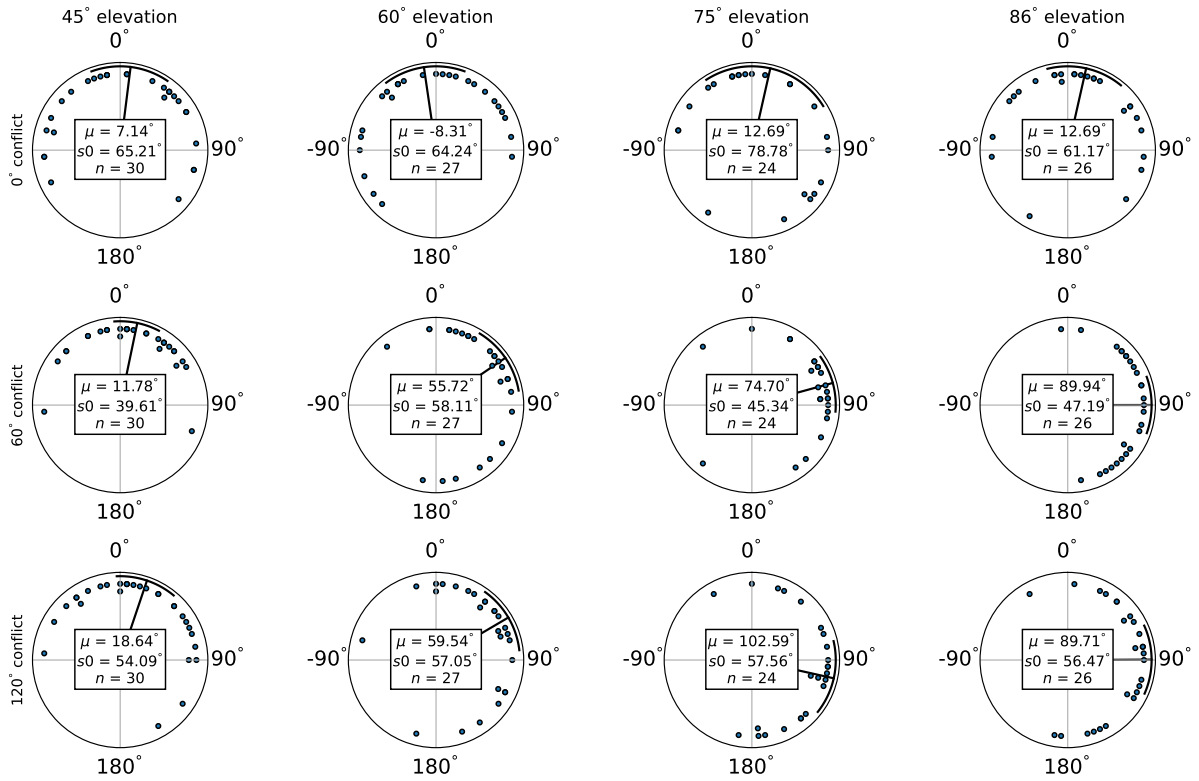


Figure S6: The same simulated experiment as Figure S5, this time using the relative reliability weight model from Shaverdian et al. [20] (Figure S4 - bottom row). Different conflicts are shown in each row and different elevations are shown in each column. Rayleigh statistical significance ($p < 0.05$) is achieved in all cases. μ is the angle of the mean vector, s_0 is the circular standard deviation, and n is the number of agents included in the analysis. Some agents are excluded from the analysis as they are not orientated when cues are aligned; we use the same exclusion criterion as Shaverdian et al. [20]. Arcs indicate the 95% confidence interval for the mean.

2.3.2 Unequal numbers of R neurons

The network response changes if we use different numbers of R neurons per group. Greater influence is given to the cue with more R neurons. It is not entirely clear why this is the case. More R neurons could simply give more opportunities for synaptic strengthening and thus a greater overall influence. This was an unexpected consequence of the learning rule and ends up making an interesting prediction. If R neuron groups do not contain equal numbers of neurons then we would expect groups which respond to ‘primary’ cues (cues to which an animal is known to be more responsive) to have more neurons than secondary cues. This need not be the case, but adds an interesting extra component to the idea of cue influence. Further, this highlights the importance of examining the mechanism (even an abstract neural model) by which a process may be taking place.

2.4 Effect of existing snapshots

2.4.1 Noiseless cues, updated relationship

In all basic testing, the model is initialised, and then weights are flattened prior to learning. This was to ensure previous snapshots were not simply being retained and that we could actually construct usable mappings where none existed. However, it is also useful to consider how the

1.25m/s wind speed - Contrast (vector projection) model

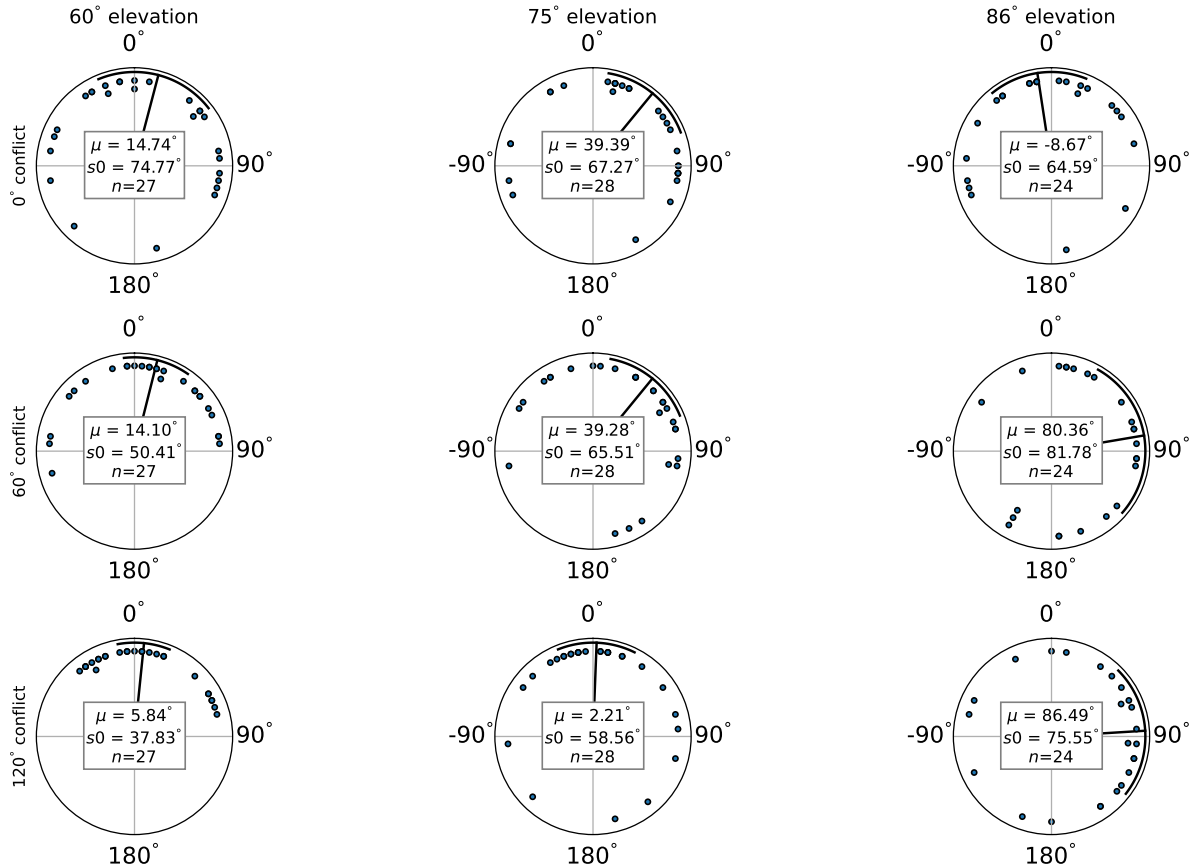


Figure S7: Cue conflict results for 1.25m/s wind speed. Different conflicts are shown in each row and different elevations are shown in each column. Rayleigh statistical significance ($p < 0.05$) is achieved in all cases. μ is the angle of the mean vector, s_0 is the circular standard deviation, and n is the number of agents included in the analysis. Some agents are excluded from the analysis as they are not orientated when cues are aligned; we use the same exclusion criterion as Shaverdian et al. [20]. Arcs indicate the 95% confidence interval for the mean.

model behaves when a mapping is actually present (as is the case in our behavioural simulations).

We first allow the network to learn a mapping between two cues, both are aligned but cue one is stronger. We then flip the weight relationship and move cue two by 179° . The network is then allowed to experience a number of learning iterations to see how many it takes for the snapshot to update to reflect the new cue state Figure S12 shows the results.

It appears that the spatial relationship is updated on the first learning iteration, however the increased synaptic strength of cue one allows it to retain disproportionate influence despite the updated inputs. It takes an additional iteration to encode an accurate weight relationship in the snapshot. There are two potentially interesting follow-ups here.

First of all, existing information does affect the information that is learned but only the relative influence. This means that an agent which already has a strong mapping for a particular cue would need multiple exposures to an alternative highly-weighted cue to fully update the relationship. This could be viewed as a form of prior, the cue was weighted favourably in the past so the agent ‘assumes’ it will be good in the future. The perceptability of cues in nature is unlikely to change as suddenly as we have modelled here; thus, the fact that the relative influence of the first cue remains high after the first learning iteration could point to a relatively strong influence for prior experience.




Figure S8: The same simulated experiment as Figure S7, this time using the relative reliability weight model from Shaverdian et al. [20] (Figure S4 - bottom row). Rayleigh statistical significance ($p < 0.05$) is achieved in all cases. μ is the angle of the mean vector, s_0 is the circular standard deviation, and n is the number of agents included in the analysis. Some agents are excluded from the analysis as they are not orientated when cues are aligned; we use the same exclusion criterion as Shaverdian et al. [20]. Arcs indicate the 95% confidence interval for the mean.

Secondly, though it is minor, the snapshot does appear to solidify with repeated exposures to the same stimulus (compare the 2nd and 3rd iterations). Assuming cues stay consistent in the environment, this may indicate that multiple learning events can improve the mapping that the agent has, thereby improving orientation performance.

2.4.2 Repeated learning with noisy cues

Figure S13 shows how snapshots change over multiple learning rotations using noisy cues. Both cues were given equal weight, cue one (1st and 3rd columns) had a concentration of $\kappa_1 = 2$ and cue two (2nd and 4th columns) had a concentration of $\kappa_2 = 0.5$ (to match the rest of our experiments). We found that the effect of noise on the snapshot was highly variable so show a ‘good’ (left) and a ‘bad’ (right) example.

On the left-hand side, the repeated learning rotations appear to filter noise out of the snapshot, in essence, repeated learning averages the motion experienced to form a clean map from R to E-PG neurons. On the right, the effect can still be seen, however, the final snapshot is still corrupted for the noisy cue (Figure S13Dii, right-hand plot). Further, the synaptic connections appear to be weaker overall suggesting that this mapping would give less influence




Figure S9: Companion to Figure S10. Visualisation of R-to-E-PG connections and resultant conflict behaviour for different numbers of R neurons. In all cases there are eight E-PG neurons. **(a)** Three R neurons per group, connected to eight E-PGs. Minor distortion can be seen in all cases. The only major difference in output is for equal weights (purple) which stabilises to follow one of the cues as opposed to resting between them. **(b)** Six R neurons per group, connected to eight E-PGs. Minor distortion can be seen for all weight combinations.

to the noisy cue.

Thus, where cues are sufficiently noisy, they may receive less synaptic weight in the final snapshot. However, the effect was not consistent enough to provide a reliable weighting effect across a population (Figure 3, main text). This inconsistency could help to explain the highly variable results we observed when examining the weighting effects of noise in a simulated conflict experiment.

2.5 Basic characteristics

Figures S14 and S15 demonstrate two basic characteristics of the network; namely that the E-PG ‘bump’ persists in the absence of external cue input, and that external cues primarily dictate the bumps position where they are present.




Figure S10: Companion to Figure S9. Visualisation of R-to-E-PG connections and resultant conflict behaviour for different numbers of R neurons. In all cases there are eight E-PG neurons. (a) Twelve R neurons per group. Minor distortion can be seen in all cases. (b) Twenty-four R neurons. Minor distortion can be seen in all cases.

2.6 Partial rotations

In Figure S16 we plot $R \rightarrow E$ -PG mappings and conflict outputs for different degrees of rotation.

2.7 Continuous learning

Here we tested whether or not the network would operate sensibly under continuous plasticity. While this paper was being finalised, Fisher et al. [11] demonstrated that self-motion is actually required for activation of ExR2 neurons; dopaminergic neurons (DANs) which appear to drive the plasticity in the $R \rightarrow E$ -PG connections. This aligns nicely with our results which indicate that self-motion is required to learn useful information; however, at the time we did not have any candidate for plasticity regulation and so decided to test what would happen if plasticity were enabled continuously. After all, many insects do not dance, and $R \rightarrow E$ -PG connections had previously been shown to update in freely walking fruit-flies [10, 17] (with no specific ‘learning rotation’).

We speculated that plasticity could be up-regulated or down-regulated but would be continuously




Figure S11: Variation in network behaviour when different numbers of R neurons are used for each group. In both cases the network shows preference to R_2 , the group with the greater number of neurons.

active at varying rates. We allowed an agent to perform random walks with two (noiseless) cues present. During these walks, plasticity was enabled constantly and we varied the learning rate to see if there was a sweet-spot where the connections would remain stable during normal behaviour, but update where cues were put in conflict. The agents all perform the same walk but their learning rates differ which causes deviation in the heading (read from the E-PG neurons) over time. The results are shown in Figure S17.

Where the learning rate is high ($\eta = 0.1$), it can be seen that the map is quite unstable and changes as the agent moves. While there is always a usable head direction signal, it is important to note that the instability in the $R \rightarrow E$ -PG mappings would make any navigational behaviour which depends on external cues (e.g. path integration or straight-line orientation) impossible.

Decreasing the learning rate ($\eta = 0.01$) improves stability somewhat. The mapping is still clearly perturbed but remains fairly consistent where cues are not being manipulated. Where cues are shifted relative to one-another, the mapping updates almost in lockstep with the cue shift (Figure S17; R maps for $\eta = 0.01$, segment between the magenta and second red line). As both cues are equal in weight, both mappings are shifted to give a final separation of around 180° which reflects the updated cue relationship. Finally, for an immediate conflict (green line), the map update is similarly immediate. Interestingly the R_1 map is not shifted (as in the gradual conflict case) despite the cues having equal weight; this could be to do with the fact that the first cue remains consistent with the bump position, meaning that the existing R_1 map is reinforced rather than updated.

Finally, the slowest learning rate we tested ($\eta = 0.001$) gave the most stable mappings but still allowed mappings to be updated where cue relationships changed. This case is what we would consider the ideal for continuous baseline plasticity. In natural conditions, it is unlikely that relative cue positions would change as quickly as presented here. Thus, any large change in input should be treated with caution in the light of previously consistent information. This again indicates that an existing mapping can act as a form of expectation based on prior experience. Change is permitted and useful, however, mappings are only updated once enough evidence has accrued that the positions really have changed.

It would be interesting to revisit this experiment using a variable learning rate coupled to self-motion though it is likely we would observe similar results to those of Dan et al. [7].




Figure S12: **Summary:** how long does it take to learn a new mapping? A base relationship is learned to begin with. Cue two has $\frac{1}{4}$ the weight of cue one (as in Figure 5C, main text), we then allow the agent to perform a sequence of dances to see how quickly the mapping can be updated. **(a)** Cues are aligned and C1 has four times the weight of C2, this is the same snapshot as is learned in Figure 5D, main text. **(b)** 1st iteration with updated cues. Here the weight relationship is flipped and C2 is moved to be 179 degrees away from C1. The spatial relationship changes to match the cue positions within one learning event. However, the relationship in the map does not fully reflect the change in inputs. After one learning event, the cues appear to have approximately even synaptic strength. **(c)** 2nd iteration with updated cues. The updated weight relationship is now stored in the snapshot and mappings appear a little less spread. **(d)** 3rd iteration with updated cues. The relationship now appears to be stable, the old snapshot has been completely overwritten.




Figure S13: Demonstrating (1) how snapshots evolve given repeated exposure to noise, and (2) how variable this evolution can be depending on the noise experienced. On the left we show a ‘good’ example where noise is essentially filtered out over repeated learning events. On the right is a ‘bad’ scenario where the noisy cue ends up with a corrupted snapshot. The only difference between these simulations was the noise experienced; the parent distributions were the same. Noise does not have a reliable effect on snapshots.

2.8 Random walk learning

A previous model by Cope et al. [5] used a random walk to drive the learning process. Here we verify that our model can similarly learn using a random walk and that experiencing cue rotation in combination with angular velocity is all that is required to build usable maps.

Agents performed a random walk (driven by a Gaussian process), mappings were wiped, and plasticity was enabled for a period of the walk. We tried varying the time over which plasticity was enabled as well as the scale parameter (standard deviation) of the Gaussian process which directed the walk. Cues were equally contrasted, useful, and noiseless. In each case we show




Figure S14: E-PG bump persists in the absence of cue input. Some drift is visible; at the point where external cues are re-introduced, the E-PG bump does not exactly match the position given by the R neurons.

three snapshots, one at the start of learning (always blank) one half-way through the learning period, and one at the end. The learning rate was always $\eta = 0.1$, as in a normal dance. While varying standard deviation, the learning period was kept constant at 50 timesteps. Similarly, while varying learning period, standard deviation was kept constant at 20.

