A central complex model of dung beetle celestial menotaxis

Ramsey H. El-Naggar

THE UNIVERSITY
of EDINBURGH

MInf Project (Part 2) Report

Master of Informatics
School of Informatics
University of Edinburgh

2018

Acknowledgements

First and foremost, I'd like to thank my supervisor Professor Barbara Webb for her continuous support, patience and motivation throughout both years of the project. I might not always have been the best student, but I couldn't have asked for a better supervisor.

Next, I'd like to thank Professor Marie Dacke of Lund University for organising the field trip to Johannesburg, and for allowing me to accompany her and her team. Professor Dacke also provided insightful conversation and constant cheer and motivation in the face of unfortunate weather conditions. The thanks extend to the rest of the team: James Foster, Adrian Bell, and Lana Khaldy of Lund University, Assistant Professor Orit Peleg of the University of Colorado Boulder, and in particular Professor Marcus Byrne of the University of the Witwatersrand, for insightful discussions, assistance with experiments, and taking fantastic photographs and videos. The working environment is only as fun as the people in it, and I'd like to thank everyone there for making the field trip one of the major highlights of my university experience.

The field trip to Johannesburg to collect results for this project was financially supported by alumni and friends of the University of Edinburgh through an Innovation Initiative Grant, and I'd like to thank all those who make such grants possible by donating to the Edinburgh Fund.

Finally, a special thanks to Garry Ellard for his patient and tireless technical support throughout both years of the project – there would be no robot without Garry.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

This is a report detailing the second year of a two year project applying the biorobotics methodology to research concerning dung beetle celestial menotaxis – the ability of dung beetles to use visual information from their view of the sky to navigate in a straight-line path in a chosen direction. In the first year of the project [1], a biorobotic platform for the project was constructed and the robot was shown to be able to produce menotaxis behaviour using simple implementations of mechanisms derived from biological research. The central task of this year of the project was to produce a neurologically plausible model (based on [2]) of the same behavioural mechanisms, focusing on the central complex (CX) region of the dung beetle brain – a region that has been shown to be involved in sensorimotor control of different navigational tasks in many insects – and implement and test this model on the existing biorobotic platform.

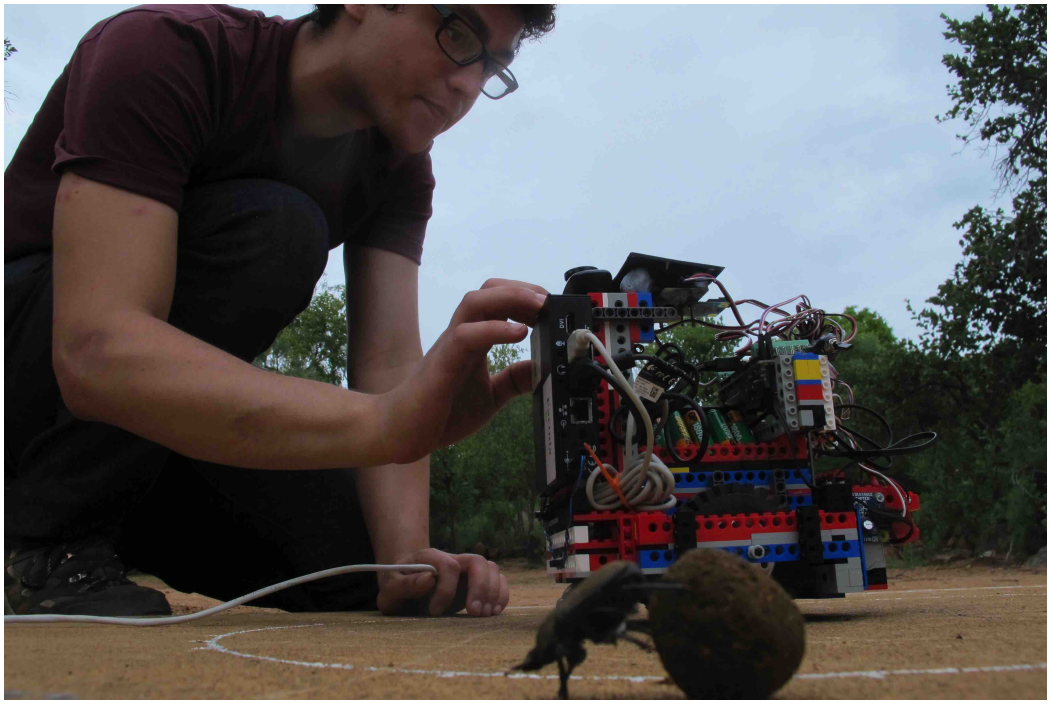

Figure 1.1: Who's more intelligent?

The remainder of Chapter 1 is structured as follows: Section 1.1 gives a brief account of the most relevant work carried out during the first year of the project; Section 1.2 develops the research questions for this year of the project, ending with a statement of concrete aims, and for each aim an indication is given of whether or not that aim was met successfully; Section 1.3 gives a concise list of contributions made; and Section 1.4 provides an overview of the remaining chapters.

1.1 Previous work carried out

This section gives a brief account of the relevant work carried out during the first year of the project [1]. Before reading this section, the reader may first wish to read Section 2.1 of this report for an overview of relevant background concepts concerning the environment and behaviour of dung beetles.

The main achievement of the first year of the project was the construction of a biorobot that was successfully able to menotax using an algorithmic 'snapshot' model inspired by the literature [3] using either a polarisation cue or an intensity-gradient cue. Proof-of-concept experiments demonstrated that the robot could consistently and accurately (to within less than 6° absolute error) maintain a randomly chosen straight-line bearing using either of these cues, both in an indoor arena and outside under the Scottish spring-time weather.

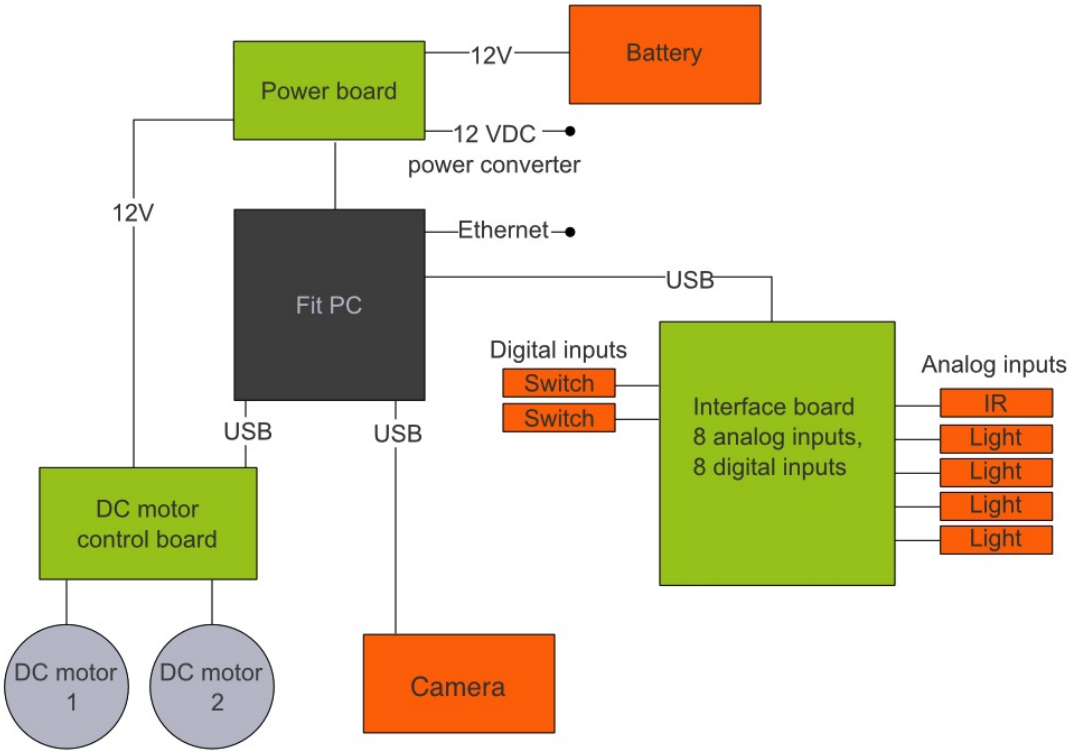

Figure 1.2: Robot hardware schematic, modified from [4]

1.1.1 Robot construction

The robot, $20cm \times 20cm \times 24cm$, was constructed from LEGO and electronic hardware components. The main hardware components were a fit-PC2 (1.6GHz single-core processor, 1GB RAM) running Tiny Core Linux, a PhidgetMotorControl LV (1060_0) motor board and a PhidgetInterfaceKit 8/8/8 (1018_2) interface board, powered by a 12V battery (or mains power supply) via a custom-built power board (built by Garry Ellard). An Ethernet connection between the fit-PC (running the Tiny Core distribution of Linux) and a laptop allowed the system to be developed on the laptop, and then uploaded to and tested on the fit-PC, running on-board and wirelessly (necessary as the control task of menotaxis relies on real-time feedback from the environment).

1.1.2 Motor control

For the purposes of simple menotaxis, the complex details of hexapedal movement were simplified to a two-wheeled robot design (with a single ball bearing), capable of turning on the spot (mimicking the dung beetle dance), moving straight, or curving to the side whilst moving forward (mimicking the way dung beetles correct for deviations to their straight-line path). A pretend dung ball was considered, but the physical dynamics of ball rolling were not believed to be necessary for successful menotaxis. Each LEGO wheel was driven by a LEGO DC Electric Technic Mini-Motors (9V) connected to the motor board.

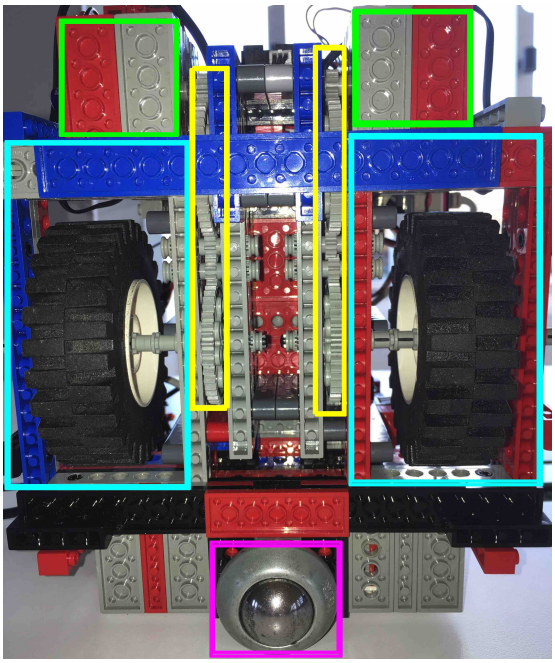

Figure 1.3: Robot wheel base – view from below, showing positions of wheel-frames (cyan), motors (green), connecting gear-chains (yellow), and the ball bearing (pink); modified from [1, Figure 3.2]

The motor-wheel gearing ratio was 1:5 such that the robot moves 15cm/s at full speed on a flat smooth surface. Pulse-width modulation of the 100% power motor signal, alternating at 7Hz, was used to reduce this to a speed more comparable to dung beetles.

1.1.3 The polarisation cue

Visual sensors, placed on the top of the robot and facing skyward, were chosen to capture information from the sky thought to be used by dung beetles during menotaxis, without considering the exact means that dung beetles acquire this information.

To sense polarisation information, the robot has four Phidgets Precision Light Sensors (1127_0) connected to the interface board. Inspired by the Sahabot [5], each light sensor was covered by a linearly polarising filter set at a different relative bearing from $\{0^\circ, 45^\circ, 90^\circ, 135^\circ\}$ such that the first three Stokes parameters [6][7] can be calculated experimentally in real-time (we denote light sensor intensity readings as I_α, the total light intensity as I, the degree of polarisation as $p \in [0,1]$, and the ellipse of polarisation of the largest polarised component as described by angles Ψ and $\chi \in [0,\pi)$):

$$S_0 = I = I_0 + I_{90} = I_{45} + I_{135} = \frac{I_0 + I_{90} + I_{45} + I_{135}}{2} \ ; \tag{1.1}$$

$$S_1 = Q = Ip\cos(2\Psi)\cos(2\chi) = I_0 - I_{90} \ ; \tag{1.2}$$

$$S_2 = U = Ip\sin(2\Psi)\cos(2\chi) = I_{45} - I_{135} \ . \tag{1.3}$$

Note that raw light values are used throughout – the 'calibration' procedure devised in the first year of the project was not found to be necessary for successful menotaxis. This follows the theme that the scope of the project has been narrowed this year to no longer attempt to model the dance of the dung beetle.

We can then use the assumption of the Rayleigh sky model [8] that there is little non-linearly polarised light from the sky to set $\chi = 0$, turning the ellipse of polarisation into a line. The angle of **linear** polarisation of the largest **linearly** polarised component, Ψ, can then be calculated using:

$$\Psi = \frac{\arctan\frac{U}{Q}}{2} \tag{1.4}$$

The assumption of the Rayleigh sky model also allows us to set the fourth Stokes parameter $S_3 = V = Ip\sin(2\chi) = 0$, allowing the intensity of the largest linearly polarised component (and degree of linear polarisation, $p \in [0,1]$ by dividing by I) to be calculated with:

$$Ip = \sqrt{Q^2 + U^2 + V^2} = \sqrt{Q^2 + U^2} \tag{1.5}$$

Note that the calculation of degree of polarisation was implemented and its use as a metric for cue preference was conceived and briefly explored during the first year of the project, but its use was not finalised until this year of the project (see Sections 3.1 and 3.2.2).

1.1.4 The centroid vector (CV) intensity-gradient cue

To sense light intensity information, the robot has a Logitech C270 webcam connected via USB to the fit-PC2, with the field of view extended to about 100° using a fish-eye lens. The camera image (visual field) is first converted into greyscale, after which two methods were devised to extract an intensity-based cue.

Inspired by the assumption in the literature (e.g. [9]) that point-sources (celestial bodies such as the sun and the moon) and weak intensity-gradient information are treated as separate cues by the beetle, a method to extract directional information from the

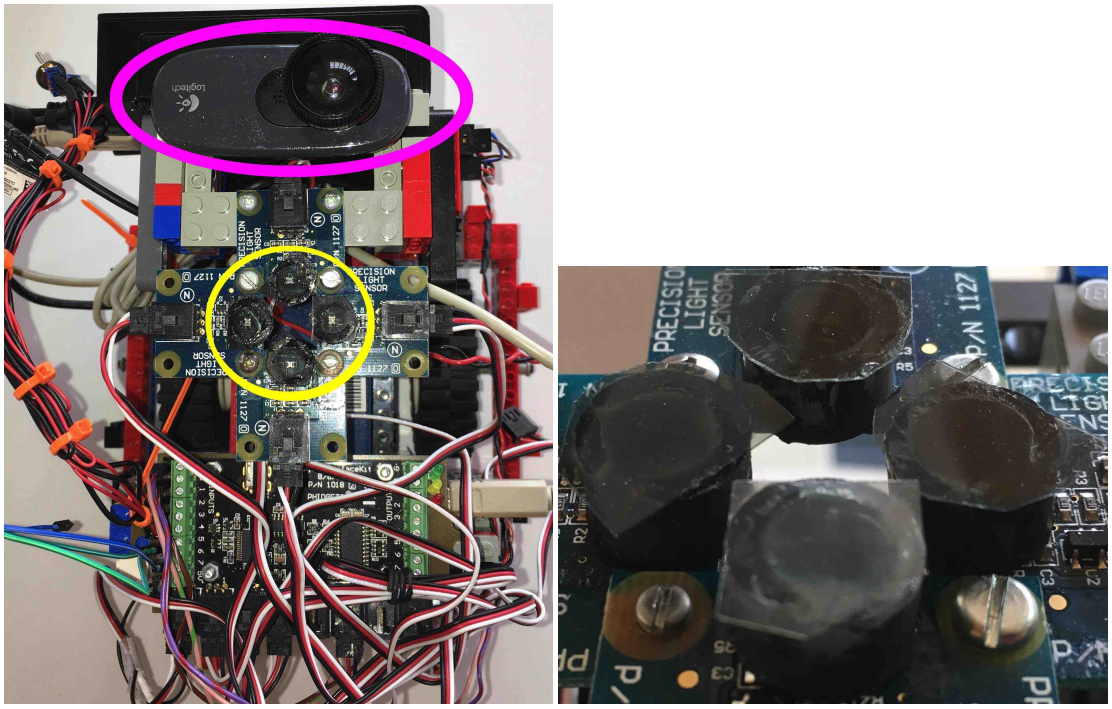

Figure 1.4: Robot visual sensors – view from above the robot (left), showing the polarised light sensors (yellow) and camera (pink), modified from [1, Figure 3.5]; close-up of the polarised light sensor (right), modified from [1, Figure 3.7]

view was to identify the brightest point (location of the point-source) in the sky (Gaussian convolution was applied to extend this idea to the brightest area, as in [10]) and construct a vector (called the brightest vector BV) from the center of the field of view to this brightest point. However, issues to do with saturation and multiple brightest points, and noise due to visual speculation and cloud movement, resulted in the experimental failure of attempts to menotax using the BV cue.