We generally found that the quality of the mapping depended on the amount of motion experienced which could be affected by learning period duration or having a greater amount of rotation over a smaller learning period (greater standard deviation). In all cases presented, the maps are not great. This is a combination of the fact that the learning rates were high and that plasticity continues with limited rotation, meaning cue relatively static cue positions are over-learned [11].

E-PG bump primarily driven by cue input




Figure S15: R neurons primarily drive the position of the compass bump. Cues are fixed in place relative to the agent while self-motion continues to be received as normal via the P-ENs. The E-PG bump remains stable. When cues are again allowed to move, the bump moves with them.

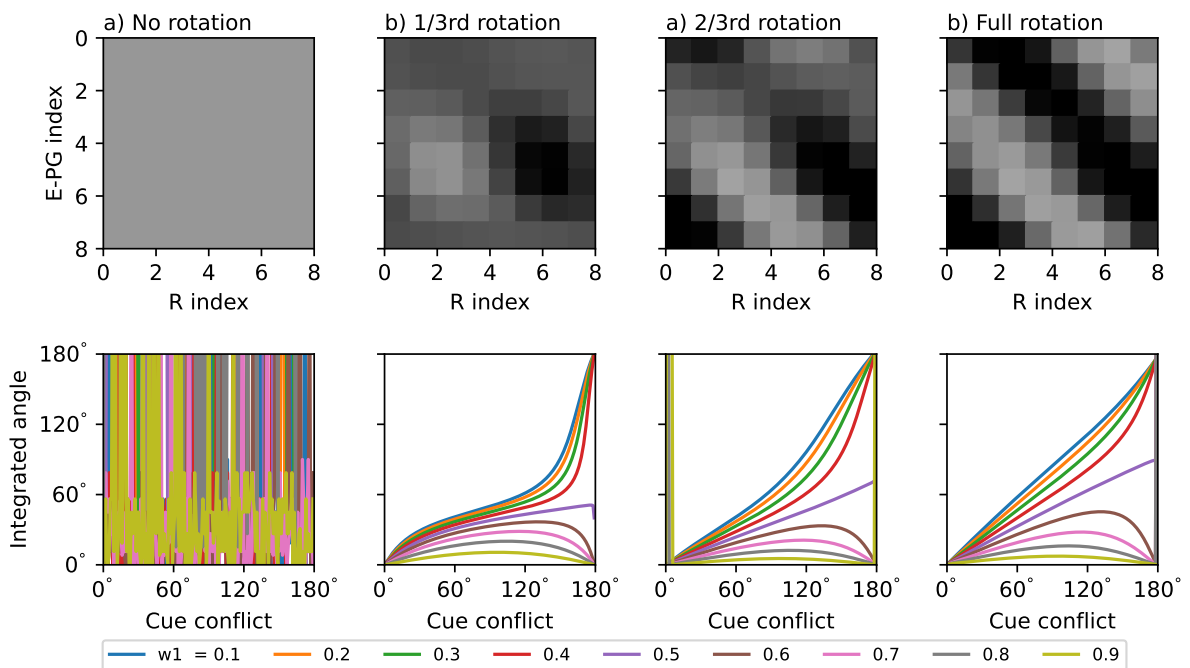


Figure S16: Conflict/mapping output for different degrees of rotation.

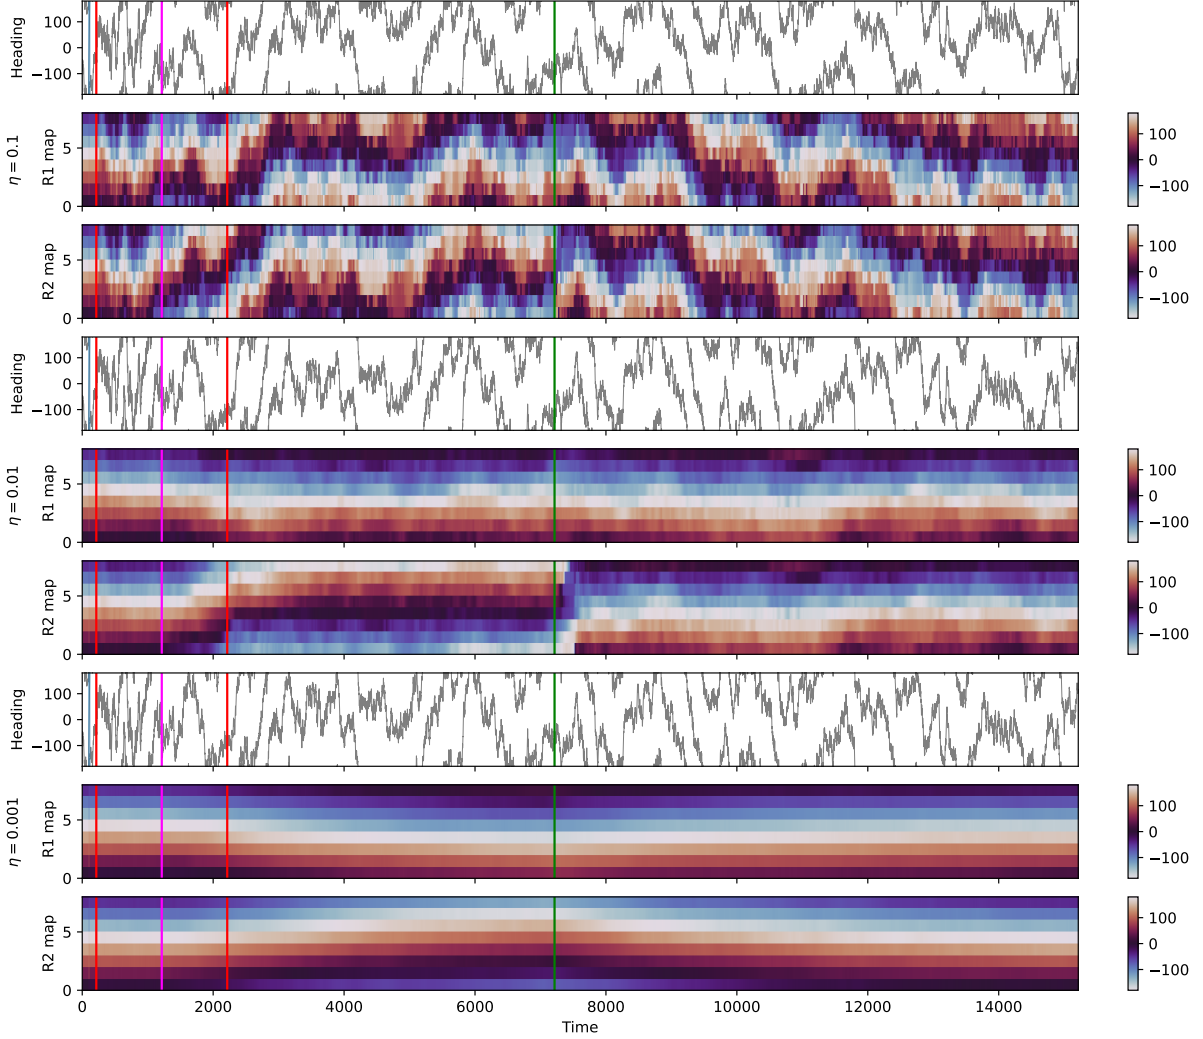


Figure S17: Plasticity constantly enabled at different learning rates. Each walk gives three outputs: Heading, R1 map, and R2 map (i.e. this figure shows three separate walks). The heading shows the population vector average of the E-PG neurons at each timestep. The mappings indicate the average $R \rightarrow$ E-PG mapping. Each row represents an R neuron and the colour indicates where, on average, each neuron maps to on the E-PG ring. Note that this is an indicative measure and does not guarantee that the mapping forms the desired diagonal pattern (see Figure 4, main text). The vertical coloured lines indicate points where the cues are manipulated. Cues start aligned, have equal weight, and are noiseless. **Sequence:** The agent starts with a default mapping and moves for some time. The blue line (heading traces) indicates a dance, the first red line indicates that plasticity is enabled (continuous learning begins here). The magenta line indicates the introduction of a gradual cue conflict; between here and the second red line cue two shifts to be 180° opposite cue one. The second red line indicates that cues are again fixed with respect to with respect to each other. Finally the green line indicates an immediate conflict, where cue two is instantly shifted by 180° (back to where it started w.r.t. cue one).

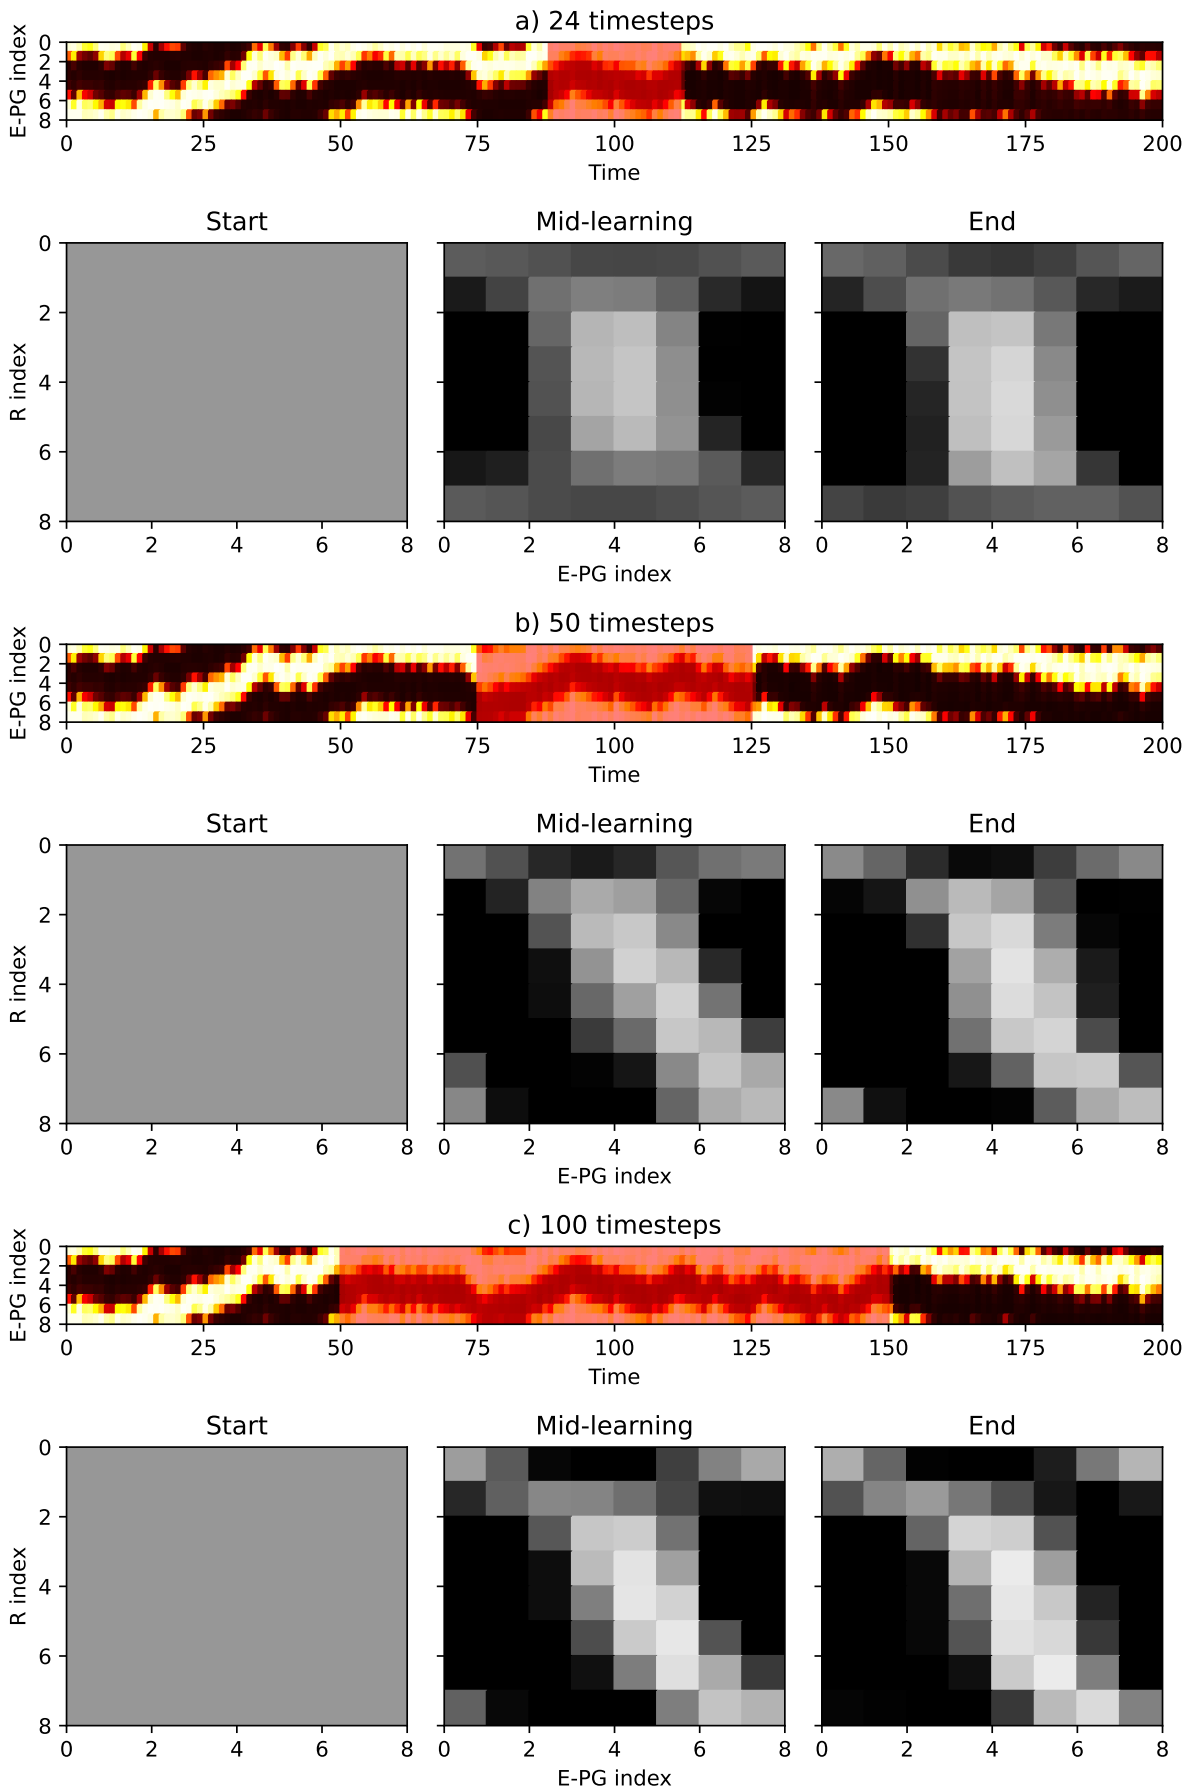


Figure S18: Random walk learning with varying learning times.