The contrary approach was therefore taken to treat point-sources and weak intensity gradient as a single intensity-gradient cue. Inspired from the literature by a visual homing algorithm known as COMALV (center-of-mass average landmark vector) [11][12][13], the centroid-vector (CV) is simply the vector pointing from the center of the visual field to the centroid (aka center-of-mass) of the visual field. In an image with uniform brightness, the centroid lies in the center of the image, whereas if one side of the image is brighter than the other, the centroid lies on the brighter side. The bearing of the CV therefore combines information from a point source and the weak light intensity-gradient to give a single direction. Additionally, the magnitude of the CV can be used as a metric for cue preference (see Sections 3.1 and 3.2.2).

The failure of the BV method and success of the CV method lead to the conclusions that (i) dung beetles probably do not treat point-source and weak intensity-gradient cues separately, and (ii) a CV-like method is the most plausible explanation so far.

1.1.5 The algorithmic menotaxis model (Model A)

The algorithmic menotaxis model (Model A) was developed for a single cue. We define the `current_bearing` as the directional information returned by the cue currently in use: for the polarisation cue, this is the angle of linear polarisation of the largest linearly polarised component (Ψ above); and for the intensity-gradient cue, this is the bearing of the centroid vector, denoted α_{CV}.

The inputs to Model A each iteration are the `current_bearing` and a boolean (denoted `snapshot_trigger`) indicating whether or not the current bearing has been chosen to be the new snapshot bearing, which will subsequently be maintained. The output of Model A is an error signal (denoted `error`) indicating the error between the current bearing and the snapshot bearing. This error signal can then be used by any standard control mechanism to maintain the snapshot bearing during forward movement, resulting in menotaxis. Thus the following is repeated every iteration:

- If `snapshot_trigger` is true set `snapshot_bearing` equal to `current_bearing` and reset `snapshot_trigger` to be false for subsequent iterations.
- Return `error` equal to the angular difference between `current_bearing` and `snapshot_bearing` (taking into account that the range of bearings $Psi \in [0,\pi)$ and $\alpha_{CV} \in [0,2\pi)$).

In the first year of the project, the control mechanism used to maintain the snapshot angle, denoted `control_straighten` was as follows (repeated every iteration):

- Inputs: `error`; an angular threshold θ; the current speed s (percentage of full speed).
- If $|\texttt{error}| > \theta$, determine the rotation required (clockwise or anticlockwise) to reduce the error signal, and rotate appropriately at speed s.
- Otherwise, move forward at speed s.

1.2 Aims

The research questions for this year of the project were chosen to match current lines of investigation in the biological study of dung beetle celestial menotaxis, particularly with relation to larger neuroethological investigations concerning the role of the CX in insect navigational behaviours (e.g. [2]). Taking the recommendations from the first year of the project, the scope of this year of the project was narrowed as far as possible to focus on three particularly interesting and interrelated questions:

- How are celestial visual cues, particularly the intensity-gradient and polarisation cues, encoded in the dung beetle CX for menotaxis purposes?
- How are different cues combined in the CX during menotaxis?

- By what neurologically plausible mechanism does the dung beetle CX produce menotax behaviour?

These general questions constitute the rational for all work carried out during this year of the project. The focus on neurological plausibility reflects the key shift in direction of this this year of the project compared to the first year – within a sufficiently narrow scope (explained in Section 3.2), whenever possible only 'neurologically plausible' mechanisms and implementations were considered.

In focusing on these three research questions, previous topics and questions naturally fell out of the scope of the project, the most relevant of which are listed here:

- the use of cues other than polarisation and intensity-gradient (in particular, the spectral cue is no longer considered); and
- the triggering factors, neurological mechanism and purpose of the dung beetle dance, and its exact relation to the snapshot mechanism (outwith its designation as the time during which the snapshot is taken).

Considering the research questions above, the central task of the project (developing, implementing on the biorobot, and testing a neurologically plausible CX model of dung beetle celestial menotaxis) was broken down to give more concrete aims:

1. Develop a high-level hypothesis that explains how cue combination occurs in the context of relevant dung beetle behavioural results, in particular results which seem to demonstrate a preference ordering of cues.
2. Develop a CX model, as much as possible in the style of [2], that:
   (a) encodes polarisation and intensity-gradient cues in a neurologically plausible manner;
   (b) integrates these cues in a neurologically plausible manner, taking into account the hypothesis developed for aim 1; and
   (c) produces, in a neurologically plausible manner, control signals that if followed will result in menotaxis behaviour comparable to the behaviour of dung beetles.
3. Implement this CX model on the biorobotic platform developed during the first year of the project.
4. Test and compare the menotaxis abilities of the biorobot (using Model A and Model CX) with actual dung beetle experimental results. This will require (most preferred option first):
   (a) travelling with the robot to a location where dung beetles can be sourced in order to perform robot behavioural experiments directly alongside dung beetle behavioural experiments;
   (b) importing or otherwise sourcing dung beetles in order to perform robot behavioural experiments directly alongside dung beetle behavioural experiments; or

   (c) designing and performing experiments such that results can be directly compared to those in the literature.

5. Discuss the implications of model building and the experimental results on hypothesised answers to the three research questions.

Each of these aims were met successfully at different points throughout the project. Briefly described here is the general approach to meeting each aim:

1. Section 3.1 develops a detailed explanation of cue combination, via usefulness weighting and reliability weighting, that accounts for all behavioural results considered, including results from [3], [14] and [15].

2. Section 3.2 presents the CX model of dung beetle menotaxis developed during this year of the project, including:

   (a) in Section 3.2.1, a framework for the neural encoding of cues, based on neural encodings from [2], that also incorporates the idea of usefulness weighting developed for aim 1; in Section 3.2.2 an exact specification of the assumed neural encodings of the intensity-gradient and polarisation cues; and

   (b) and (c) in Sections 3.2.3 and 3.2.4, mechanisms to (i) extend the range of the polarisation cue to match the range of the intensity-gradient cue, (ii) integrate cues, and (iii) produce menotaxis control signals, using plausible neural operations throughout, such as those utilised within [2].

3. Chapter 3, in particular Sections 3.2 and 3.6, describes the implementation of the model and how it fits within the existing biorobotic platform.

4. (a) Chapter 4 describes the field trip to Johannesburg, where directly comparable dung beetle and robot (using Model A and Model CX) behavioural experiments were performed under South African skies.

5. Chapter 5 discusses the implications of model building and experimental results on answers to the three research questions (Sections 5.1, 5.2 and 5.3), and concludes with some possible answers to aspects of these questions (Section 5.5).

1.3 Contributions

The main contributions of this project are summarised here:

- A hypothesis for cue combination during dung beetle celestial menotaxis that successfully explains all known relevant dung beetle behavioural experimental results.
- The extension of a neurologically plausible CX model of path integration in bees to account for celestial menotaxis in non-homing ball-rolling dung beetles, including:

  - an interpretation of the neural encoding and mechanisms in the context of the above-mentioned hypothesis;
  - an extension of the range of the polarisation cue;
  - suggested metrics for the weighting of cues;
  - a mechanism for taking a snapshot; and
  - a mechanism for combining snapshot errors for the purposes of menotaxis
- The implementation of this CX model on an existing biorobot, including a user interface, executable scripts and behavioural routines appropriate to carry out behavioural experiments on the robot under high-pressure time-limited situations.
- The implementation of a control mechanism more closely matching realistic dung beetle paths.
- A demonstration that the implemented model and system successfully produces beetle-like menotaxis under certain South African sky conditions.
- A novel mechanism for encoding the intensity-gradient cue (the RV cue), implemented within the context of the CX model, including including a demonstration that this cue can be used successfully for successful menotaxis under certain South African sky conditions.
- A novel interpretation of existing dung beetle behavioural results with respect to how dung beetles apparently ignore terrestrial landmarks.

1.4 Overview

The remainder of this report is structured as follows: Chapter 2 introduces key background information and concepts necessary to understand and appreciate the contributions and results of this year of the project, including dung beetles and their celestial environment and behaviour, the central complex (CX) in general, and a CX model for path integration in bees that forms the basis of the work carried out during this year of the project; Chapter 3 presents the actual work carried out, including the development of a neurologically plausible CX model of dung beetle menotaxis (Model CX) and its implementation on the biorobot; Chapter 4 reports the scientific results of experiments performed during a field trip to Johannesburg, including experiments comparing the behaviour of dung beetles and the behaviour produced by Model CX; and finally Chapter 5 discusses the implications of the methods and results on the three research questions, briefly evaluates the limitations of the project, and draws overall conclusions in relation to the research questions.

Chapter 2

Background

This chapter introduces key background information and concepts necessary to understand and appreciate the contributions and results of this year of the project. Summarising and updating the background chapter of the first year of the project [1, Chapter 2], Section 2.1 introduces dung beetles, their celestial environment, and behavioural experiments concerning their navigational abilities; Section 2.2 introduces the neurological structure known as the central complex (CX) and explains its navigational function for insects and in particular for dung beetles; and Section 2.3 briefly describes the CX model for path integration in bees [2] on which the CX model for menotaxis in dung beetles is based (presented in Chapter 3).

2.1 Dung beetle celestial menotaxis

Figure 2.1: A South African non-homing ball-rolling diurnal dung beetle

Diurnal (day-active), crepuscular (twilight-active) and nocturnal (night-active) non-homing ball-rolling dung beetles exhibit celestial menotaxis behaviour – the ability to use visual cues from the sky to maintain a straight-line bearing. The natural behaviour of these beetles is to first locate a dung source and form a dung ball. Then, before rolling their ball, or if separated from the dung ball or otherwise disturbed during ball rolling [16], the beetles perform a 'dance' [17], whereby they climb onto and rotate themselves on top of the dung ball. Beetles then roll their ball in a (seemingly randomly [18]) chosen straight-line direction (menotaxis) directly away from the dung source and the competition of potentially kleptoparsitic rivals [19].

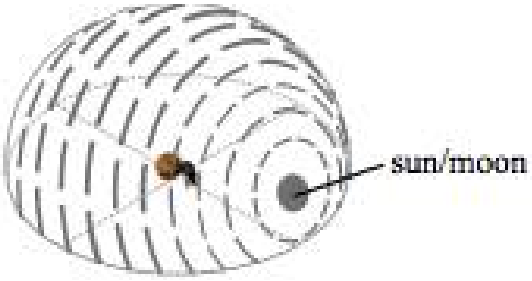

Figure 2.2: Schematic indicating how the angle of polarisation depends on the altitude of the sun/moon according to the Rayleigh sky model; modified from [15, Figure 1]

The sky presents a large variety of directional information in the form of the intensity pattern, and also the polarisation pattern, which is available at varying degrees under most conditions [20][21][22][23][24] except a moonless night. Intensity and polarisation information originates from and changes with the position of celestial bodies, particularly the sun and the moon. For both cues, the quality of information is worse when the sun or moon are at higher altitudes e.g. the zenith (directly above the animal) compared to lower altitudes, and quality is also affected by overall light levels and cloud cover. Physiological evidence suggests that dung beetles are suitably equipped and well-adapted to capture this information. In particular, nocturnal dung beetles have incredibly sensitive eyes [25], and are able to reliably utilise the moon [26], the polarisation cue at night [27], and even the intensity pattern of the Milky Way galaxy [28].

Over the course of many investigations for over a decade, a paradigm for behavioural experiments involving ball-rolling menotaxis has been developed and adopted by biologists at Lund University. A willing dung beetle is first placed in the center of a circular arena along with its precarved dung ball. Then if the beetle successfully proceeds to roll the dung ball to the edge of the arena, its exit bearing is recorded (usually to within 5°). This process is then repeated with the same beetle/ball and the exit bearing again recorded. The difference between the two exit bearings for a single beetle (denoted CIB for change in bearing) then indicates if the same bearing was maintained over successive rolls (menotaxis). Due to the unpredictable and noisy nature of behavioural experiments, this is normally repeated a number of times for a number of beetle/ball pairs, with the distribution of CIB and (circular) mean CIB reported.

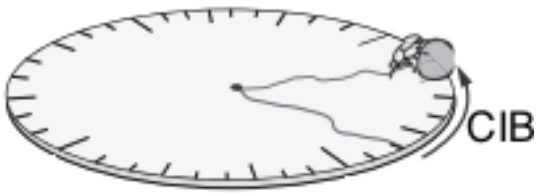

Figure 2.3: Schematic depicting the dung beetle behavioural experiment paradigm, and how a CIB result is calculated experimentally; modified from [14, Figure 1]

By carefully manipulating the environmental conditions (which visual cues are available, celestial or artificial), this paradigm has been successful in demonstrating a number of results. Dung beetle behavioural experiments have confirmed beyond doubt that in order to menotax, they integrate information from the azimuth of celestial bodies [26][27][14], the direction of the skylight intensity gradient [9][29], the angle of linear polarisation at the zenith [25], the intensity gradient of the Milky Way galaxy [28][29], as well as spectral information [30] and in more recent experiments even wind direction (unpublished data collected by Adrian Bell of Lund University). A dance-based snapshot mechanism [3], whereby a snapshot of available cues is taken during the initial dance and subsequently compared against the

currently available information, is the leading candidate mechanism by which menotaxis is achieved using this integrated visual information. An apparently innate preference ordering of cues (denoted by $\succ$) has been identified, whereby point-source and polarisation cues $\succ$ the spectrum, weak intensity-gradient and (probably) wind cues, however cue preferences also appear to depend on the empirically perceived accuracy of each cue [14] (discussed further in Section 3.1).

2.2 The central complex (CX)

The name 'central complex' (CX) refers to a group of interconnected neuropils (dense networks of interwoven neuronal branches and synapses), the major components of which are found in the brains of all known insects (some CX structures are even found in certain crustaceans and other arthropods [31]). In fact, the similarity between the CX for different insects and other animals has been shown to be remarkably high, even down to the level of individual neurons [32]. For this reason, the CX is said to be highly 'conserved'. This means that natural selection has seemingly maintained (conserved) this brain region over millions of years of evolution, indicating the importance of the CX and its functions for survival. Matching the importance of the the CX, its structures are located centrally in the brain of insects, with input and output connections to both hemispheres. Figure 2.4 shows the main neuropils of the CX in a diurnal dung beetle, named the PB (protocerebral bridge), CBU (upper division of the central body, aka the fan-shaped body), the CBL (lower division of the central body, aka the ellipsoid body), and the NO (two paired noduli).