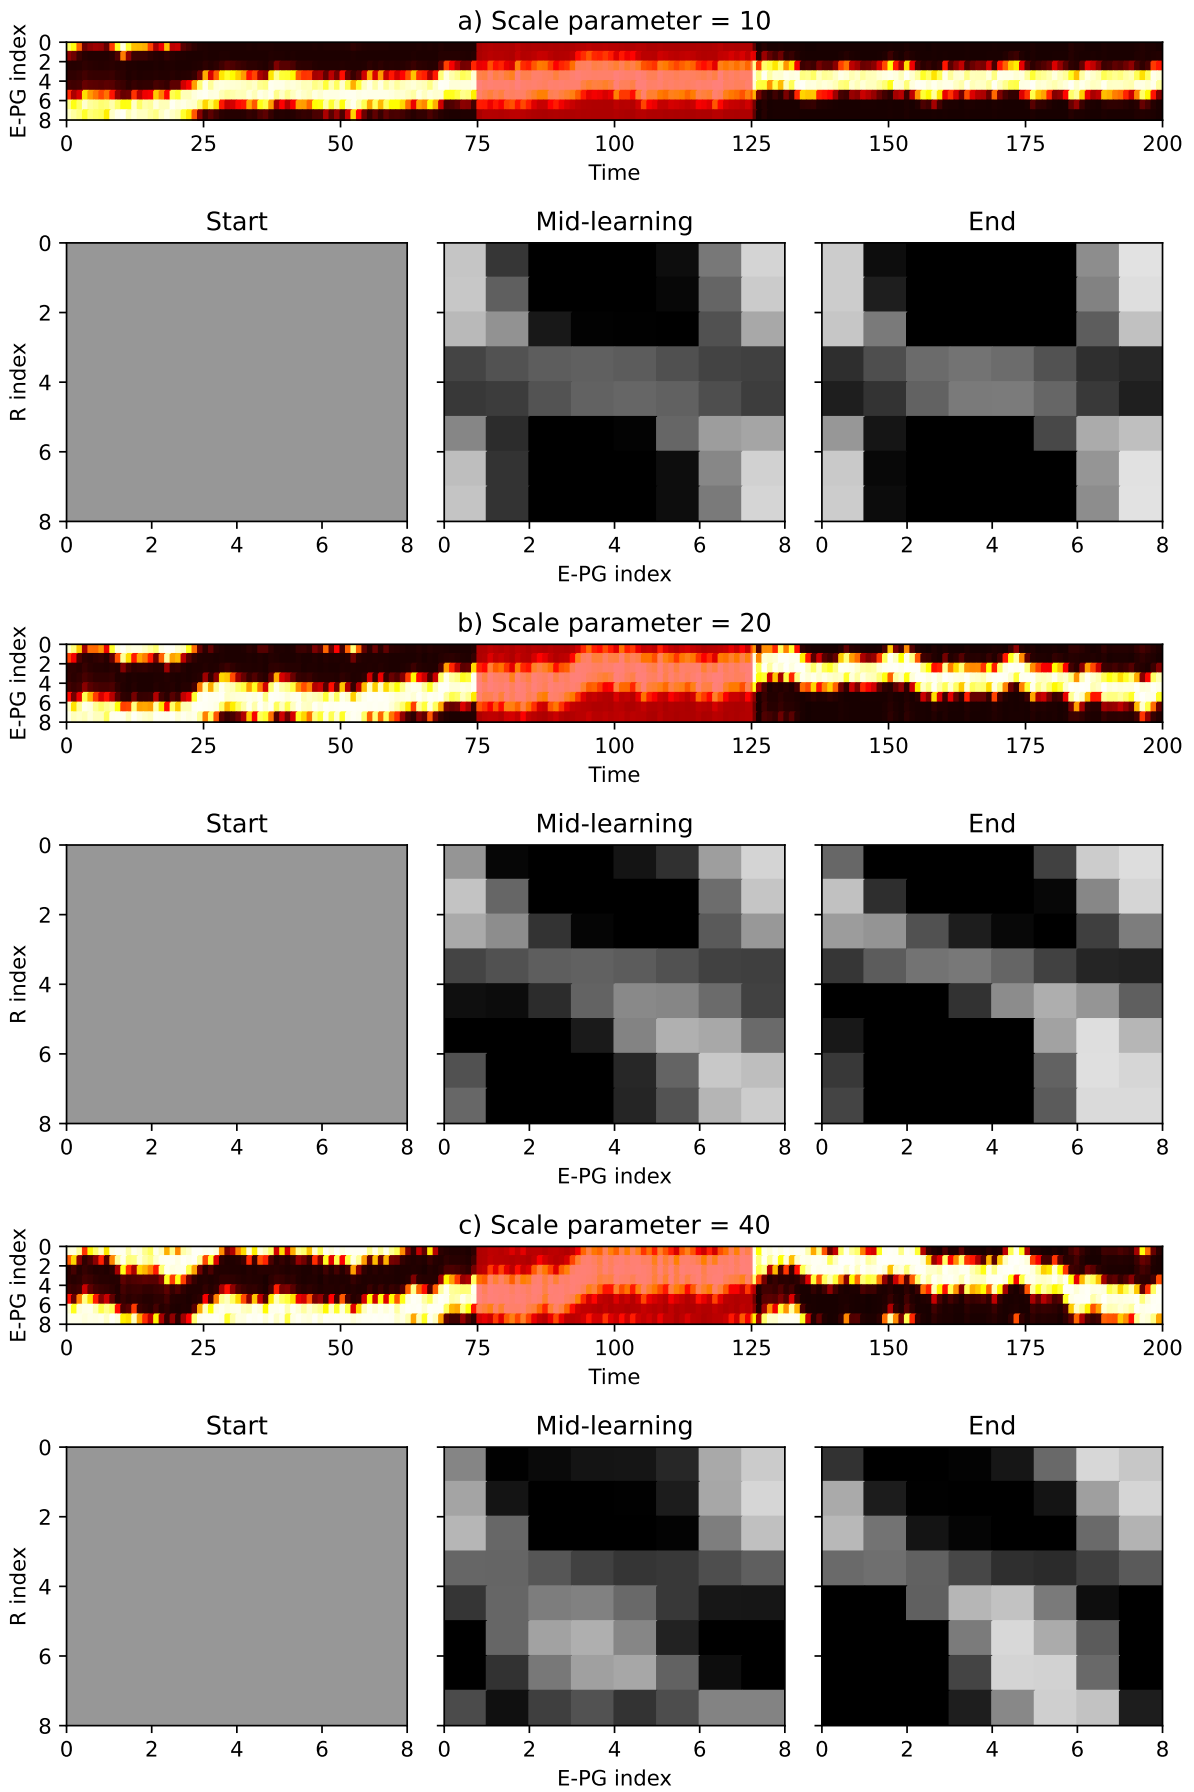


Figure S19: Random walk learning with varied standard deviation.

References

- [1] Emily Baird, Marcus J. Byrne, Clarke H. Scholtz, Eric J. Warrant, and Marie Dacke. Bearing selection in ball-rolling dung beetles: Is it constant? *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, 196(11):801–806, 2010. ISSN 03407594. doi: 10.1007/s00359-010-0559-8.
- [2] Emily Baird, Marcus J. Byrne, Jochen Smolka, Eric J. Warrant, and Marie Dacke. The dung beetle dance: An orientation behaviour? *PLoS ONE*, 7(1), jan 2012. ISSN 19326203. doi: 10.1371/journal.pone.0030211.
- [3] Edward Batschelet. Circular statistics in biology. *ACADEMIC PRESS, 111 FIFTH AVE., NEW YORK, NY 10003, 1981, 388*, 1981.
- [4] Philipp Berens. Circstat: A matlab toolbox for circular statistics. *Journal of Statistical Software*, 31(10):1–21, 2009. doi: 10.18637/jss.v031.i10. URL <https://www.jstatsoft.org/index.php/jss/article/view/v031i10>.
- [5] Alex J. Cope, Chelsea Sabo, Eleni Vasilaki, Andrew B. Barron, and James A.R. Marshall. A computational model of the integration of landmarks and motion in the insect central complex. *PLoS ONE*, 12(2):1–19, 2017. ISSN 19326203. doi: 10.1371/journal.pone.0172325. ISBN: 1111111111.
- [6] Marie Dacke, Adrian T.A. Bell, James J. Foster, Emily J. Baird, Martin F. Strube-Bloss, Marcus J. Byrne, and Basil El Jundi. Multimodal cue integration in the dung beetle compass. *Proceedings of the National Academy of Sciences of the United States of America*, 116(28):14248–14253, 2019. ISSN 10916490. doi: 10.1073/pnas.1904308116.
- [7] Chuntao Dan, Ramya Kappagantula, Brad K. Hulse, Vivek Jayaraman, and Ann M. Hermundstad. Flexible control of behavioral variability mediated by an internal representation of head direction. *bioRxiv*, page 2021.08.18.456004, August 2022. doi: 10.1101/2021.08.18.456004. URL <https://www.biorxiv.org/content/10.1101/2021.08.18.456004v2>. Publisher: Cold Spring Harbor Laboratory ISBN: 10.1101/2021.08.1.
- [8] Basil el Jundi, James J J. Foster, Lana Khaldy, Marcus J J. Byrne, Marie Dacke, and Emily Baird. A Snapshot-Based Mechanism for Celestial Orientation. *Current Biology*, 26(11):1456–1462, 2016. ISSN 09609822. doi: 10.1016/j.cub.2016.03.030. URL <http://dx.doi.org/10.1016/j.cub.2016.03.030>.
- [9] Yvette E. Fisher. Flexible navigational computations in the Drosophila central complex. *Current Opinion in Neurobiology*, 73:102514, April 2022. ISSN 09594388. doi: 10.1016/j.conb.2021.12.001. URL <https://linkinghub.elsevier.com/retrieve/pii/S0959438821001537>.
- [10] Yvette E. Fisher, Jenny Lu, Isabel D’Alessandro, and Rachel I. Wilson. Sensorimotor experience remaps visual input to a heading-direction network. *Nature*, 576(7785):121–125, 2019. ISSN 14764687. doi: 10.1038/s41586-019-1772-4. URL <http://dx.doi.org/10.1038/s41586-019-1772-4>.
- [11] Yvette E Fisher, Michael Marquis, Rachel I Wilson, Chan Zuckerberg Biohub, and San Francisco. Dopamine promotes head direction plasticity during orienting movements. *Nature 2022*, pages 1–7, 11 2022. ISSN 1476-4687. doi: 10.1038/s41586-022-05485-4. URL <https://www.nature.com/articles/s41586-022-05485-4>.

- [12] Roman Goulard, Cornelia Buehlmann, Jeremy E Niven, Paul Graham, and Barbara Webb. A unified mechanism for innate and learned visual landmark guidance in the insect central complex. *PLoS computational biology*, 17(9):e1009383, 2021.
- [13] Jerrold H Zar. *Biostatistical analysis*. Pearson Education, 2010.
- [14] Brad K Hulse, Hannah Haberkern, Romain Franconville, Daniel B Turner-Evans, Shin-ya Takemura, Tanya Wolff, Marcella Noorman, Marisa Dreher, Chuntao Dan, Ruchi Parekh, et al. A connectome of the drosophila central complex reveals network motifs suitable for flexible navigation and context-dependent action selection. *Elife*, 10:e66039, 2021.
- [15] Lana Khaldy, Orit Peleg, Claudia Tocco, L. Mahadevan, Marcus Byrne, and Marie Dacke. The effect of step size on straight-line orientation. *Journal of The Royal Society Interface*, 16(157):20190181, aug 2019. ISSN 1742-5689. doi: 10.1098/rsif.2019.0181. URL <https://royalsocietypublishing.org/doi/10.1098/rsif.2019.0181>.
- [16] Sung Soo Kim, Hervé Rouault, Shaul Druckmann, and Vivek Jayaraman. Ring attractor dynamics in the Drosophila central brain. *Science*, 356(6340):849–853, 2017. ISSN 10959203. doi: 10.1126/science.aal4835.
- [17] Sung Soo Kim, Ann M. Hermundstad, Sandro Romani, L. F. Abbott, and Vivek Jayaraman. Generation of stable heading representations in diverse visual scenes. *Nature*, 576(7785):126–131, 2019. ISSN 14764687. doi: 10.1038/s41586-019-1767-1. URL <http://dx.doi.org/10.1038/s41586-019-1767-1>.
- [18] Jaison Jiro Omoto, Bao-Chau Minh Nguyen, Pratyush Kandimalla, Jennifer Kelly Lovick, Jeffrey Michael Donlea, and Volker Hartenstein. Neuronal Constituents and Putative Interactions Within the Drosophila Ellipsoid Body Neuropil. *Frontiers in Neural Circuits*, 12:103, November 2018. ISSN 1662-5110. doi: 10.3389/fncir.2018.00103. URL <https://www.frontiersin.org/article/10.3389/fncir.2018.00103/full>.
- [19] Ioannis Pisokas, Stanley Heinze, and Barbara Webb. The head direction circuit of two insect species. *eLife*, 9:1–49, 2020. ISSN 2050084X. doi: 10.7554/eLife.53985.
- [20] Shahrzad Shaverdian, Elin Dirlik, Robert Mitchell, Claudia Tocco, Barbara Webb, and Marie Dacke. Weighted cue integration for straight-line orientation. *iScience*, page 105207, 2022. ISSN 2589-0042. doi: <https://doi.org/10.1016/j.isci.2022.105207>. URL <https://www.sciencedirect.com/science/article/pii/S2589004222014791>.
- [21] Thomas Stone, Barbara Webb, Andrea Adden, Nicolai Ben Weddig, Anna Honkanen, Rachel Templin, William Weislo, Luca Scimeca, Eric Warrant, and Stanley Heinze. An Anatomically Constrained Model for Path Integration in the Bee Brain. *Current Biology*, 27(20):3069–3085.e11, 2017. ISSN 09609822. doi: 10.1016/j.cub.2017.08.052. URL <https://doi.org/10.1016/j.cub.2017.08.052>.
- [22] Daniel Turner-Evans, Stephanie Wegener, Hervé Rouault, Romain Franconville, Tanya Wolff, Johannes D Seelig, Shaul Druckmann, and Vivek Jayaraman. Angular velocity integration in a fly heading circuit. *eLife*, 6:e23496, May 2017. ISSN 2050-084X. doi: 10.7554/eLife.23496. URL <https://elifesciences.org/articles/23496>.
- [23] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold,

Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.

[24] D Wilkie. Miscellaneous rayleigh test for randomness of circular data, 1983.

Appendix E

BeetleBot start-up guide

E.1 Software availability and usability

The software is split across two GitHub repositories, with a third providing dependencies. These repositories are `beetlebot_software` (main repository), `bb_sensors`, and `beetlebot_dependencies`. The main repository (see ([BeetleBot Software](#))) contains links to the other two repositories.

The `beetlebot_software` repository contains all software which is designed to be built and run on the Host PC. The `beetlebot_dependencies` repository contains a snapshot of the source code for the compile and run-time dependencies (so as to ensure version compatibility).

All software contained in the `bb_sensors` repository is designed to be built and run on the robot. A full system image is provided for the robot which includes all dependencies. This image should work on any Raspberry Pi 3 but will not work on a Raspberry Pi 4.

E.2 Practical considerations for usage

Most of the software on which the robot codebase is based has passed, or close to, its end of life. The Raspberry Pi is running Raspbian 9 with ROS Kinetic. Support for Debian 9 (on which Raspbian is based) ended in 2022 and support for ROS Kinetic ended in 2021. The Host PC software (Ubuntu 20.04 and ROS Noetic) is supported until 2025.

The ideal solution would be to migrate everything to ROS2, however, this presents other challenges. ROS2 uses a different API for both C++ and Python meaning that all nodes would need to be updated. In addition, using ROS2 requires using Ubuntu 22.02 which strains the older Raspberry Pi 3B+. ROS releases are tightly coupled to Ubuntu releases. Breaking this coupling is not easy or advisable. Upgrading the Rasp-

berry Pi may, in turn, break the interaction with the polarisation or wind sensing hardware as these interface directly with the Raspberry Pi GPIO. In short, there is not a quick and easy way to bring everything up to date. It would not be impossible, I just wish to highlight that it would not be trivial. This may not be a problem but is worth bearing in mind.

E.3 Start-up guide

To run the simulation demo, you will need:

- A computer running Ubuntu 20.04 to act as the Host PC
- An internet connection

For the hardware demo at the end, you will need:

- The BeetleBot
- A fan
- A light source
- A testing space
- WD-40 or an equivalent product

This guide assumes your Host PC can generate a WiFi hotspot; it is also useful to have a wired connection to permit internet access when the hotspot is active. I do not know if an Ubuntu 20.04 virtual machine will work; in my experience, virtualisation tends to cause problems with ROS1 networking. If you do not have access to the physical BeetleBot, the code can still be installed and built so as to run a simulated demo (follow the instructions until 'Configure the WiFi'). The BeetleBot should contain an SD card with the required system image ready-to-go, but if you experience any problems, a system image is provided which can be flashed onto an empty MicroSD card ($\geq 8\text{GB}$)¹.

Install build tools On the Host PC, run:

```
$ sudo apt install build-essential
```

Install ROS Noetic On the Host PC, follow the instructions on the ROS Wiki ([Quigley et al., 2009](#)) for installing ROS Noetic. Make sure you set up your catkin workspace.

¹Do not flash the card which is in the robot, use a fresh card.

Set up Python

```
$ sudo apt install python-is-python3 python3-pip
$ pip install matplotlib==3.7
```

These commands will point the python command to Python 3, install the pip package manager and then install the required version of matplotlib.

Install Host PC beetlebot software

```
$ cd ~/catkin_ws/src
$ git clone https://github.com/refmitchell/beetlebot_software
$ git clone https://github.com/refmitchell/beetlebot_dependencies
$ cd ~/catkin_ws/src
$ catkin_make
```

This sequence of commands will clone both the beetlebot_software and beetlebot_dependencies repositories into your catkin workspace, then build the workspace.

Simulation demo

At the terminal, run:

```
$ roslaunch bb_launchers sim_ci_demo.launch
```

This should launch a TurtleBot simulation using RViz along with a window which visualises the neural populations in the cue integration model from ([Paper II](#)). In a **new terminal** run:

```
$ roslaunch turtlebot3_teleop turtlebot3_teleop_key.launch
```

Entering movement commands in this terminal will cause the simulated TurtleBot to move. If you apply some angular velocity by pressing the 'd' key, you will see the simulated TurtleBot rotate, simultaneously, the neural populations should update to reflect the rotation and the R → E-PG mappings will update in response to the rotation (Figure E.1).

Configure the WiFi The TurtleBot communicates with the Host PC via WiFi. A convenient way to set this up is to have the Host PC generate a WiFi hotspot to which the TurtleBot connects. The BeetleBot image is configured to automatically connect to a WiFi network with the SSID 'rob-fs' and password 'robotnetwork'. If you set up your Host PC to generate a WiFi hotspot with this SSID and password, the robot should automatically connect when turned on (this can easily be done in the Ubuntu WiFi settings).