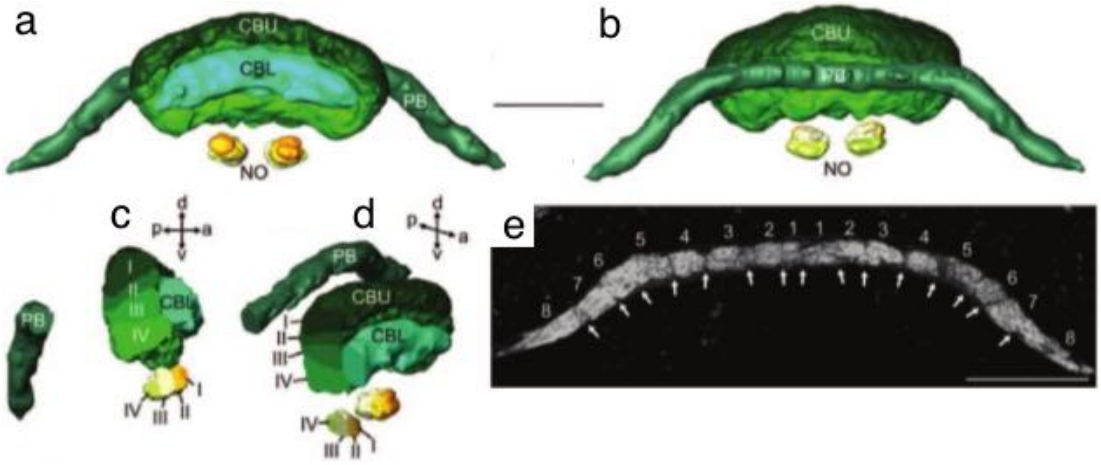

Figure 2.4: (a,b,c,d) 3D model of central complex (CX) neuropils in the diurnal dung beetle and (e) frontal optical slice showing the neuropils of the protocerebral bridge (PB); anterior (a), posterior (b) and side slice (c,d) views; CBU = upper division of the central body, CBL = lower division of the central body, NO = noduli; a = anterior (front-facing), p = posterior (rear-facing), d = dorsal (back-facing), v = ventral (belly-facing); scale bar 100μm; modified from [33, Figure 6].

The CX has been linked to a range of navigational abilities in insects, including sun-compass navigation in Monarch butterflies and locusts [34], obstacle navigation in cockroaches [35], path integration in desert ants and honey bees [2], and of course celestial orientation in dung beetles [15]. The neurons of the CX of different insects have been shown to be tuned to visual (both celestial and terrestrial), proprioceptive and mechanical (puffs of air to the head) sensory signals [32], and to be predictive of motor outputs for walking and flying animals [32][35]. The CX has also been identified to be crucial to successful spatial memory in fruit flies [36], where neurons act as a compass encoding the current bearing of the fly as it rotates.

All this and other evidence from converging lines of research point to the CX as the location in the brain where sensory signals and past experience are integrated into a current state of the world (usually in the form of a 'compass' of sorts) whereby motor outputs can be determined as required by the state and current goal of the insect. In other words, the CX can be thought of as the location of sensorimotor control – where output motor signals are determined from input sensory signals.

In dung beetles, neurons in the CX have been shown to be sensitive to, and encode, both polarisation and sun/moon azimuth celestial cues [15]. In a recent anatomical study [33], there was no evidence to suggest the CX in dung beetles has very different functions compared to the CX in other insects, and so in dung beetles this brain region remains the most likely location for the celestial compass and for the mechanism responsible for menotaxis.

2.3 The CX path integration model

This section briefly describes the CX model for visual path integration in bees [2] on which the CX model for visual menotaxis in dung beetles (presented in Section 3.2) is based. Path integration is a mechanism for homing, a different task from menotaxis.

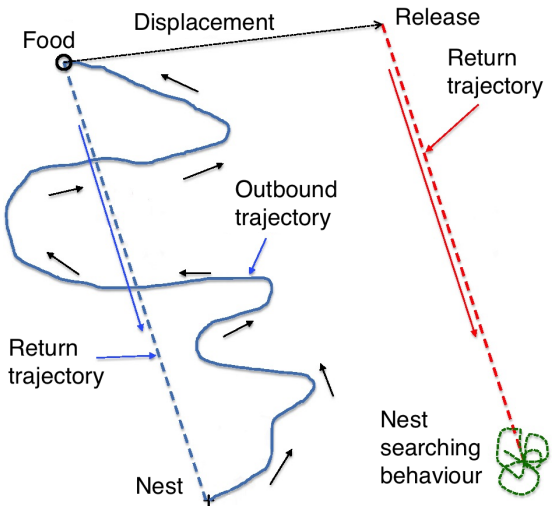

Figure 2.5: Schematic showing how path integration can be used for homing; modified from [37, Figure 1]

Homing is the the task of taking the most direct route back to some home location (e.g. the nest) after taking some unpredictable route away from home (e.g. to find food). Path integration is the strategy of continuously calculating the return trajectory (in the form of a 'home vector') by summing over the vectors of movement (directions and distance traveled) during the outbound path. This can be visualised in Figure 2.5 which shows an example experiment performed on *Cataglyphis fortis* desert ants demonstrating that the path integration strategy is used. If the nest is not located after following the home vector, 'nest searching behaviour' occurs – a strategy of circling around the estimated position of the nest. In *Megalopta genalis* honeybees the home vector has been shown to be calculated based on visual cues only [37]. Direction of movement is estimated using a celestial compass including polarised light as in dung beetles. Distance traveled in a particular direction is estimated using optical flow ("the pattern of apparent motion of objects, surfaces, and edges in a visual scene caused by the relative motion between an observer and a scene" [38]).

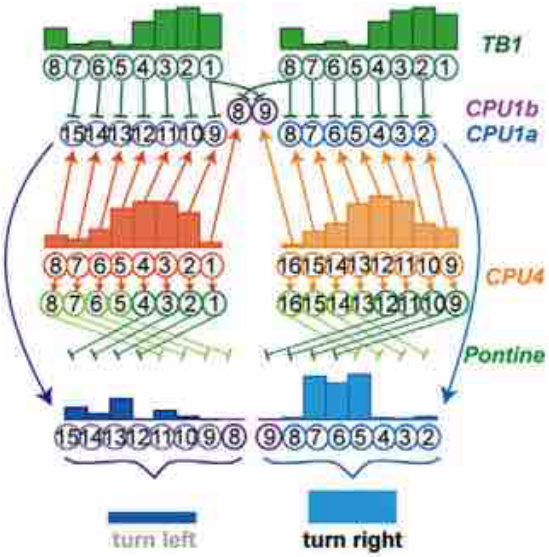

Figure 2.6: CX path integration model – example of steering mechanism in terms of neural activity values – the neural activity (firing rate) of each cell is indicated by the height of the bars; 8 TB1 cells encode current bearing (green cells); 16 CPU4 cells encode desired bearing shifted 45° to the right (orange cells) and 45 ° to the left (yellow cells); 16 CPU1 cells combine TB1 and CPU4 cells to produce a decision to turn right or left; modified from [2, Figure 5I]

In the model [2] it is assumed that directional information from the celestial compass is encoded in the neural activity (firing rate) of eight 'TB1' ('compass') neurons (cells) in the PB. It is hypothesised that the complex inhibitory connections between TB1 cells form a 'ring attractor' – the neurons are connected such that their complex dynamics consistently stabilise to patterns of activity in 8-dimensional space that correspond to points around a ring i.e. bearings. Specifically, it is hypothesised that the eight cells stabilise to a sinusoidal pattern of activity, where the encoded direction corresponds to the peak of the sinusoid. As the underlying sinusoidal function is represented only by these eight values, the peak of the sinusoid is not explicitly represented – the eight cells directly correspond to bearings 45° apart, but together can encode the full range [0°,360°). An example of the sinusoidal bearing encoding is shown in Figure 2.6 (green cells), where the peak of the sinusoid is somewhere between the second and third TB1 cell. If we assume that a peak at the third TB1 cell corresponds to an encoding of 90° and a peak at the second TB1 cell corresponds to an encoding of 135° then the TB1 values shown encode a bearing of somewhere between 90° and 135°. All bearings in the model are encoded in this way, referred to as the neural bearing-encoding. Further mathematical details and examples of this encoding are given in Section 3.2.1.

In the model [2] directional information encoded by TB1 cells is continuously combined with optical flow information from 'TN' ('speed') cells in the NO (TN cells together encode the instantaneous speed of travel in the TB1 direction, calculated using optical flow) to produce home vector updates. The home vector is encoded in 'CPU4' ('memory') cells which are also hypothesised to be recurrently connected, resulting in a simple form of memory. In the model, the home vector updates are continuously summed with the CPU4 cells such that the CPU4 cells encode the home vector at all times. The details of the combination of TB1 and TN cells (as well as the normalisation effect of Pontine cells) are not relevant to dung beetle menotaxis.

More relevant to dung beetle menotaxis is the general purpose mechanism proposed by [2] of how current and desired bearings are compared in the CX to output motor signals which will steer the animal towards its desired bearing. This is accomplished by the steering network shown in generalised form in Figure 2.7 and included in full form in Figure 2.9. The 16 CPU4 cells effectively encode the home vector shifted by 45° to the left (anticlockwise, denoted to be negative) and 45° to the right (clockwise, denoted to be positive). Note that in the actual model [2] and in Figures 2.6 and 2.9 the CPU4 cells encode ±90° of the home vector, related to the different expansion points of optical flow encoded by the TN cells in each hemisphere, but the details of this are not relevant to dung beetle menotaxis – the comparison of current to desired bearing only requires a CPU4 shift of ±45°.

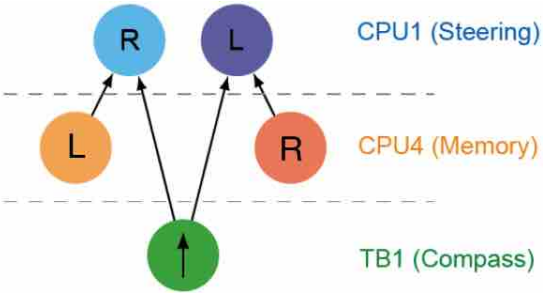

Figure 2.7: CX path integration model – network depicting the steering mechanism in terms of neuronal connections – the network shows that the current bearing (8 TB1 cells, green) is combined with the desired bearing (16 CPU4 cells, yellow and orange) to produce motor signals (encoded in 16 CPU1 cells, blue and purple); modified from [2, Figure S6B]

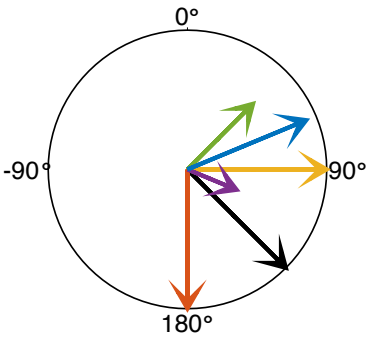

Figure 2.8: CX path integration model – example of steering mechanism in terms of encoded direction vectors – explained in the text

Figure 2.8 shows an example in terms of direction vectors of how the current bearing (TB1 green arrow) and home vector (black arrow) shifted 45° to the left (CPU4-L yellow arrow) and to the right (CPU4-R orange arrow) are combined to produce values for the 16 'CPU1' ('steering') cells in the CBU (CPU1-R blue arrow and CPU1-L purple arrow). Summing sinusoidally encoded cells has the same effect as taking the circular mean vector, meaning CPU1-R encodes the circular mean vector of TB1 and CPU4-L and CPU1-L encodes the circular mean vector of TB1 and CPU4-R. As the magnitude of the vector encodes the spread of bearings included in the sum, a CPU1-R vector longer than a CPU1-L vector indicates that the TB1 vector is closer to the CPU4-L vector than the CPU-R vector, and therefore that a **right** turn will bring the current bearing

closer to the desired bearing. The magnitude of CPU1-R and CPU1-L vectors can be compared simply by comparing the sum of the activity values of the eight cells encoding each vector (depicted in Figure 2.6).

Figure 2.9 shows the CX path integration network in full. In summary, the green TB1 cells encode the current bearing. This information is combined with TN cells to produce the home vector, encoded 45° to the left in the yellow CPU4-L cells and 45° to the right in the orange CPU4-R cells. The TB1 current bearing is then compared to the CPU4 home vector such that if the sum of all CPU1-L cells is greater than the sum of all CPU1-R cells, the home vector will be matched by rotating to the left, and vice versa. The anatomical difference highlighted between CPU1a and CPU1b cells in Figures 2.6 and 2.9 is not relevant to understanding the steering mechanism. This concludes the presentation of the CX model for path integration in bees – most of Chapter 3, in particular Section 3.2, concerns how this model was modified and included within the robotic system to produce experimentally testable menotaxis behaviour. A particular omission of [2] is any suggestion of how the 180° range of the polarisation cue might be extended to form the 360° range of the TB1 bearing-encoding. A mechanism for this is proposed in Section 3.2.3.

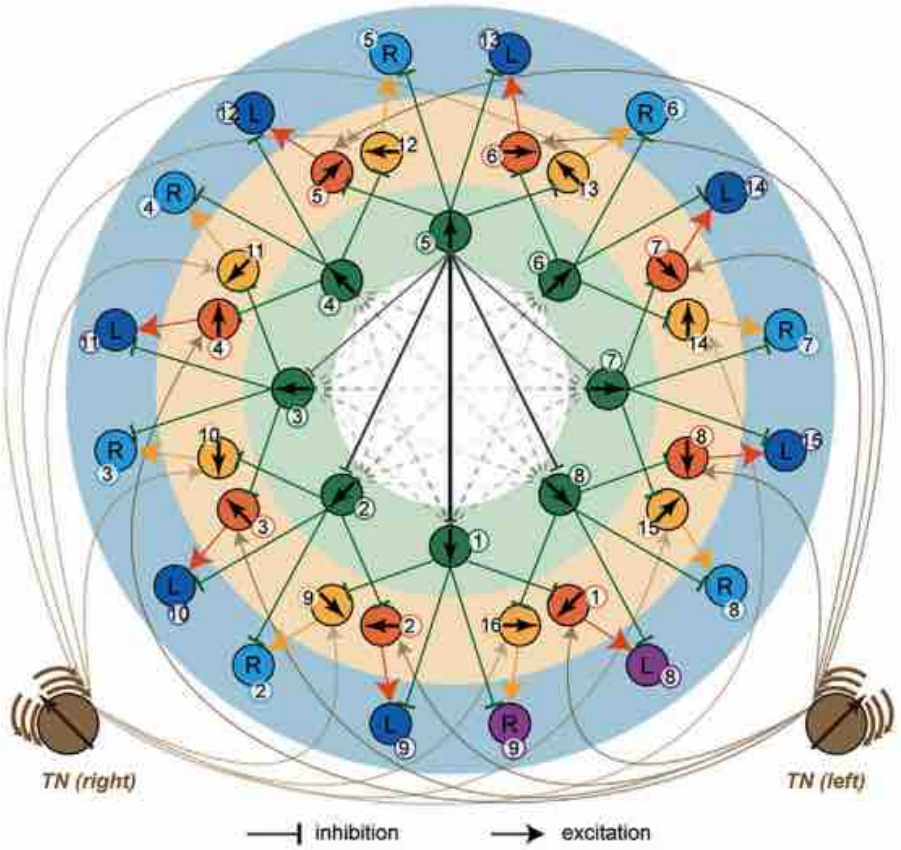

Figure 2.9: CX path integration model – full anatomically constrained network, showing TB1 compass cells (green), TN speed cells (brown), CPU4 memory cells (yellow – shifted to the left by 45° and orange – shifted to the right by 45°) and CPU1 steering cells (blue/purple); modified from [2, Figure 5G]

Chapter 3

Methods

This chapter presents the original work carried out during this year of the project, resulting in a neurologically plausible CX model of dung beetle menotaxis (Model CX) that uses and combines polarisation and intensity-gradient cues. Section 3.1 develops a general hypothesis for the mechanism by which different cues are combined, developing the necessary definitions and concepts to do so, and aiming to explain the results of dung beetle behavioural experiments involving an apparent preference ordering of cues; Section 3.2 presents Model CX in full, using the CX model of path integration described in Section 2.3 as a starting point, and implementing the general hypothesis for cue preference from Section 3.1 in a neurologically plausible way, giving concrete details for how the dung beetle brain might combine polarisation and intensity-gradient cues in particular; the details of the neural encoding of cues naturally gives rise to a novel mechanism for utilising directional information within the intensity-gradient cue, presented in Section 3.3; Section 3.4 describes how the control mechanism for both Model A and Model CX was improved to more realistically reproduce beetle-like paths during behavioural experiments; Section 3.5 presents a minor but theoretically interesting restriction made to the visual field of the robot; and finally Section 3.6 explains how all of the features developed during the first year and this year of the project fit together, giving a high-level view of the biorobotic system as a whole.