Turn on the WiFi hotspot, then turn on the robot. You can check to see if the robot has connected by using:

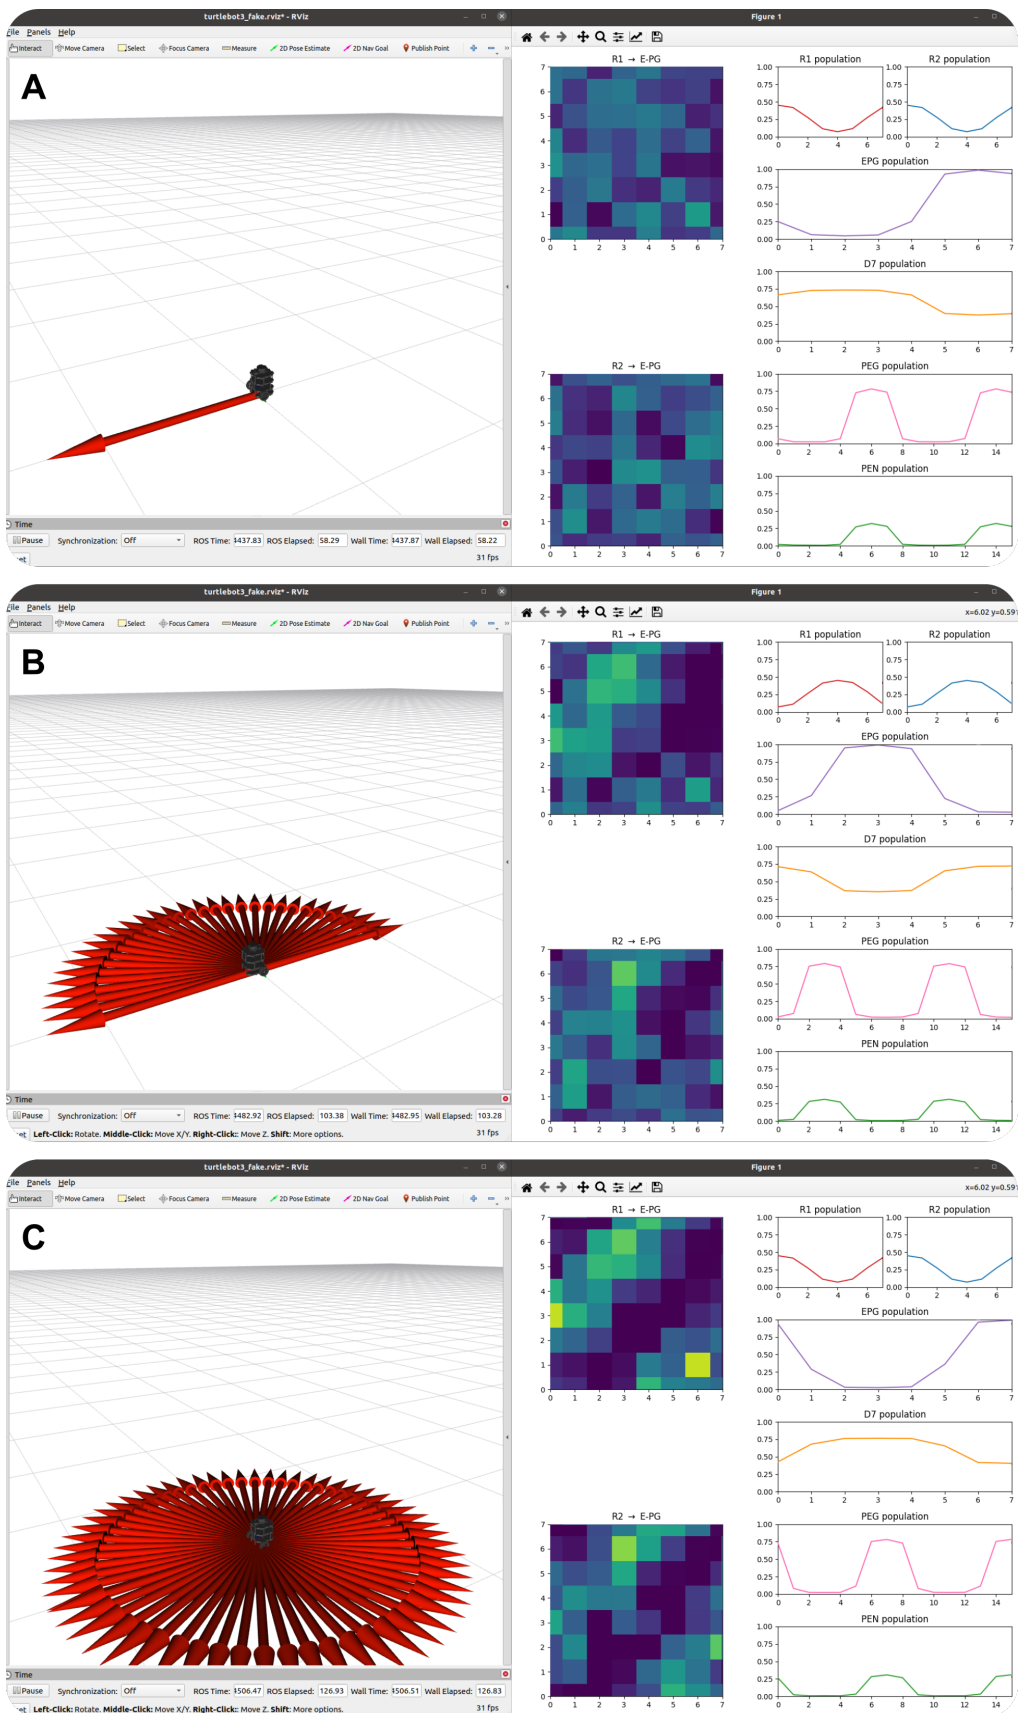


Figure E.1: Expected simulation output when running `ci_dummy.launch`. **A** (Left) The RViz window showing the Initial position of the simulated TurtleBot. (Right) The cue integration model visualisation window. **B** Network and simulated robot state after 180° using teleoperation. **C** Network and simulated robot state after a complete rotation.

```
$ sudo apt install net-tools
$ arp -an
? (10.42.0.229) at b8:27:eb:1e:7b:dd [ether] on wlo1
```

The last line here indicates that something has connected to the WiFi hotspot with the IP address 10.42.0.229; this should be the BeetleBot. Note that the specific IP address may be different. If nothing has connected, give it time; it can take a couple of minutes (but it usually helps if the hotspot is available before the robot is turned on). If the BeetleBot fails to connect to the hotspot, you will need to connect a screen, keyboard, and mouse to the Raspberry Pi and troubleshoot the WiFi problems.

Assuming the robot did connect, you should now be able to SSH into the robot using:

```
$ ssh pi@10.42.0.229
```

Substitute 10.42.0.229 for whatever IP was returned by `arp` on your system. Enter the password 'turtlebot' when prompted². You should now be logged into the Raspberry Pi:

```
pi@raspberrypi:~ $
```

Setting up the BeetleBot environment variables The BeetleBot system image should have everything you need to get going but you will need to check that the ROS environment variables are correct. On the Raspberry Pi, open the `~/.bashrc` file and look for the following environment variables:

```
export ROS_MASTER_URI=http://10.42.0.1:11311
export ROS_HOSTNAME=10.42.0.229
```

Update them such that³:

```
export ROS_MASTER_URI=http://<IP OF HOST PC>:11311
export ROS_HOSTNAME=<IP OF TURTLEBOT OBSERVED USING arp -an>
```

The IP of the Host PC can be found by using `ifconfig` (or similar) when the WiFi hotspot is running. Save and close the `~/.bashrc` file and run

```
pi@raspberrypi:~$ source ~/.bashrc
```

²If you are planning to work with the robot regularly, it is a good idea to set up an SSH key pair so you do not need the password every time.

³You can configure Ubuntu to assign a static IP to the BeetleBot whenever it connects to the hotspot but I never found this necessary.

Similarly, on the Host PC, open the `~/.bashrc` file and add the following lines:

```
export TURTLEBOT3_MODEL="burger"  
export ROS_MASTER_URI=http://<IP OF HOST PC>:11311  
export ROS_HOSTNAME=<IP OF THE HOST PC>
```

Save and close the file, then do:

```
$ source ~/.bashrc
```

Reboot both the Host PC and the robot at this point.

Running the robot The system image should contain all of the software pre-built, but changes may have been made to `bb_sensors` since this guide was written.

```
pi@raspberrypi:~$ cd ~/catkin_ws/src/bb_sensors  
pi@raspberrypi:~$ git pull origin main  
pi@raspberrypi:~$ cd ~/catkin_ws  
pi@raspberrypi:~$ catkin_make
```

This will fetch any changes made to the `bb_sensors` package made after the system image was made and re-build the catkin workspace on the robot.

On the **Host PC**, run:

```
$ roscore
```

In a **new terminal**, run:

```
$ roslaunch turtlebot3_teleop turtlebot3_teleop.launch
```

Then on the **Raspberry Pi** run:

```
pi@raspberrypi:~$ roslaunch bb_sensors beetlebot.launch
```

This will start the nodes which read from the camera and wind sensor along with the TurtleBot3 core⁴.

If you switch to the terminal with the teleoperation input, you should be able to drive the robot by following the instructions on screen. This confirms that the TurtleBot and Host PC can communicate. You can now kill the ROS process on the turtlebot (by using `ctrl+c`) and then shut it down using:

```
pi@raspberrypi:~$ sudo shutdown now
```

Kill any ROS nodes/instances running on the Host PC before proceeding.

⁴Specifically, `turtlebot3_bringup turtlebot3_robot.launch`; see ([TurtleBot Manual](#))

BeetleBot demo Finally, we will replicate the simulated demo, but with real cues and a real robot. **Make sure the robot is off and that no ROS instances are running.**

Spray a small amount of WD-40 into the anemometer bearing. Spin the anemometer gently with your hand or by blowing on it to make sure it is spinning freely⁵. Set up the robot in front of the fan and set up the light source over the robot⁶. Turn on the fan and check that the anemometer and vane are moving freely in response to the air current.

Turn on the WiFi hotspot on the Host PC, then turn on the robot. On the **Host PC**, run:

```
$ roscore
```

and in a new terminal run:

```
$ roslaunch bb_launchers robot_ci_demo.launch
```

Open another terminal, SSH into the BeetleBot and then run:

```
pi@raspberrypi:~$ roslaunch bb_sensors beetlebot.launch
```

On the Host PC, two windows should appear. One shows the camera view with the location of the brightest spot in the visual field highlighted, the other shows the cue integration model state (as in the simulated version above).

If you open another terminal and start the keyboard teleoperation system, then rotate the robot as before, you should see the network state updating with respect to the real cues in the world.

Further information For more information on the TurtleBot platform, see ([TurtleBot Manual](#)). If you want to write code for the robot then you will need an understanding of ROS1 ([Quigley et al., 2009](#)), and in addition, you will need to check the existing documentation available in the software repository ([BeetleBot Software](#)).

I am more than happy to be contacted regarding the robot (hardware or software). I can be reached by email using r.mitchell@ed.ac.uk or refmitchell@outlook.com.

⁵If it is not spinning freely, then repeat with a little more WD-40 but use this sparingly

⁶Make sure you remove the lens cap from the robot camera.






Appendix F

Additional paper: Celestial compass design mimics the fan-like polarisation filter array of insect eyes

This paper presents an evaluation of a polarisation sensor prototype based on a simulated sensor. I constructed and field-tested the sensor which involved two 3-4 week field trips and some rudimentary analysis. For field testing, the sensor was integrated into the BeetleBot and is therefore of relevance to Chapter 6.

This work is included as an appendix as it does not directly contribute to the cue integration discussion in the thesis. The article is available under CC BY 4.0 and is included with the permission of the lead author, Evripidis Gkanias.

Celestial compass sensor mimics the insect eye for navigation under cloudy and occluded skies

Evrpidis Gkaniyas ^{1✉}, Robert Mitchell ¹, Jan Stankiewicz¹, Sadeque Reza Khan ², Srinjoy Mitra ³ & Barbara Webb ¹

Insects use the sun's position (even when concealed) as a compass for navigation by filtering celestial light intensity and polarisation through their compound eyes. To replicate this functionality, we present a sensor that imitates essential aspects of insect eyes, particularly the fan-like arrangement of polarised light receptors in their dorsal rim area. Our sensor comprises a ring of eight pairs of photodiodes (evaluating two orthogonal orientations of polarised light) to analyse the skylight coming from different directions. Because the layout of our sensor aligns with the polarised light pattern in the sky, a circular-mean model that integrates information spatially across the analysers can estimate the solar azimuth. When using the same sensor design, our model achieves lower compass errors than alternative (and computationally more complex) algorithms, especially under cloudy and occluded skies. Thus, the morphology and processing of the insect celestial compass provide an efficient and robust directional input for navigation.

¹School of Informatics, University of Edinburgh, Edinburgh, United Kingdom. ²School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh, United Kingdom. ³School of Engineering, University of Edinburgh, Edinburgh, United Kingdom. ✉email: ev.gkaniyas@gmail.com

A compass is used in navigation to provide robust orientation estimates. As travel distance increases, idiothetic orientation information—from proprioception or an inertial measurement unit (IMU)—will eventually succumb to noise¹. Typical allothetic compass solutions such as magnetometers often suffer from electronic interference, and the alternative of relying on a global positioning system (GPS) creates dependence on a full infrastructure of satellites which may not always be available. Ideally, a compass should be independent of external support, impervious to disturbance, lightweight, cheap, and energy efficient. The celestial compass (inspired by insects^{2–4}) fits all these criteria.

A celestial compass uses the properties of light coming from the sky to estimate the position of the primary celestial body (herein, the sun) or, with appropriate time compensation, the North. The light from large celestial objects (e.g., the sun or moon⁵) creates a regular pattern of intensity and polarisation across the sky (Fig. 1a). Under clear skies, skylight intensity peaks at the visual position of the sun (solar azimuth and elevation), while the degree of polarisation (DoP) is strongest in the opposite direction (at 90° from the sun and across the zenith). At a specific point, the angle of polarisation (AoP) is always perpendicular to the arc between the sun and that point, forming concentric rings of polarisation around the visual position of the sun^{6–8}. These stereotyped skylight patterns allow for various methods to recover the sun's position.

A range of celestial compass sensors has been developed that integrate AoP estimations from different positions in the sky to locate the sun (using photodiodes^{9–13}, or cameras^{14–21}). Theoretically, one can locate the sun by measuring the AoP in two points of the concentric-rings pattern. These sensors focus on the accurate extraction of the AoP and ignore the remaining (useful) properties of light. Estimating the AoP in silico is straightforward²² (using at least three polarisation filters and photodiodes), but the complication arises when attempting to compute the solar azimuth, which requires computing eigenvectors^{9–21}, a discrete Fourier transform (DFT)^{23–25}, or finding the straightest line of AoPs passing through the zenith^{26,27}.

The eyes and brains of many insects evolved to provide an alternative solution^{28,29}. Each of their compound eyes has near 180° panoramic vision (desert ants³⁰, fruit flies³¹) and is formed by hundreds of facets called ommatidia (Fig. 1b). The ommatidia that occupy the dorsal rim area (DRA) of the eyes are sensitive to polarisation. Each dorsal rim ommatidium analyses the incident light in two polarisation axes: one roughly perpendicular to its meridian from the dorsal-most position on the eye (orange in Fig. 1b–d; I_{90}), and one parallel to this meridian (blue in Fig. 1c, d; I_0). The fan-like arrangement of DRA ommatidia³ (see Fig. 1b, e) forms what Wehner called a matched filter³² that imitates the sky AoP spatial pattern and partially solves the navigational problem by transferring the complexity from algebra to geometry.

Lambrinos et al.^{33,34} first developed a sensor that could perceive and process light in a similar way to the DRA ommatidia (Fig. 1f). They replicated an ommatidium by using two photodiodes behind orthogonal polarisation filters (similar to Fig. 1d) and integrating their signals to the response of polarisation-opponent (POL-OP) units (Fig. 1g). The response of three units (pointing at the sky zenith but oriented at 120° intervals) was used to estimate the AoP at the zenith, predicting two possible solar azimuths that are 180° apart. They resolved this ambiguity by comparing the intensity of light in these two directions from surrounding polarisation-insensitive photodiodes. A similar bio-inspired approach has since been followed by others^{22,24,25,35–40}. However, Smith⁴¹ observed that rotating a (simulated) POL-OP unit at an angle from the zenith breaks the 180° ambiguity and its

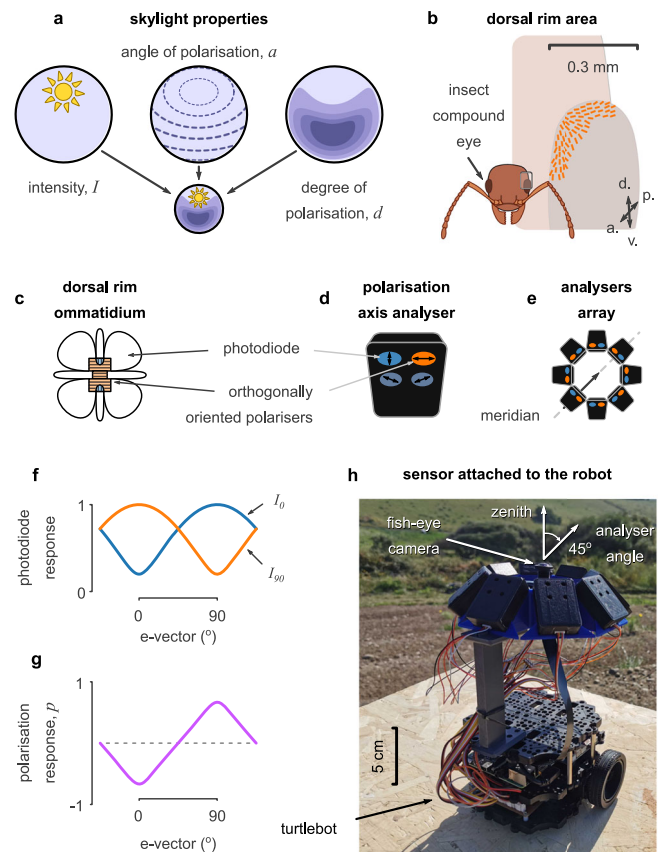


Fig. 1 Overview. **a** Skylight varies in light intensity (I), angle (α) and degree (d) of polarisation. **b** The compound eyes of insects filter light using specialised ommatidia in the dorsal rim area (DRA). These form a fan-like distribution, pointing in different directions in the sky, and covering up to 120° field of view. d.: dorsal, v.: ventral, p.: posterior, a.: anterior. **c** In each ommatidium, the skylight is filtered by orthogonally oriented microvilli before it is captured by two groups of photoreceptor cells. **d** Our polarisation axis analyser (PAA) mimics this function by filtering the skylight with two orthogonal polarisers before being captured by the photodiodes. **e** Our sensor has eight PAAs in a ring arrangement, elevated by 45°. **f** The responses of the orthogonal photodiodes (I_0 and I_{90}) over the e-vector of polarised light. **g** The normalised difference between the orthogonal responses (p) matches the response of the polarisation-sensitive neuron in the insect optic lobes. **h** Our robot uses a panoramic camera and our designed sensor to collect skylight data. **c**, **f**, and **g** were adapted and modified from²⁹.

response is strongest in the opposite direction of the solar azimuth (corresponding with the DoP pattern in the sky). Therefore, they suggested that the fan-like arrangement of the dorsal rim ommatidia could have evolved to take parallel measurements of the sky polarisation at different positions, obtaining the information from a single reading (without scanning) and inferring an instantaneous and unambiguous estimate of the sun's direction. Gkaniats et al.²⁹ independently arrived at the same idea and provided proof of principle for such a celestial compass using a simulated sensor array that mimicked the layout of the insect DRA.

Here, we extend this line of work and provide a hardware prototype of the celestial compass sensor of insects. This sensor may be used to determine relative solar azimuth using either polarised light or intensity information. We validated the performance of our compass model by recording sensor data under a wide range of atmospheric, weather, and sky occlusion

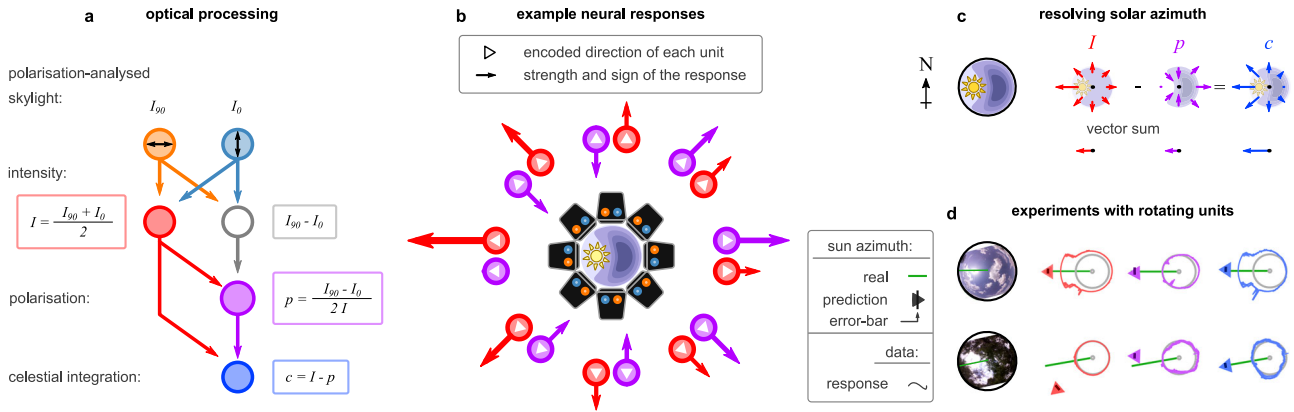


Fig. 2 Description of the computational model. **a** Adding the pulses from two photodiodes (PDs) with orthogonal polarised light filters can approximate the light intensity (I); subtracting them approximates the polarisation Fresnel ratio (PFR). Normalising the PFR with light intensity provides a good approximation of how well the polarity of light is aligned with the orientation of the polarisation axis analyser (PAA, p). The difference between I and p results in their celestial integration (c), which indicates how well the PAA is aligned with the brightest celestial body in the sky. **b** An example of the intensity and polarisation measurements of all the PAAs in a clear sky condition. Each measurement can be multiplied with a unit vector directed towards the azimuth of the PAA (see red and purple arrows), which can also flip by negative measurements (e.g., for p). **c** The simple addition of all the vectors of each unit type results in a vector that has the direction of the solar azimuth (anti-solar for p). **d** The measured responses of the different types of units (solid lines) in two different sky conditions (clear and partially cloudy with occlusions from trees) and the estimations of our model using only I (red), p (purple), or c (blue). Small triangles denote the median estimation across a set of 360 homogeneously distributed facing directions (one per degree) of the sensor; circular lines denote the first and third quartiles of the distribution; the green line highlights the solar azimuth.

conditions. We also compared its performance with several alternative models that could use similar input. Our results confirm that the proposed sensor provides a robust compass system that could be used by autonomous vehicles, and also provides insight into the neural processing stages underlying the celestial compass of insects.

Results

The celestial compass sensor. The sensor consisted of an array of eight polarisation axis analysers (PAAs) distributed evenly around a ring, then tilted to give an inclination of 45° (Fig. 1e, h). This arrangement (loosely) mimics the fan-like arrangement of the dorsal rim ommatidia across both insect eyes. Each PAA measured the voltage from four ultraviolet (UV) sensitive photodiodes (Fig. 1d), which were covered by linear polarisers in a different orientation (0° , 45° , 90° , and 135°). Readouts from all four photodiodes were recorded, but to maintain equivalence to the insect eye, our model used only two—those sensitive to light polarised at 0° and 90° , which correspond to $I_0 \in [0, 1]$ and $I_{90} \in [0, 1]$ respectively (Fig. 2a).

Following Stürzl¹⁶, the overall light intensity is

$$I = (I_{90} + I_0)/2. \tag{1}$$

Polarisation information is calculated as the intensity-normalised POL-OP²⁹, which ensures that $p \in [-1, 1]$,

$$p = (I_{90} - I_0)/(2I). \tag{2}$$

This value will be maximal when the PAA is perpendicularly aligned with the AoP of the light source (negative when parallelly aligned) and is proportional to the DoP (see Fig. 1g). Subtracting p from I gives what we call the celestial integration,

$$c = I - p. \tag{3}$$

Under clear sky conditions, I is higher when the PAA points towards the sun (solar azimuth) and is almost uniform across the remaining directions (Fig. 2b). In contrast, p is highest at the anti-solar azimuth (where the DoP is highest and the AoP is aligned with the filter of I_{90}), and negative at the sides and closer to the solar azimuth. A local maximum in p occurs at the solar azimuth

but is usually lower than the one in the anti-solar azimuth. Thus, the celestial integration (c) is higher towards and around the solar azimuth and gradually decreases towards the anti-solar azimuth (Fig. 2c). Although c looks similar to I under clear sky conditions, it becomes more informative than I about the solar azimuth under cloudy or occluded skies (Fig. 2d).

To generate a compass signal from the array, the values from each PAA were integrated according to the model suggested by Gkaniats et al.²⁹. Values (c) were transformed into vectors with their polar angles equal to the azimuth of their recording PAAs and lengths equal to the values themselves. The vectors were then averaged, yielding a mean vector that points towards the solar azimuth. If the celestial integration of the k^{th} PAA is c_k , the mean vector can be calculated as a complex number,

$$z_c = \frac{1}{K} \sum_{k=1}^K c_k e^{i2\pi(k-1)/K}, \tag{4}$$

where $2\pi(k-1)/K$ represents the azimuth of the k^{th} PAA for K PAAs evenly distributed around the ring. The response coming from each PAA is interpreted as its estimation of whether the sun is in its facing direction. Negative responses would indicate that the sun is estimated to be in the opposite direction (see Fig. 2b, c, d). Note that the mean vector also provides a measure of dispersion in its magnitude^{42,43}. The mean solar azimuth indicated by each PAA (α_c) and the angular deviation of these values (σ_c) can therefore be recovered as

$$\alpha_c = -i \ln \frac{z_c}{||z_c||}, \quad \sigma_c = \sqrt{2(1 - ||z_c||)}, \tag{5}$$

where $||z_c||$ is the length of the mean vector. By replacing c_k with I_k or $-p_k$ in (4), we may also estimate the solar azimuth based only on intensity (z_I) or polarisation (z_p) respectively. Depending on the values used we could have an intensity, polarisation, or celestial compass.

Figure 2d shows examples of I , p , and c values, along with their respective estimates of the solar azimuth for two different sky conditions. Under an almost clear sky, all provided similar estimates. However, with canopy occlusion, the intensity and polarisation estimates show a consistent deviation towards

opposite sides of the foliage opening, which was corrected by the celestial integration.

A dataset of varying skylight conditions. Our sensor was mounted on a robot to collect data from remote sites (Fig. 1h and Materials and Methods). Data were collected from field sites in Italy and South Africa (in May and November 2022 respectively), and our dataset comprises a variety of solar elevations, weather conditions, and occlusions. The collection was organised into sessions, each of which consisted of twelve complete (360°) rotations under a given condition. At the beginning of each recording session, the robot was initialised facing magnetic North ($\pm 5^\circ$) and IMU was reset to zero. On each rotation, the robot would log all photodiode responses and IMU data (to measure the deviation from the starting direction; Supplementary Fig. 1a). We computed the angle of the facing direction of the robot with respect to the solar azimuth by taking the difference between the IMU and the theoretical solar azimuth (for the particular location, time and day of the recording). This raw dataset was then transformed into a pooled dataset which additionally included the responses of each optical unit (*I*, *p*, and *c*; Supplementary Fig. 1b).

Using this pooled dataset we were able to reconstruct the performance of the same basic design using different numbers of PAAs distributed evenly around the sensor ring (Supplementary Fig. 1c). This was achieved by determining the preferred direction of the reconstructed PAAs, then taking the median of the five responses from the dataset which were most closely clustered around that preferred direction (Supplementary Fig. 1d).

Spatial sampling. Using the interpolation process described in Supplementary Fig. 1, we compared the performance of our sensor in predicting the solar azimuth, for varying numbers of (hypothetical) PAAs. We refer to solar azimuth prediction error as the global reference error as it represents how well we could predict the actual solar azimuth from one compass sensor reading (from the given number of PAAs). However, this measure is subject to the $\pm 5^\circ$ error incurred when initialising the robot (note that attempting to use the fish-eye images to determine the actual solar azimuth was subject to even more error, due to lens distortion, lens flare and occlusions). We therefore also evaluate the local reference error which is how well the sensor can estimate the robot's angular displacement, for any point in its rotation, relative to its starting direction, where the IMU measurement of displacement (which has negligible error over the short time scale) is taken as ground truth.

To assess the effect of PAA number, we selected the easiest scenarios from our dataset (little to no cloud, solar elevation at least 10°); the results are shown in Fig. 3a, b and Supplementary Table 1. With the minimum of three PAAs, the root mean square error (RMSE) for the sun direction (global reference) was relatively high (10.53°). The lowest error was achieved with our maximum of sixty PAAs (RMSE 2.65°), i.e., within the range of the initialisation error. Adding PAAs always improved the performance. The local reference error was high for fewer than six PAAs but dropped substantially for six or more (3.78° , see Supplementary Table 1). Beyond thirty-five PAAs, performance improvements were negligible (RMSE 0.53° for 36 PAAs, compared to 0.43° for 60 PAAs). In all subsequent results, we use only eight PAAs as in our hardware implementation.

Solar elevations. Under (relatively) clear sky conditions, we examined the effect of solar elevation on sensor performance. Data were sorted by their solar elevation and binned in 5° intervals from -5° (below the horizon) to 85° (near zenith). Figure 4a shows examples of the sensor performance for specific

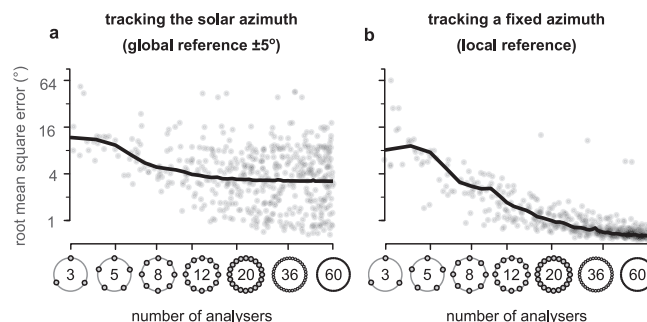


Fig. 3 Sensor performance over variable spatial sampling. **a** Root mean squared error (RMSE) of estimating the solar azimuth ($\pm 5^\circ$ error). **b** RMSE of tracking any fixed direction. Vertical axes are in the \log_2 scale. Horizontal axes are in \log_6 scale, and random noise (0.5) was added to the (whole) number of polarisation axis analysers, making them appear as a continuous distribution. Ticks without labels denote (from bottom to top) 0° , 2° , 8° , 32° , and 90° RMSE. Although the spread appears to be increasing with higher spatial resolution, this is an illusion of the logarithmic scale of the vertical axis. Standard deviation (SD) in **a** is 21.77° for 3 analysers and it falls in the range from 4.41° to 5.15° from 4 to 60 analysers (randomly distributed). The SD in **b** is 21.92° for 3 analysers and it falls in the range from 2.82° (4 analysers) to 0.48° (60 analysers; decreasing for higher spatial resolutions) for the rest of the cases.

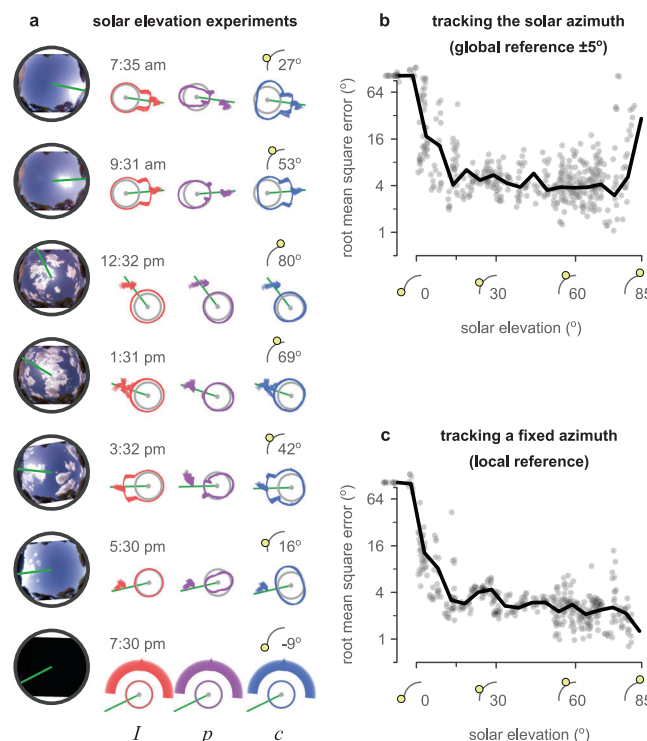


Fig. 4 Sensor performance over different solar elevations. **a** Examples of different solar elevations in the sky. Left: fish-eye images of the sky for each example. Right: *I*, *p*, and *c* values in each example, the prediction of the model using these responses (intensity, polarisation, and celestial compass, respectively). Above the responses are the time of day and the solar elevation during the session. Each example shows the median (arrowhead) and quantiles (circle segment) of the 360 predictions produced (at homogeneously distributed orientations) during each of the twelve full rotations of the sensor. **b** Root mean squared error (RMSE) of estimating the solar azimuth ($\pm 5^\circ$ error). **c** RMSE of tracking any fixed direction. Vertical axes are in the \log_2 scale. Ticks without labels denote (from bottom to top) 0° , 2° , 8° , 32° , and 90° RMSE.

solar elevations. Figures 4b and c show respectively the global and local reference errors for all the sessions, which are also summarised in Supplementary Table 2. There was no noticeable difference in the performance of our sensor for solar elevations above 10° (average RMSE, global: 5.89°, local: 2.77°). Lower elevations caused the error to increase rapidly; at and below 0° elevations the sensor was unable to provide useful predictions. There was also a sharp increase in the global error at high elevations that was not reflected in the local error; we believe this might reflect the increasing inaccuracy of our (hand-selected) estimate of the true sun direction in the sky images as the sun nears the zenith, rather than inaccuracy in the compass per se. Overall, the sensor performance was stable throughout the day when the sun was higher than 10° from the horizon.