3.1 Developing a hypothesis to explain cue preferences

In the first year of the project, 'reliability' of a cue was defined as a measure of belief ($\in [0,1]$) representing how likely a visual cue is to be useful for menotaxis [1, Section 3.4.2]. Here the concepts of usefulness and reliability are more fully defined, and the idea of usefulness predictors introduced, allowing a general hypothesis for cue-combination to emerge that successfully accounts for the results of relevant dung beetle behavioural experiments.

If a cue provides sufficient information to determine the current bearing in some consistent frame of reference (i.e. relative to some consistent reference bearing) to a sufficient level of accuracy (maximal difference between a sensor measurement and the

actual value), then this bearing could be used to menotax (e.g. as the current bearing in Model A). The **usefulness** of a cue can therefore be defined as the accuracy of the bearing measurement using that cue. Whether or not a cue is **available** is then defined as a binary decision – whether or not the cue meets some minimal usefulness threshold.

Empirically measuring the accuracy of bearing measurements would require independent bearing measurements to verify the 'actual' bearing. Odometry was considered for this, but would have required the invasive insertion of proprioceptive sensors, and likely would have had its own accuracy issues without regularly synchronising with an external reference bearing. Overhead camera tracking was also considered, but the camera would have likely obstructed the robot's view, interfering with polarisation and intensity-gradient measurements. Manual measurement was the best option for directly measuring accuracy, but for experimental purposes an automated means was preferred. Instead of accuracy then, usefulness is reported by measuring **empirical granularity** of bearing change, defined as the size of the mean non-zero increment between consecutive bearing measurements during a full rotation (180° for the polarisation cue, 360° for the intensity-gradient cue), assuming a sufficiently high frame rate. The insight here is that accuracy is limited by the resolution of the sensor (the smallest detectable incremental change), and empirical granularity is a more appropriate empirical measure than resolution when the effective resolution is not consistent (e.g. bearing measurement increments using the polarisation cue vary as the robot rotates).

According to the snapshot hypothesis [3], cues not available during the dung beetle dance (when the snapshot takes place) are not used for menotaxis in the immediately following rolls. One hypothesis considered to explain this is that the usefulness of cues is determined during the dance, and that the format the snapshot is stored in encodes a weighting factor for each cue ($\in [0,1]$) that when close to zero can effectively 'switch off' the cue. Measuring the empirical granularity seems a reasonable way to determine these **snapshot usefulness weightings**. Indeed, the dung beetle dance is a perfect mechanism for measuring empirical granularity (or accuracy with respect to odometry) – although the duration of the dance varies considerably [16] and does not normally last as long as a full rotation, granularity (or accuracy with respect to odometry) could just be measured over however long the dance does take place.

A snapshot usefulness weighting does not seem satisfactory to explain further dung beetle behavioural experiments however. From [14] we have the result that when the directional information given by the intensity-gradient cue is altered during ball rolling (the sun was mirrored such that it appears rotated 180° about the dung beetle's field of view), while the directional information given by the polarisation cue stays the same, the effect on the CIB is dependent on the altitude of the sun. If the beetle was following the sun cue only and ignoring the polarisation cue it would turn 180°, whereas if it was following the polarisation cue only and ignoring the sun cue, it would not turn. If the beetle considers both cues equally, it would turn 90°, but the results from [14] show that on average the beetle will turn somewhere between 90° and 180°. Crucially, the effect on CIB, corresponding to the weighting given to the intensity-gradient cue compared to the polarisation cue, was shown to **depend on the altitude of the sun**, corresponding to the relative usefulness of the cues ($128.0 \pm 5.78^\circ$ at high altitudes, $120.6 \pm 6.7^\circ$ at medium altitudes, and $104.2 \pm 7.48^\circ$ at low altitudes).

These results suggest that dung beetles are also applying a usefulness weighting to measurements taken during ball rolling. The specific results above could actually be explained just by the snapshot usefulness weighting (applied to every measurement until the next snapshot), but in general the usefulness of a cue is not guaranteed to stay the same over the course of ball rolling (e.g. the sun could be shaded) and so it seems more likely that a usefulness weighting is determined at regular intervals during ball rolling. The hypothesis suggested here is that every measurement is coupled with a **per-frame usefulness weighting**. Furthermore, if we assume that per-frame weighting occurs during ball rolling, it is reasonable to assume that per-frame weighting is used during the snapshot process as well, and thus from now on we assume that snapshot usefulness weightings are equal to the per-frame usefulness weightings at the time of the snapshot.

Taking the assumption that there is such a per-frame weighting, the question becomes how to weight each measurement. Empirical granularity or accuracy with respect to odometry could not be used per frame, as these can only be calculated over a number of measurements while the beetle is moving. Indeed, accuracy cannot be calculated per frame without impossible knowledge of the true reading. Instead, the hypothesis suggested here is that the beetle relies on **usefulness predictors** which can be measured per frame. A usefulness predictor for a cue is thus any measurement that correlates with the accuracy of the bearing measurement for that cue. Usefulness predictors for the intensity-gradient and polarisation cues are investigated in Section 3.2.2.

With usefulness fully defined, the **reliability weighting** $\in [0,1]$ of a cue is then better defined as how consistently available/useful the cue is within the dung beetle's environment. The reliability of a cue is thus environment-dependent, and is actually best revealed by dung beetle behavioural experiments (the assumption is that natural selection will have preferred the use of more reliable cues even when both cues are equally useful). A further consideration is the limitation in range of the polarisation cue – all bearing measurements are in the twice-repeated range $[0^\circ, 180^\circ)$ compared to the full range $[0^\circ, 360^\circ)$ of the intensity-gradient cue (and most other conceivable cues). The reliability of the polarisation cue should therefore be defined a priori as lower than the intensity-gradient cue assuming both are otherwise equally useful.

Evidence from [15] could support the existence of an evolutionarily determined reliability weighting of cues – preferences are seemingly different for diurnal dung beetles (low-intensity point-source $\succ$ low-intensity polarisation) and nocturnal dung beetles (low-intensity polarisation $\succ$ low-intensity point-source) when cues are seemingly equally useful. This could be explained by a reliability weighting favouring the intensity-gradient cue for diurnal beetles more than for nocturnal beetles. A simpler explanation that does not require an explicit reliability weighting is that the tunings of the usefulness predictors for the cues are different between the species (this could take into account both physiological differences between the eyes of the animal and differences between the reliability of cues at day versus at night). Both explanations are behaviourally equivalent, and so for the sake of clarity an explicit reliability weighting is assumed.

To summarise, the hypothesis presented is that for all cues, every bearing measurement is coupled with a weighting reflecting a cue-specific prediction of the accuracy of that measurement. This weighting is stored in the snapshot, explaining how cues not available during the snapshot can be subsequently ignored. These weightings also explain how intermediate bearings are taken when cues change independently mid-roll, where the relative weighting determines how closely the intermediate bearing matches each cue. Conceptually separate from the measurement weighting, measurements for each cue are also weighted by the reliability of the cue within the dung beetle's environment (tuned by natural selection), which explains why the intensity cue seems to be a priori preferred over the polarisation cue.

Section 3.2 incorporates this hypothesis within the framework of a CX menotaxis model. In particular, a potential neural encoding combining usefulness weightings with bearing measurements is presented in Section 3.2.1 and potential usefulness predictors are investigated in Section 3.2.2. Section 5.1.1 suggests dung beetle behavioural experiments which could test this hypothesis.

3.2 The CX menotaxis model (Model CX)

This section presents the CX menotaxis model (Model CX) developed during this year of the project. Model CX was developed using the CX path integration model [2] described in Section 2.3 as a starting point. In particular, the steering circuit was kept intact, bearings were represented in the same eight-point sinusoidal encoding, and the vector comparison mechanism was reused to extend the range of the polarisation cue.

Starting from the inputs to the model and working towards its outputs, with reference throughout to Figure 3.6 which describes Model CX in full: Section 3.2.1 gives further explanation and examples of the way cues are encoded, including how the novel idea of weighting the cues by usefulness (described in Section 3.1) was included in this encoding; Section 3.2.2 investigates possible usefulness predictors for the polarisation and intensity-gradient cues, ending with a full specification of the inputs to Model CX; Section 3.2.3 describes how the range of the polarisation cue was extended using the intensity-gradient cue; Section 3.2.4 considers two alternative but behaviourally equivalent implementations concerning (i) how the polarisation and intensity-gradient cues might be combined, and (ii) how the memory component of the circuit might be altered to implement the snapshot mechanism such that the model produces menotaxis behaviour instead of homing.