Atmospheric conditions. As well as different sky conditions, our data were collected in three locations that represent different atmospheric conditions: Sardinia (Italy), Vryburg (South Africa), and Bela Bela (South Africa)—see Materials and Methods for more details. Sardinia is an island and the data were collected close to the shore, where the humidity was generally high. Vryburg is located in the African Savannah, where the climate is dry, and Bela Bela is in the woodlands, which provides an average climate. Supplementary Table 3 shows that the performance of our compass sensor and model was not affected by the climates of these different locations. Interestingly, the panoramic images of the sky that we collected from Sardinia look surprisingly similar to the ones reported in oceanic atmospheric conditions⁴⁴. Thus, they could approximate off-shore atmospheric conditions and suggest that our sensor could be used in intercontinental missions.

Cloudy skies. The examples in Fig. 4a suggested that (under clear skies) the predictions of solar azimuth made by our compass model were unaffected by the value used (*I*, *p*, or *c*). Under a clear sky, the relationship between solar azimuth, light intensity, and polarisation remains stable. However, the presence of clouds disrupts this relationship. The thickness of clouds determines how sunlight is scattered, and affects both the polarisation and intensity distribution in the sky. Our dataset includes recording sessions under a variety of sky conditions which we classified based on their cloud cover (for examples see Fig. 5a). Global and local error under varying cloud conditions can be seen in Fig. 5b, c and Supplementary Table 4. Sensor function was not substantially disrupted by thin or broken cloud cover but deteriorated for thicker cloud cover (uniform and solid).

Occluded skies. Some insects which are known to use polarisation for orientation live in densely wooded areas, so an insect-inspired celestial compass should function under canopy cover. The same principle extends to robots working in built-up areas; a celestial compass should be robust to local occlusion due to buildings. We tested our sensor with different degrees of canopy occlusion (examples can be seen in Fig. 6a). Results are shown in Fig. 6b, c and Supplementary Table 5. In general, the intensity and polarisation compasses showed a strong dependence on whether the solar or anti-solar side of the sky was blocked respectively. The celestial integration could effectively use information from whichever was most useful at the time. Where trees occupied the central area or full panorama (with relatively open foliage), polarisation tended to be more robust than the intensity or celestial compasses. Additional PAAs generally reduced both global and local RMSE (especially for the intensity compass, and for scenarios close to trees and with dense cover; see Supplementary Table 5).

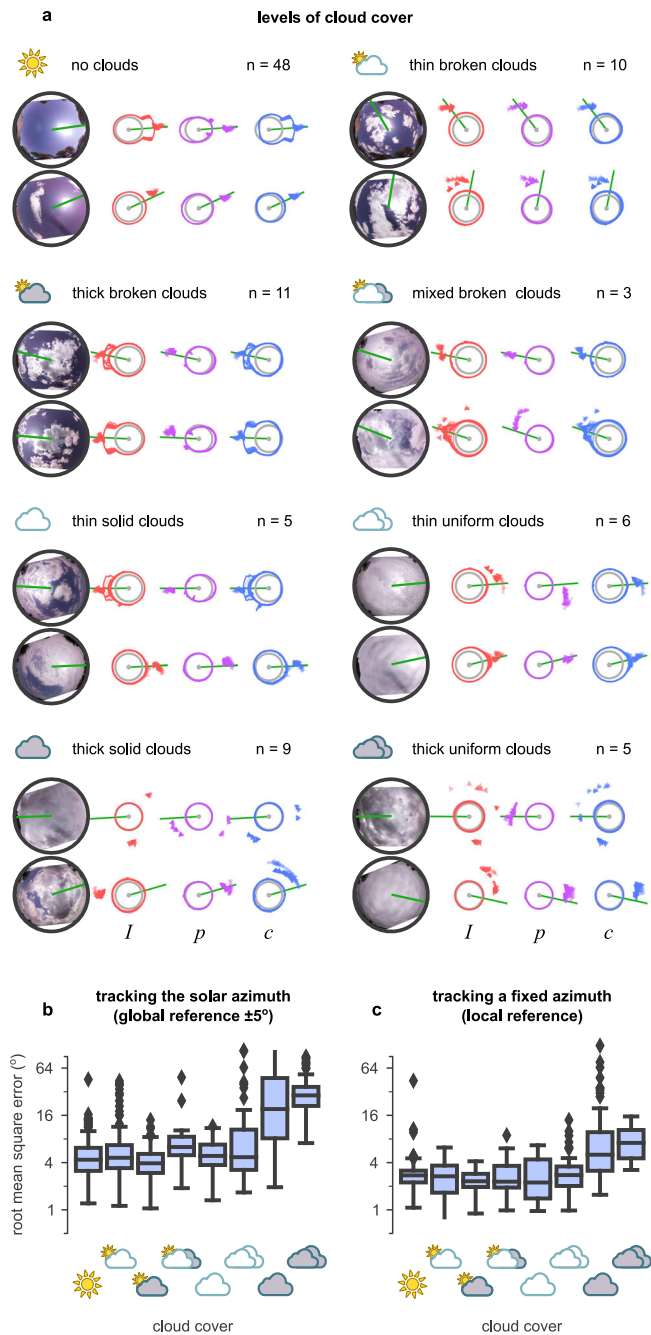


Fig. 5 Sensor performance over cloudy skies. **a** Examples of different levels of cloud cover in the dataset, in increasing difficulty: no clouds, thin broken clouds, thick broken clouds, mixed broken clouds, thin solid clouds, thin uniform clouds, thick solid clouds, and thick uniform clouds. Each block represents a different condition with two examples and reports the total number of available examples in the dataset (*n*). Left: fish-eye images of the sky for each example. Right: *I*, *p*, and *c* values for each example, the prediction of the solar azimuth using the different models (intensity, polarisation, and celestial compass). Each example shows the median (arrowhead) and quantiles (circle segment) of the 360 predictions produced (at homogeneously distributed orientations) during each of the twelve full rotations of the sensor. **b** Root mean squared error (RMSE) of estimating the solar azimuth (±5° error) using the celestial compass. **c** RMSE of tracking any fixed direction. Box-plot: centre line, median; box limits: upper and lower quartiles; whiskers: 1.5 × interquartile range; points: outliers. Vertical axes are in the log₂ scale. Ticks without labels denote (from bottom to top) 0°, 2°, 8°, 32°, and 90° RMSE.

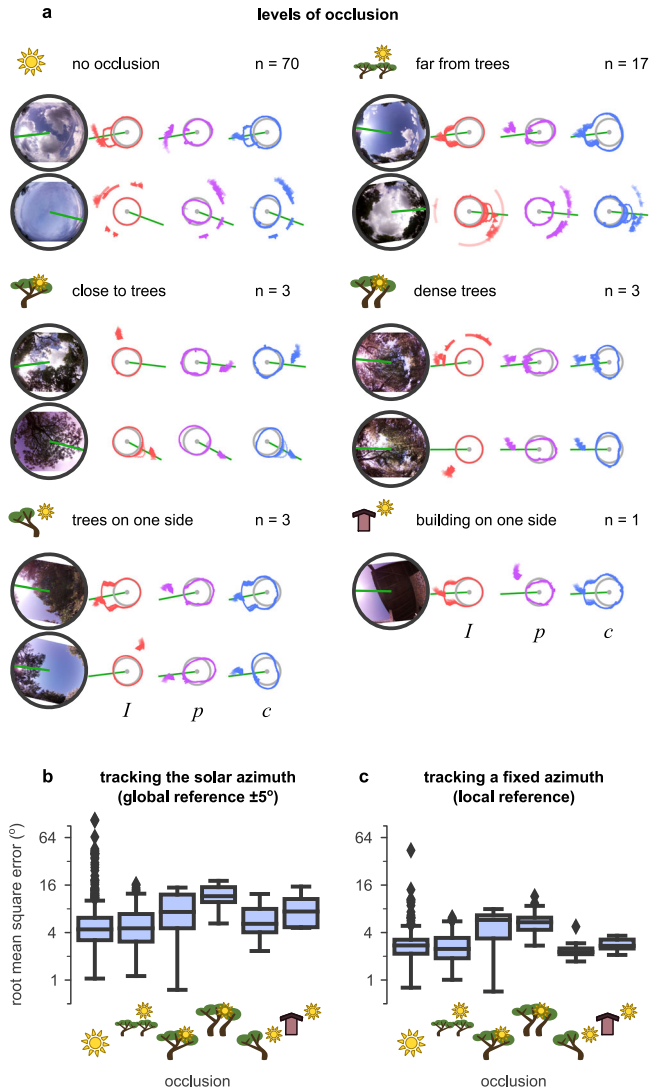


Fig. 6 Sensor performance over occluded skies. **a** Examples of different levels of occlusion in the dataset: no occlusion, occlusion based on sensor distance from trees (far or close), dense tree cover with openings, and trees or buildings on one side. Each block represents a different condition with up to two examples and reports the total number of available examples in the dataset (n). Left: fish-eye images of the sky for each example. Right: I , p , and c values in each example, the prediction of the solar azimuth using the different models (intensity, polarisation, and celestial compass). Each example shows the median (arrowhead) and quantiles (circle segment) of the 360 predictions produced (at homogeneously distributed orientations) during each of the twelve full rotations of the sensor. **b** Root mean squared error (RMSE) of estimating solar azimuth ($\pm 5^\circ$ error). **c** RMSE of tracking any fixed direction. Box-plot: centre line, median; box limits: upper and lower quartiles; whiskers: $1.5 \times$ interquartile range; points: outliers. Vertical axes are in the \log_2 scale.

Alternative compass models. We have noted several instances above in which the celestial compass ($c = I - p$) provides a better estimate than a compass using only intensity I or polarisation p . Here we provide a more direct comparison between these three alternatives and two further models that use different methods to estimate the solar azimuth from polarisation measurements. Considering first the use of c vs. I or p (Fig. 7), in nearly every condition, the celestial integration provided better estimates of the solar azimuth than the other two compasses. An exception was when the sun was completely hidden, in which case the

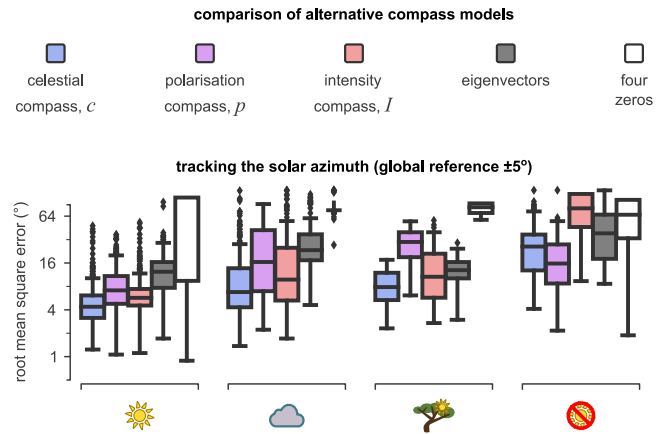


Fig. 7 Performance of alternative models. Boxes represent the distribution of root mean square error (RMSE) across the subset of the data where the sky was almost clear (sun), with thick clouds (cloud), with severe occlusions (tree), or where the sun was completely covered by clouds or canopies. Blue represents the celestial compass, purple represents the polarisation compass, red represents the intensity compass, dark grey represents the eigenvectors model, and white represents the four zeros model. We report the root mean square error (RMSE) of predicting the solar azimuth. The data shown are for solar elevations of at least 15° . Box-plot: centre line, median; box limits: upper and lower quartiles; whiskers: $1.5 \times$ interquartile range; points: outliers.

polarisation compass was more accurate (Supplementary Fig. 2a, b and Supplementary Table 6). In the remaining cases, the intensity compass and celestial integration had a better performance than the polarisation compass alone. Increasing the number of PAAs (from eight to thirty-six) had almost no effect on the estimation of the intensity compass, but positively affected the estimations of the polarisation and celestial compasses in occluded skies. When tracking a fixed azimuth (local RMSE), the intensity and celestial compasses achieved the most accurate estimation of orientation (Supplementary Fig. 2c and Supplementary Table 6).

We also implemented two alternative models that can use the signals from our PAAs to extract the solar azimuth. In the first model, we calculated a unit vector for each PAA, which was directed perpendicularly to the AoP (as computed by the complete set of four photodiodes per PAA, at 0° , 45° , 90° , and 135° , and following the computations suggested by Zhao et al.⁴⁰). However, the AoP ranges in $[-90^\circ, +90^\circ]$, which creates a 180° ambiguity of the direction. To resolve this, we assumed that all the unit vectors point towards the interior of the compass sensor (see Supplementary Fig. 3a). Then we followed the approach of Stürzl and Carey¹⁵ and used the covariance of these vectors to calculate the eigenvector with the lowest eigenvalue, which should point towards the average direction of all these vectors, and we refer to this as the eigenvectors model. Numerous other celestial compasses used variations of this method^{9–21}. Smith and Stewart proposed the second model^{41,45,46} and suggested that, when rotating a tilted PAA, the polarisation responses (p) as calculated by equation (2) form a curve that takes both positive and negative values along the rotating axis and that this curve is zero at exactly four directions (see Supplementary Fig. 3b). The two directions that are the closest to each other should also be the closest to the solar azimuth. Thus the solar azimuth can be calculated as the mean direction of these two directions (described by the four zeros), and we refer to this as the four-zeros model. A detailed description of how we implemented these models can be found in Materials and Methods.

Figure 7 suggests that our proposed model outperformed the other two models described above in all the tested sky conditions. In the easiest scenario, where the compass was composed of 36 PAAs and used a local reference under almost clear sky conditions, our celestial compass model achieved (on average) 0.59° RMSE, while the eigenvector and four-zeros models achieved 2.70° and 70.38° respectively (Supplementary Fig. 2c and Supplementary Table 6). The four-zeros model was very fragile even for clear sky conditions, as any solar elevations higher than the PAA (>45°) produced only two zeros instead of four (Supplementary Fig. 3a—first example) and elevations below 15° produced evenly spaced zeros (Supplementary Fig. 3a,b—second example) resulting in a random direction choice. In most conditions, the eigenvector model performance was better than the four-zeros model but not as good as the celestial compass model. The performance of the eigenvector model was affected less by occlusions, but substantially affected by clouds, and more substantially when the sun was hidden completely. The main reason for this performance drop was the 180° ambiguity of the unit vectors (which were based on the AoP) and the assumption that they always face towards the interior of the compass. Impressive results reported previously for this method rely (instead of this assumption) on a calibration process to minimise the RMSE^{11–21}. However, such a process does not seem biologically plausible and can also be expensive from a technical perspective.

Discussion

Inspired by insect vision, we physically implemented the sensor design and computational model proposed by Gkaniyas et al.²⁹, verified that it can be used as a robust celestial compass sensor, and that even better estimations can be obtained by a simple integration of the skylight polarisation and intensity signals. The performance of the model was tested in various sky conditions, with different types of clouds and occlusions. Compared to the performance of two other compass models, it demonstrated superiority in its estimations of both the (global) solar azimuth and of a (local) point of reference.

Our computational model was also developed as an in-principle way to process the information from an array of PAAs. However, it is interesting to compare the algorithm to known processing pathways in the insect brain (illustrated in Fig. 8 for the *Drosophila melanogaster* fruit fly). The orthogonally placed polarisers of our PAAs imitate the microvilli of the photoreceptors in the dorsal rim ommatidia. Respectively, I_{90} and I_0 values are analogous to the responses of R7 and R8 photoreceptor cells in the insect retina. Interestingly, R7 axons are inhibited by R8s, implementing the $I_{90} - I_0$ node in Fig. 2a^{47,48}. Intensity information (I) is probably encoded by the Dm9 neurons, which collect R7 and R8 responses and feedback to the R7 axons. Polarisation information (p values in our model) could then be encoded by the DmDRA1 interneurons and is projected to the anterior optic tubercle (AOTu) through a subset of medulla-tubercule (MeTu) neurons. This is because the R7-R8 difference encoded by the R7 axon is normalised by the Dm9 activity before connecting to the DmDRA1, which results in a response that is increased by R7 (I_{90}) and decreased by R8 activity (I_0)^{47,49,50} or unpolarised light (I , feedback from Dm9)^{48,51}. We predict that the final intensity information (I value) is encoded by a different subset of MeTu neurons that receive direct input from Dm8 neurons (downstream of non-DRA R7 neurons that respond to unpolarised light)⁴⁸. The responses of all these neurons are collected by the AOTu and communicated by the tubercular-bulb (TuBu) neurons to central brain regions. Note that the representation in the AOTu keeps the fan-like retinotopic structure of




Fig. 8 Pathway of the polarised light in the insect brain (here *Drosophila melanogaster*) and the suggested parallels to the processing in our model. R7/8: retina neuron 7/8, Me: medulla, Lo: lobula, DmDRA1: distal medulla dorsal rim area neuron 1, Dm8/9: distal medulla neuron 8/9, MeTu: medulla-tubercule, AOTu: anterior optic tubercle, TuTu: tubercular-tubercle, TuBu: tubercular-bulb, Bu: bulb, Bu_s: superior bulb, Bu_a: anterior bulb, ER: ellipsoid-body from ring neurons, EB: ellipsoid body, CX: central complex. Adapted and modified from⁵¹; data from^{48,51}.

the DRA, but pooled and summarised in 10 columns per hemisphere⁵¹. Also, it establishes an inter-hemisphere communication^{48,51–55}, which we predict connects the two AOTus in a complete ring that represents the solar azimuth. Thus, our celestial integration values (c) should be analogous to the responses of the TuBu neurons and collectively represent a vector pointing towards the solar azimuth. However, these responses should also be affected by spectral and optic flow inputs⁵¹, which are omitted from our model. The celestial compass of insects also corrects for the sun’s movement during the day, which might happen through a circadian mechanism (modelled²⁹) or synaptic plasticity. Therefore, downstream ring neurons roughly correspond to our z_c , but their responses might also reflect other inputs relevant to the sun’s position, such as spectral cues and even circadian corrections to create a true compass.

Fruit flies (*D. melanogaster*) have been shown to integrate their absolute compass with self-motion in two different stages of processing: first in the AOTu (optic flow input), and later in the ellipsoid body (feedback from motor neurons and optical flow)^{48,51}. In an interesting parallel, the use of Kalman filters^{12,13,56} or recurrent neural networks^{57,58} has been explored to improve the performance of some celestial compass sensors. This might therefore be a biologically plausible way to improve the performance of our compass. Insects also demonstrate colour opponency in the medulla, and this approach has been explored by Stürzl¹⁶ showing an advantage in distinguishing sky from tree branches to improve celestial compass performance under canopies. Thus, adding photoreceptors that respond to different wavelengths could also improve the performance of our compass. Celestial compass sensors can also be used for navigation at night⁴⁴ when the moon replaces the sun and forms a similar pattern of AoP and DoP⁵. The photodiodes of our sensor responded in twilight (when the moon had a weak effect on the polarisation pattern of the sky; Supplementary Fig. 4) and revealed a potential for a nocturnal (as well as diurnal) function. This aligns with the abilities of insects^{4,59} and, thus, a more

systematic test of the nocturnal abilities of our sensor in the future would be interesting.

Although our sensor might appear bulky (especially compared to camera approaches), we designed this prototype to be easily customised. In principle creating a miniaturised version seems straightforward, by integrating all the photodiodes onto one printed circuit board (PCB); and the first stages of optical processing to obtain I , p and c could happen onboard. An alternative could be a complementary metal-oxide-semiconductor (CMOS) with carefully tuned polarisers on the top to follow a fan-like arrangement (as in⁶⁰) or the full dome as described in²⁹. Nevertheless, to approach the size, speed, and efficiency of the insect brain, stretchable electronics^{61,62}, nanowire technology and photonic computations might be an interesting way forward.

Materials and Methods

Array of polarisation axis analysers. Each PAA was made of 3D-printed housing, a PCB, 4 photodiodes that were sensitive to UV light (221–358 nm), and 4 linear polarisers that allowed UV light (280–450 nm) with AoP at (1) 45°, (2) 135°, (3) 0°, and (4) 90° with respect to the vertical axis of the PAA. The specific materials of the photodiodes, filters, and PCB parts are summarised in Supplementary Table 7.

Figure 9a shows the block diagram of the designed PCB for each PAA. It includes four photodiodes (SG01D-18, SGLUX) followed by their trans-impedance amplifiers (TIAs). A Schematic of the designed TIA is provided in Fig. 9b. Linear technology LTC6082 quad low offset (60 μV), bias current (1 pA), voltage noise (13 nV · Hz^{-0.5}) and current noise (0.5 fA · Hz^{-0.5}) CMOS operational amplifier (opamp) was used where the photodiode was operating in photoconductive mode. The output voltage, V_{out} of the TIA was $I_{PD} \times R_F$, where I_{PD} is the photodiode reverse

current which is proportional to the wavelength of the light and R_F is the feedback resistor. A higher gain was achieved by using a 30 MΩ feedback resistor, which also helped to achieve a higher signal-to-noise ratio. A gain bandwidth product (GBWP) of 15 kHz was achieved by using a 6.8 pF feedback capacitor. This helped to improve the noise performance of the TIA and also provided a good transient response for the analogue-to-digital converter (ADC). The Texas Instruments ADS112C04 16-bit 4-channel precision delta-sigma ADC was selected. It was configured as differential mode ADC, where one channel was the input of the TIA output signal and another channel was fed with V_{Bias} of 2.5 V. A multiplexer was used to select the differential input signals from a particular TIA and the bias voltage. A programmable gain amplifier was acting as a differential amplifier with a gain of 1. This configuration improved the dynamic range of the ADC, which was configured to run in normal mode, with a data rate of 45 Hz (samples per second). The ADC was set to continuous conversion mode and the input multiplexer was configured as necessary during operation. All other configurations used default options. Two ADS112C04 chips were used to accommodate 4 TIAs in the PCB, which were configured with different inter-integrated circuit (I²C) addresses. The power management unit was built with the linear technology ADP7118-2.5V and ADP7118-5V for the bias and power supply voltage respectively. A 6-pin JST-GH connection was used to connect the PCB of each PAA to the battery and processing unit. Figure 9c shows the result PCB.

As each PCB uses the same configuration, we included an I²C multiplexer (TCA9548A, Adafruit) which allows sequential communication with each PCB. The I²C multiplexer was mounted in a custom breakout board that used Grove Universal 4-pin connectors (Seeed Studio) to interface with connected I²C PAAs, with a custom adaptor cable to link the 6-pin PCB

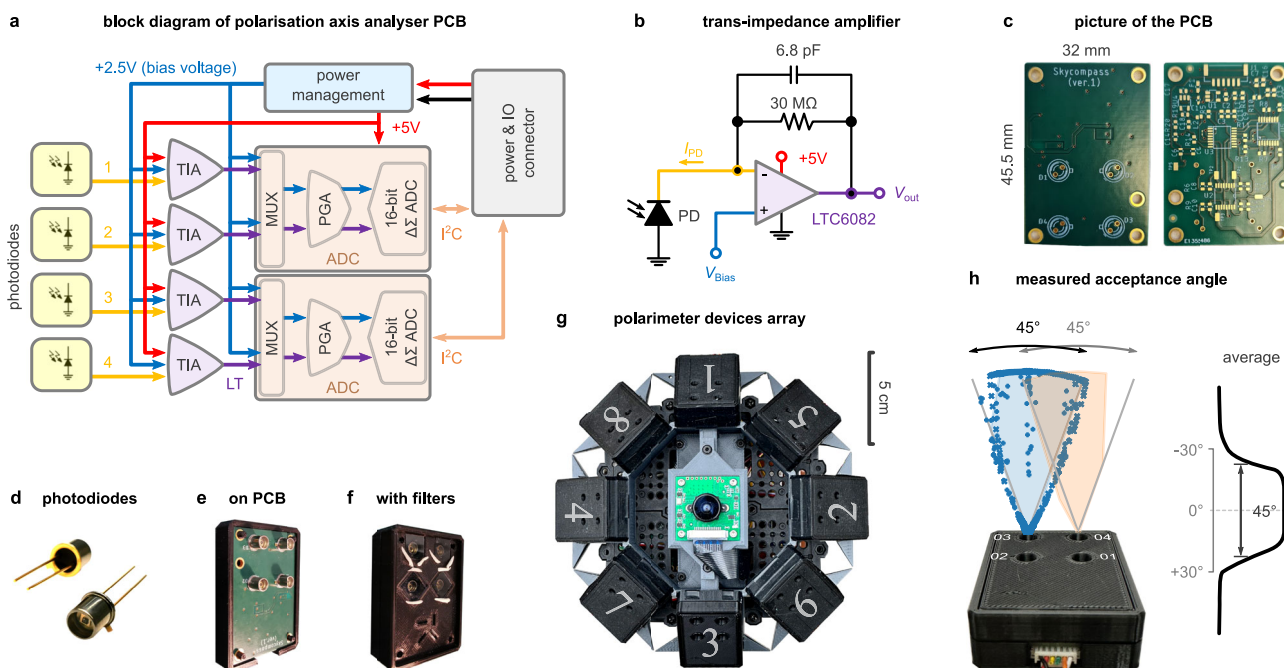


Fig. 9 Assembly of the celestial compass sensor. **a** Block diagram of the polarisation axis analyser (PAA) printed circuit board (PCB). Red lines represent positive voltage, blue lines are bias voltage, and the black line is ground; the rest are as indicated in the figure. TIA: trans-impedance amplifier, MUX: multiplexer, PGA: programmable gain amplifier, ADC: analogue-to-digital converter, I²C: inter-integrated circuit. **b** Schematic of the TIA circuit. PD: photodiode. **c** Front (left) and back (right) images of the PCB. **d** Picture of the photodiodes. **e** Four photodiodes on the PCB, and **(f)** in the 3D printed casing with polarisation filters attached. **g** The array of PAAs and the camera mounted on the 3D-printed caddy. Numbers denote the identity of each PAA and the order in which their output was read. **h** The acceptance angle was measured using a torch; the average response distribution of the photodiodes was estimated by using data from photodiodes 3 and 4.

connector to the breakout board. To communicate with the multiplexer (and PCBs) we used the I²C interface of Raspberry Pi and custom software using the Linux I²C/system management bus subsystem (adapted from existing Arduino libraries).

Our sensor was 3D printed using polylactic acid (parts were printed on Ultimaker 2+ extended and Ultimaker S3 3D printers). Each PCB had its own housing, comprised of a base, mid, and top plate. The base and mid plates closed around the PCB. The mid plate has square insets over each photodiode, holding the polarisers (OUV2525, Knight Optical) and giving each photodiode its polarisation tuning (see Fig. 9d–f). A small amount of adhesive putty (e.g., Blu Tack, Bostik, or similar) was used to fix the polarisation filters in place (Fig. 9e). The PCB with filters and housing creates one PAA (Fig. 9f). Eight PAAs were constructed and arranged on a 3D-printed caddy which positioned the units with an elevation of 45° (Fig. 1h). The caddy was mounted to the robot using two different configurations (due to external constraints on the robot at the time). While in Sardinia, the caddy was mounted to the robot using a pillar at the rear of the robot base. While in South Africa, the sensor was mounted directly to the top of the robot base. We noted no difference in operation or performance using the different mounts—both sets of computer-aided design (CAD) files are available in Supplementary Data 1. The camera module was mounted in the centre of the caddy (Fig. 9g).

To measure the effective acceptance angle of each PAA, we placed it vertically on a protractor and flashed UV light from the side with a torch (Alonefire SV10, at 365 nm wavelength). We targeted a single photodiode at a time from different angles, and logged the apparent angle and recorded pulses from the photodiode—blue and orange shaded areas in Fig. 9h for photodiodes 3 and 4 respectively. We then recorded from the photodiode 3 while swinging the torch with roughly stable angular velocity and inferred the apparent angle of the torch using the carefully collected samples from before and linear interpolation (these are illustrated with blue crosses in Fig. 9h). There was a small deviation in the centre of the acceptance between photodiodes 3 and 4, probably because of a small deviation in their placement angle and was not intentional (Fig. 9h). The average acceptance angle was measured roughly to 45°; the photodiodes responded more weakly for up to 60°.

The robot platform. The robot was constructed using a TurtleBot3 Burger kit (ROBOTIS). The kit contained a Raspberry Pi 3B+, OpenCR1.0, two Dynamixel XL430-W250 actuators, and a light detection and ranging (LIDAR) subsystem which we did not use (full details are available in the Turtlebot e-manual provided by ROBOTIS). In addition, the kit contained several standardised structural plates and beams to construct the body of the robot as appropriate. The kit was augmented with a 3rd party Raspberry Pi Camera module (B0103, Arducam) which was mounted such that it pointed towards the zenith. Figure 1h illustrates the final appearance of the robot. The IMU is embedded in the OpenCR1.0 board included with the TurtleBot kit and interaction with it was facilitated by libraries provided by ROBOTIS for the OpenCR1.0 platform.

The Turtlebot was controlled via the robot operating system (ROS). A ROS master node ran on a host laptop (ROS Noetic, Ubuntu 20.04) which generated a Wi-Fi hotspot. The Turtlebot (ROS Kinetic, Raspbian 9 stretch) was connected to this Wi-Fi hotspot as soon as it was available and was configured to view the host laptop IP as the ROS master. Thus, ROS nodes running on the Turtlebot and host laptop could communicate. For all data presented, the laptop ran only the recording routine (see below). On board, we ran nodes to read from the sensor and camera (as

well as all those concerned with the basic operation of the Turtlebot). On the host laptop, we ran the `roscore` and recording routine. All data were recorded using the `roscpp` C++ application programming interface (API). All interaction with the IMU was performed via ROS.

Data collection. Data was collected in three locations: Sardinia (Italy; 39.258648N, 8.440184E) in May 2022, Vryburg (South Africa; -26.398643N, 24.327144E) from 10 to 16 November 2022, and Bela Bela (South Africa; -24.714872N, 27.918972E) from 16 to 29 November 2022. During each recording session, the robot was placed on a smooth level surface facing approximately north. The robot would then record statically for five seconds, rotate slowly through 360°, and then record statically for five seconds. On subsequent rotations, the robot would correct using the IMU for any over-rotation before starting to record. The IMU was periodically re-initialised to prevent noticeable drifts over the course of a recording session. At the beginning of each recording session, a time-stamped sky photo was captured using the onboard camera. Recording sessions lasted around 12–15 minutes on average and consisted of 12 rotations under the condition of interest (different solar elevations, cloud cover, or occlusions).

Compass models. The output of the photodiodes (voltage pulses scaled down by a factor of $11 \cdot 10^3$ and clipped to $[0, 1]$) was used from our model to estimate the solar azimuth, which represented the compass direction. Using this prediction, an ephemeris function (describing the course of the sun during the day) and the coordinated universal time (UTC) we can accurately estimate the true north as well as an approximation to the coordinates of the sensor on Earth. Equation (1), (2), (3), (4), and (5) describe how our model transforms the photodiode pulses into a compass direction. For comparison, we implemented two additional compass models: the eigenvectors^{16,40} and the four-zeros⁴¹. All the models were implemented in Python 3.

For the eigenvector model, we first needed to calculate the AoP (a) and DoP (d) from the photodiode responses. We calculated these following Zhao et al.⁴⁰, whose PAAs were similar to ours,

$$a_n = \frac{1}{2} \arctan \frac{r_{2,n} - r_{1,n} + r_{1,n}r_{2,n} - 1}{r_{1,n} - r_{2,n} + r_{1,n}r_{2,n} - 1} + \phi_n, \quad (6)$$

$$d_n = \frac{r_{1,n} - 1}{(r_{1,n} + 1) \cos(2a_n)}, \quad (7)$$

$$r_{1,n} = \frac{I_{0,n}}{I_{90,n}}, \quad r_{2,n} = \frac{I_{45,n}}{I_{135,n}}, \quad (8)$$

where $\phi_n = 2\pi(n-1)/N$ is the azimuth angle of the respective PAA, and n is its identity as illustrated in Fig. 9g. The respected elevation of each PAA is always 45° ($\pi/4$). The polarisation vector associated with each PAA is perpendicular to the respective AoP and points towards the inside of our sensor,

$$\mathbf{p}_n = [\sin(a_n + \pi) \quad \cos(a_n + \pi)]. \quad (9)$$

Weighting each of these vectors with the respective DoP (d_n) did not affect any of our results. The covariance matrix of these vectors was calculated as

$$\mathbf{C} = \mathbf{P} \cdot \mathbf{P}^T, \quad (10)$$

and its eigenvectors represent the principal facing axes. The eigenvector with the highest eigenvalue ($\hat{\mathbf{e}}$) points towards the solar azimuth¹⁶. We calculated the eigenvectors and eigenvalues of the covariance using the NumPy package⁶³, and transformed this eigenvector with the highest eigenvalue into a complex

number for consistency with the other two models,

$$z_{ev} = \hat{e}_2 + i \hat{e}_1, \quad (11)$$

while the estimated solar azimuth (α_{ev}) can be calculated using equation (5). Three examples of the AoP, \mathbf{p} vectors, and solar azimuth estimations of this model are illustrated in Supplementary Fig. 3a. For the four-zeros model, we used as input the polarisation responses (p_n) of N PAAs, which were computed using equation (2). For this model, we needed to fit a curve on the responses, which could be described by the first three Fourier coefficients of the DFT. However, we were interested in centring the curve to zero, so the zero-coefficient can be omitted,

$$z_1 = \frac{2}{N} \sum_{n=1}^N p_n e^{i2\pi(n-1)/N}, \quad (12)$$

$$z_2 = \frac{2}{N} \sum_{n=1}^N p_n e^{i4\pi(n-1)/N}. \quad (13)$$

The magnitude of these coefficients is $\rho_k = \|z_k\|$ ($k \in \{1, 2\}$), and their angle (α_k) is given by equation (5). The function describing the curve of the responses is then

$$f(\theta) = \rho_1 \cos(\alpha_1 - \theta) + \rho_2 \cos(\alpha_2 - 2\theta), \quad (14)$$

which is plotted in Supplementary Fig. 3b in black, and its derivative is

$$\frac{df}{d\theta} = \rho_1 \sin(\alpha_1 - \theta) + 2\rho_2 \sin(\alpha_2 - 2\theta). \quad (15)$$

We used these equations as input to the Newton-Raphson optimiser (from the SciPy package⁶⁴) to estimate the four solutions of equation (14). We ran the optimization four times with different initialisations that were homogeneously distributed around the angle of the second coefficient (i.e., $\theta_{\text{init}} = \alpha_2/2 \pm \pi/2 \pm \pi/4$). This ensured that the optimisation starts roughly at the correct position and falls in the correct local solution. Examples of these solutions are plotted in Supplementary Fig. 3b (red dots), demonstrating the correctness of the method so far. Next, the angles of the four zeros were normalised in $[0, 2\pi)$ and sorted. The absolute difference between the consecutive solutions was calculated, and the pair of solutions with the lowest difference was identified,

$$\delta_m = \|(\theta_{(m+1 \bmod 4)} - \theta_m + \pi) \bmod 2\pi - \pi\|, \quad (16)$$

$$\hat{m}_1 = \arg \min_{m=1}^4 \delta_m, \quad (17)$$

$$\hat{m}_2 = \hat{m}_1 + 1 \bmod 4 \quad (18)$$

where \bmod denotes and modulo operation, and $\arg \min$ finds the identity that represents the minimum value of δ_m . The estimated solar azimuth is then calculated,

$$\alpha_{fz} = \theta_{\hat{m}_1} + \frac{\delta_{\hat{m}_2}}{2}. \quad (19)$$

Performance evaluation. We evaluated the performance of the models by using the RMSE as a global and local error measurement across different orientations of the sensor in the same sky condition and rotation. The RMSE for a specific model, sky condition, and rotation was calculated as

$$\text{RMSE}_{\text{model}} = \sqrt{\frac{1}{360} \sum_{s=1}^{360} \epsilon_{\text{model},s}^2}, \quad (20)$$

where $\epsilon_{\text{model},s}$ represents the (global or local) error of the model, and $s \in \{1, \dots, 360\}$ are the homogeneously distributed

orientations of the sensor during a rotation with respect to the starting orientation.

The global error represents the overall deviation of the estimated solar azimuth from the real solar azimuth during a rotation ($\pm 5^\circ$, due to the approximate initialisation towards the north),

$$\epsilon_{\text{model},s}^G = (\alpha_{\text{model},s} - \hat{\alpha}_s + \pi) \bmod (2\pi) - \pi, \quad (21)$$

where $\alpha_{\text{model},s}$ is the prediction of the model for the solar azimuth (with respect to the front of the sensor) when the sensor was in orientation s , and $\hat{\alpha}_s$ is the respective true solar azimuth ($\pm 5^\circ$).

The local RMSE is the overall deviation of the estimated solar azimuth from the average estimated solar azimuth, and it represents the consistency of the predictions of the sensor,

$$\epsilon_{\text{model},s}^L = (\alpha_{\text{model},s} - \bar{\alpha}_{\text{model}} + \pi) \bmod (2\pi) - \pi, \quad (22)$$

where $\bar{\alpha}_{\text{model}}$ is the average prediction of the model for the solar azimuth across the different tested orientations (s).

Data availability

The data regarding the sensor design (CAD files) are provided in Supplementary Data 1. The raw and summarised data that were collected using the robot are publicly available in DataShare with identifier <https://doi.org/10.7488/ds/6106>.

Code availability

The code used for the robot, data collection, analysis and plots is available on GitHub. Code for robot and data collection is available by the authors upon request. Code for the data analysis and generation of plots is publicly available in zenodo with identifier <https://doi.org/10.5281/zenodo.8393056>.

Received: 9 June 2023; Accepted: 30 October 2023;

Published online: 15 November 2023

References

- Vickerstaff, R. J. & Cheung, A. Which coordinate system for modelling path integration? *Journal of Theoretical Biology* **263**, 242–261 (2010).
- Menzel, R. & Snyder, A. W. Polarised light detection in the bee, *Apis mellifera*. *Journal of comparative physiology* **88**, 247–270 (1974).
- Labhart, T. Polarization-opponent interneurons in the insect visual system. *Nature* **331**, 435–437 (1988).
- Greiner, B., Cronin, T. W., Ribi, W. A., Wcislo, W. T. & Warrant, E. J. Anatomical and physiological evidence for polarisation vision in the nocturnal bee *Megalopta genalis*. *J. Compar. Physiol. A* **193**, 591–600 (2007).
- Barta, A. et al. Polarization transition between sunlit and moonlit skies with possible implications for animal orientation and Viking navigation: anomalous celestial twilight polarization at partial moon. *Appl. Opt.* **53**, 5193 (2014).
- Rayleigh, L. X. On the electromagnetic theory of light. *Philoso. Magaz. Series 5* **12**, 81–101 (1881).
- Rayleigh, L. XXXIV. On the transmission of light through an atmosphere containing small particles in suspension, and on the origin of the blue of the sky. *Philoso. Magaz. Series 5* **47**, 375–384 (1899).
- Strutt, J. W. XV. On the light from the sky, its polarization and colour. *London Edinburgh Dublin Philoso. Magaz. J. Sci* **41**, 107–120 (1871).
- Wang, Y., Chu, J., Zhang, R., Wang, L. & Wang, Z. A novel autonomous real-time position method based on polarized light and geomagnetic field. *Sci. Rep.* **5**, 9725 (2015).
- Wang, Y., Chu, J., Zhang, R. & Shi, C. Orthogonal vector algorithm to obtain the solar vector using the single-scattering Rayleigh model. *Appl. Opt.* **57**, 594 (2018).
- Yang, Z., Zhou, J., Huang, H., Liu, Y. & Li, Z. Measuring solar vector with polarization sensors based on polarization pattern. *Optik - Int. J. Light Electr. Opt.* **141**, 147–156 (2017).
- Du, T. et al. An autonomous initial alignment and observability analysis for SINS with bio-inspired polarized skylight sensors. *IEEE Sensors J.* **20**, 7941–7956 (2020).
- Liu, X., Yang, J., Guo, L., Yu, X. & Wang, S. Design and calibration model of a bioinspired attitude and heading reference system based on compound eye polarization compass. *Bioinspir. Biomimet.* (2021).

14. Carey, N. & Stürzl, W. An insect-inspired omnidirectional vision system including UV-sensitivity and polarisation. *2011 IEEE Int. Confer. on Computer Vision Workshops (ICCV Workshops)* **1**, 312–319 (2011).
15. Stürzl, W. & Carey, N. A fisheye camera system for polarisation detection on UAVs. In *Computer Vision - ECCV 2012. Workshops and Demonstrations. ECCV 2012*, vol. 7584 of *Lecture Notes in Computer Science*, 431–440 (Springer, 2012).
16. Stürzl, W. A lightweight single-camera polarization compass with covariance estimation. *2017 IEEE International Conference on Computer Vision (ICCV)* 5363–5371 (2017).
17. Han, G. et al. Design and calibration of a novel bio-inspired pixelated polarized light compass. *Sensors* **17**, 2623 (2017).
18. Wang, Y., Hu, X., Lian, J., Zhang, L. & He, X. Bionic orientation and visual enhancement with a novel polarization camera. *IEEE Sensors J.* **17**, 1316–1324 (2017).
19. Fan, C., Hu, X., He, X., Zhang, L. & Lian, J. Integrated polarized skylight sensor and MIMU with a metric map for urban ground navigation. *IEEE Sensors J.* **18**, 1714–1722 (2018).
20. Han, G. et al. A polarized light compass aided place recognition system. *2019 Third IEEE Int. Confer. Robot. Comput. (IRC)* **00**, 266–270 (2019).
21. Huang, J. et al. A bio-inspired array skylight compass for urban ground navigation. In *Proceedings of International Conference on Guidance, Navigation and Control*, Proceedings of International Conference on Guidance, Navigation and Control, 6599–6609 (2023).
22. Sarkar, M., Bello, D. S. S., Hoof, C. v. & Theuwissen, A. Integrated polarization analyzing CMOS image sensor for autonomous navigation using polarized light. *2010 5th IEEE International Conference Intelligent Systems* 224–229 (2010).
23. Zhang, S., Liang, H., Zhu, H., Wang, D. & Yu, B. A camera-based real-time polarization sensor and its application to mobile robot navigation. *2014 IEEE International Conference on Robotics and Biomimetics (ROBIO 2014)* 271–276 (2014).
24. Wang, D., Liang, H., Zhu, H. & Zhang, S. A bionic camera-based polarization navigation sensor. *Sensors* **14**, 13006–13023 (2014).
25. Dupeyroux, J., Viollet, S. & Serres, J. R. An ant-inspired celestial compass applied to autonomous outdoor robot navigation. *Robot. Autono. Syst.* **117**, 40–56 (2019).
26. Lu, H., Zhao, K., You, Z. & Huang, K. Angle algorithm based on Hough transform for imaging polarization navigation sensor. *Opt. Expr.* **23**, 7248–7262 (2015).
27. Wu, X. et al. Robust orientation method based on atmospheric polarization model for complex weather. *IEEE Int. Things J.* **PP**, 1–1 (2022).
28. Bech, M., Homberg, U. & Pfeiffer, K. Receptive fields of locust brain neurons are matched to polarization patterns of the sky. *Curr. Biol.* **24**, 2124–2129 (2014).
29. Gkaniats, E., Risse, B., Mangan, M. & Webb, B. From skylight input to behavioural output: a computational model of the insect polarised light compass. *PLoS Comput. Biol.* **15**, e1007123 (2019).
30. Zollkoffer, C., Wehner, R. & Fukushi, T. Optical scaling in conspecific *Cataglyphis* ants. *J. Exp. Biol.* **198**, 1637–1646 (1995).
31. Borst, A. *Drosophila's* view on insect vision. *Curr. Biol.* **19**, R36–R47 (2009).
32. Wehner, R. 'Matched filters' — neural models of the external world. *J. Compar. Physiol. A* **161**, 511–531 (1987).
33. Lambrinos, D. et al. An autonomous agent navigating with a polarized light compass. *Adap. Behav.* **6**, 131–161 (1997).
34. Lambrinos, D., Möller, R., Labhart, T., Pfeifer, R. & Wehner, R. A mobile robot employing insect strategies for navigation. *Robot. Autono. Syst.* **30**, 39–64 (2000).
35. Chu, J., Zhao, K., Zhang, Q. & Wang, T. Construction and performance test of a novel polarization sensor for navigation. *Sensors Actuators A: Phys.* **148**, 75–82 (2008).
36. Zhao, K., Chu, J., Wang, T. & Zhang, Q. A novel angle algorithm of polarization sensor for navigation. *IEEE Trans. Instrument. Measure.* **58**, 2791–2796 (2009).
37. Chahl, J. & Mizutani, A. Biomimetic attitude and orientation sensors. *IEEE Sensors J.* **12**, 289–297 (2012).
38. Xian, Z. et al. A novel angle computation and calibration algorithm of bio-inspired sky-light polarization navigation sensor. *Sensors* **14**, 17068–17088 (2014).
39. Chu, J. et al. Integrated polarization dependent photodetector and its application for polarization navigation. *IEEE Photo. Technol. Lett.* **26**, 469–472 (2014).
40. Zhao, H. & Xu, W. A bionic polarization navigation sensor and its calibration method. *Sensors* **16**, 1223 (2016).
41. Smith, F. J. Insect navigation by polarised light. In *Proceedings of the International Conference on Bio-inspired Systems*, Proceedings of the International Conference on Bio-inspired Systems and Signal Processing, 363–368 (SciTePress, 2009).
42. Zar, J. H. *Biostatistical analysis* (Pearson Education India, 1999).
43. Berens, P. CircStat: a Matlab toolbox for circular statistics. *J. Stat. Software* **31** (2009).
44. Guan, L. et al. Study on skylight polarization patterns over the ocean for polarized light navigation application. *Appl. Opt.* **57**, 6243 (2018).
45. Smith, F. J. A new algorithm for navigation by skylight based on insect vision. In *Proceedings of the First International Conference on Bio-inspired Systems and Signal Processing*, vol. 2 of *Proceedings of the First International Conference on Bio-inspired Systems and Signal Processing*, 185–190 (SciTePress, 2008).
46. Smith, F. J. & Stewart, D. W. Robot and insect navigation by polarized skylight. In *Proceedings of the International Conference on Bio-inspired Systems*, Proceedings of the International Conference on Bio-inspired Systems and Signal Processing, 183–188 (SciTePress, 2014).
47. Weir, P. T. et al. Anatomical reconstruction and functional imaging reveal an ordered array of skylight polarization detectors in *Drosophila*. *J. Neurosci.* **36**, 5397–5404 (2016).
48. Kind, E. et al. Synaptic targets of photoreceptors specialized to detect color and skylight polarization in *Drosophila*. *eLife* **10**, e71858 (2021).
49. Schnaitmann, C. et al. Color processing in the early visual system of *Drosophila*. *Cell* **172**, 318–330.e18 (2018).
50. Sancer, G. et al. Modality-specific circuits for skylight orientation in the fly visual system. *Curr. Biol.* **29**, 2812–2825.e4 (2019).
51. Hardcastle, B. J. et al. A visual pathway for skylight polarization processing in *Drosophila*. *eLife* **10**, e63225 (2021).
52. Pfeiffer, K., Kinoshita, M. & Homberg, U. Polarization-sensitive and light-sensitive neurons in two parallel pathways passing through the anterior optic tubercle in the locust brain. *J. Neurophysiol.* **94**, 3903–3915 (2005).
53. Pfeiffer, K. & Kinoshita, M. Segregation of visual inputs from different regions of the compound eye in two parallel pathways through the anterior optic tubercle of the bumblebee (*Bombus ignitus*). *J. Compar. Neurol.* **520**, 212–229 (2012).
54. Heinze, S., Florman, J., Asokaraj, S., Jundi, B. E. & Reppert, S. M. Anatomical basis of sun compass navigation II: The neuronal composition of the central complex of the monarch butterfly. *J. Compar. Neurol.* **521**, 267–298 (2013).
55. Scheffer, L. K. et al. A connectome and analysis of the adult *Drosophila* central brain. *eLife* **9**, e57443 (2020).
56. Du, T., Zeng, Y. H., Yang, J., Tian, C. Z. & Bai, P. F. Multi-sensor fusion SLAM approach for the mobile robot with a bio-inspired polarised skylight sensor. *IET Radar Sonar Navig* **14**, 1950–1957 (2020).
57. Zhao, D. et al. Attitude-induced error modeling and compensation with GRU networks for the polarization compass during UAV orientation. *Measurement* **190**, 110734 (2022).
58. Zhao, D. et al. Multiscale parallel heading error processing model for polarization compass. *IEEE Sensors J.* **22**, 17882–17892 (2022).
59. Dacke, M., Nilsson, D.-E., Scholtz, C. H., Byrne, M. & Warrant, E. J. Insect orientation to polarized moonlight. *Nature* **424**, 33–33 (2003).
60. Tokuda, T., Yamada, H., Sasagawa, K. & Ohta, J. Polarization-analyzing CMOS image sensor with monolithically embedded polarizer for microchemistry systems. *IEEE Trans. Biomed. Circuits Syst.* **3**, 259–266 (2009).
61. Song, Y. M. et al. Digital cameras with designs inspired by the arthropod eye. *Nature* **497**, 95–99 (2013).
62. Gu, L. et al. A biomimetic eye with a hemispherical perovskite nanowire array retina. *Nature* **581**, 278–282 (2020).
63. Harris, C. R. et al. Array programming with NumPy. *Nature* **585**, 357–362 (2020).
64. Virtanen, P. et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Methods* **17**, 261–272 (2020).

Acknowledgements

This work was financially supported in part by the European Union, Horizon Europe, (Project 101046790, InsectNeuroNano), the European Research Council (817535, Ultimate-COMPASS), and the UKRI Engineering and Physical Sciences Research Council (Grant No. EP/R513209/1, EP/M008479/1, and Impact Acceleration Account PIV068). We thank Marie Dacke from Lund University for facilitating field testing by providing travel and access to the field sites used for data collection. Thanks to the members of the Insect Robotics Group from the University of Edinburgh for commenting on earlier versions of the manuscript. Also, thanks to the InsectNeuroNano team for comments on the results of this work.

Author contributions

E.G.: conceptualisation, formal analysis, investigation, methodology, project administration, software, validation, visualisation, writing—original draft, writing—review & editing. R.M.: conceptualisation, data collection, investigation, methodology, software, hardware, visualisation, writing—original draft, writing—review & editing. J.S.: conceptualisation, funding acquisition, investigation, hardware. Sadeque Reza Khan: conceptualisation, hardware, visualisation, writing—original draft, writing—review &

editing. S.M.: conceptualisation, resources, supervision, writing—review & editing. Barbara Webb: conceptualisation, funding acquisition, investigation, methodology, project administration, resources, supervision, validation, writing—review & editing.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s44172-023-00132-w>.

Correspondence and requests for materials should be addressed to Evripidis Gkanias.

Peer review information *Communications Engineering* thanks Young Min Song, Uwe Homberg and the other, anonymous reviewer for their contribution to the peer review of this work. Primary Handling Editors: Yu-cheng Chen and Mengying Su, Rosamund Daw. A peer review file is available.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023