3.2.1 Neural encoding of usefulness-weighted cues

The neural bearing-encoding assumed throughout this project originates from the CX model for path integration [2], first introduced in Section 2.3. The full mathematical framework is given in [2, pp. e9-e10]. The most important details are summarised in this section, and a novel interpretation introduced to explain how a bearing measurement can be coupled with a usefulness weighting within the same encoding.

The main idea is that a bearing vector from the origin to any point $[x,y] \in \mathbb{R}^2$, represented with polar coordinates bearing angle, $\alpha \in [0, 2\pi)$, and magnitude, $A \in \mathbb{R}$, such that $[x,y] = [A\cos(\alpha), A\sin(\alpha)]$, can be encoded as a cosine function of θ:

$$\text{encode}(A,\alpha) = A\cos(\theta+\alpha) \quad . \tag{3.1}$$

So long as the same functional representation is assumed throughout, the full function for a particular parameterisation of A and α (i.e. a particular bearing vector) can be represented in a population of neurons by taking the value of the function at a sufficiently large number of different points (different values of θ), where each value of the function corresponds to the firing rate of a single neuron.

In the present implementation, eight neurons are used to encode a bearing, with θ values $\theta \in \{0, \frac{\pi}{4}, \ldots, \frac{7\pi}{4}\}$, and the cosine wave is also shifted upwards such that the function values are always positive, and scaled by $\frac{1}{2}$:

$$\text{encode}(A,\alpha) = \frac{A \times (1+\cos(\pi+\theta+\alpha))}{2} \quad . \tag{3.2}$$

For clarity, the actual implementation is given in MATLAB code below, where the resulting encoding, `point_values`, is an 8-dimensional vector, each value of which represents the firing rate of a neuron in the population:

```
function point_values = encode(A, alpha)
    thetas = 0:pi/4:7*pi/4;
    point_values = A*(1.0 + cos(pi+thetas+alpha)) / 2.0;
```

A property of this bearing-encoding is that a pointwise sum operation on the firing rates of two neuron populations (encoding two different bearing vectors) will result in a neuron population encoding a bearing vector with a predictable magnitude and bearing angle corresponding to a 'weighted circular mean' of the original bearing vectors. A proof of this, and other mathematical details can be found in [2, pp. e9-e10].

In the CX model, the magnitude of bearing vectors was important to encode distance, such as in the home vector, and in movement vectors. The original contribution of this year of the project was to interpret the magnitudes of bearing vectors as their coupled usefulness weightings. That is, each bearing measurement is encoded as above, with the magnitude corresponding to the usefulness weighting according to some cue-specific usefulness predictor.

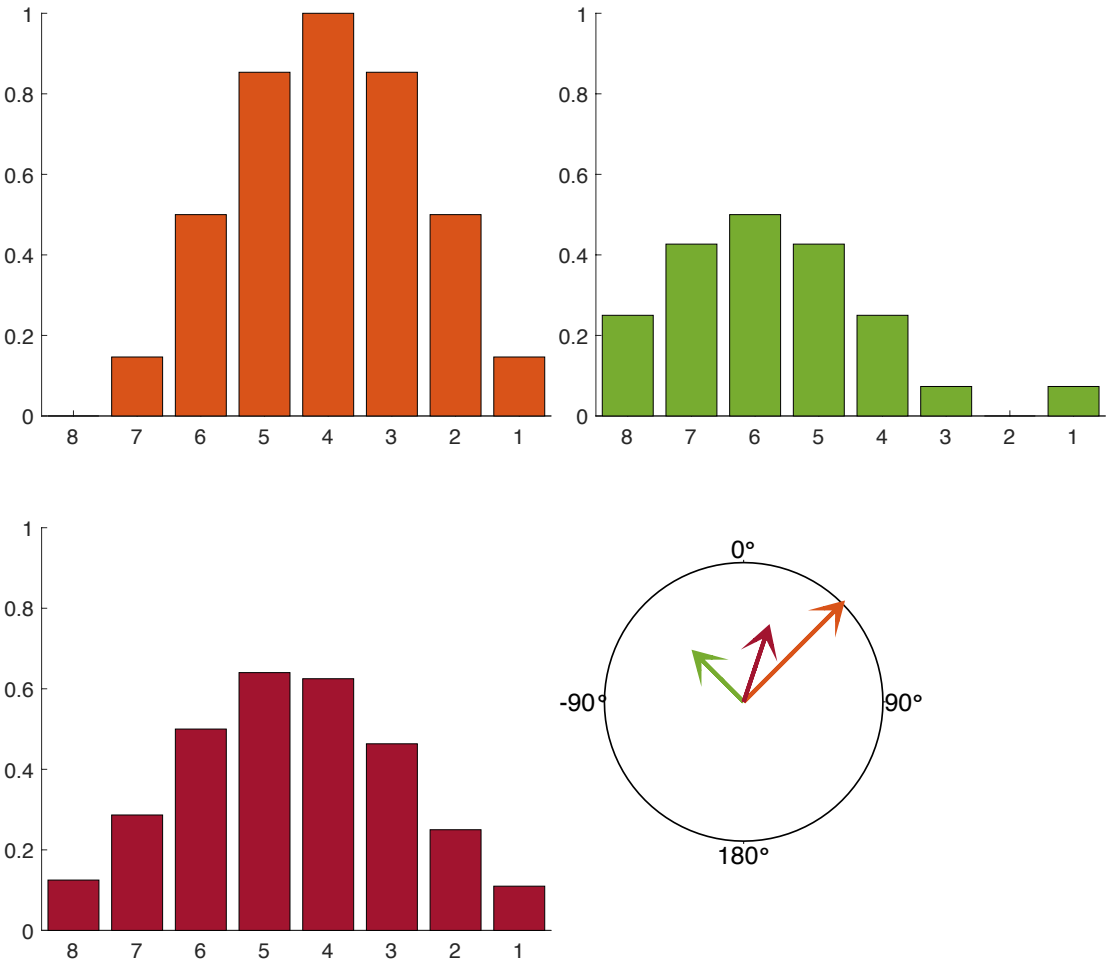

Figure 3.1: Example demonstrating the neural encoding of usefulness-weighted bearing measurements – each of the bar graphs gives the bearing-encoding for the arrow (bearing vector) of the same colour; x-axes denote the neuron index (corresponding to a value of θ), with direction reversed to match Figure 2.6; y-axes denote neuron firing rates as dimensionless real-valued numbers constrained to be in the range of the usefulness weightings; the red bearing-encoding is the pointwise sum of the orange and green bearing-encodings; $A = 1.0$, $\alpha = 45^\circ$ (orange); $A = 0.5$, $\alpha = -45^\circ$ (green); $A = 0.559$, $\alpha = 18.4^\circ$ (red)

Figure 3.1 gives an example to clarify these concepts. Consider two bearing measurements, orange and green, in the same frame of reference, each encoded with different cue-specific usefulness weightings. As can be seen, the usefulness weighting of the orange measurement ($A = 1.0$, $\alpha = 45^\circ$) is greater than the usefulness weighting of the green measurement ($A = 0.5$, $\alpha = -45^\circ$). The desired combination of these cues should therefore be closer to 45° than -45°. And this is indeed the case – when the underlying bearing-encodings are pointwise summed, although the circular mean of the angles alone is 0°, the actual result (corresponding to a 'weighted' circular mean) encodes a bearing vector with a positive bearing.

3.2.2 An investigation concerning usefulness predictors

This section details the usefulness predictors selected for the polarisation and intensity-gradient cues, including experimental justification that these do actually predict usefulness (according to the definitions in Section 3.1).

For the polarisation cue, two different metrics were considered as potential usefulness predictors – degree and intensity of the largest linearly polarised component (denoted p and Ip respectively). Section 1.1.3 explains how these values are empirically determined during a measurement from the polarisation sensor. The difference between these metrics is that p is a proportion (dimensionless) of the overall amount of linearly polarised light that is polarised in the angle Ψ, whereas Ip multiplies this proportion by the total amount of linearly polarised light, I, to give the actual amount of linearly polarised light (in lux) that is polarised in the angle Ψ.

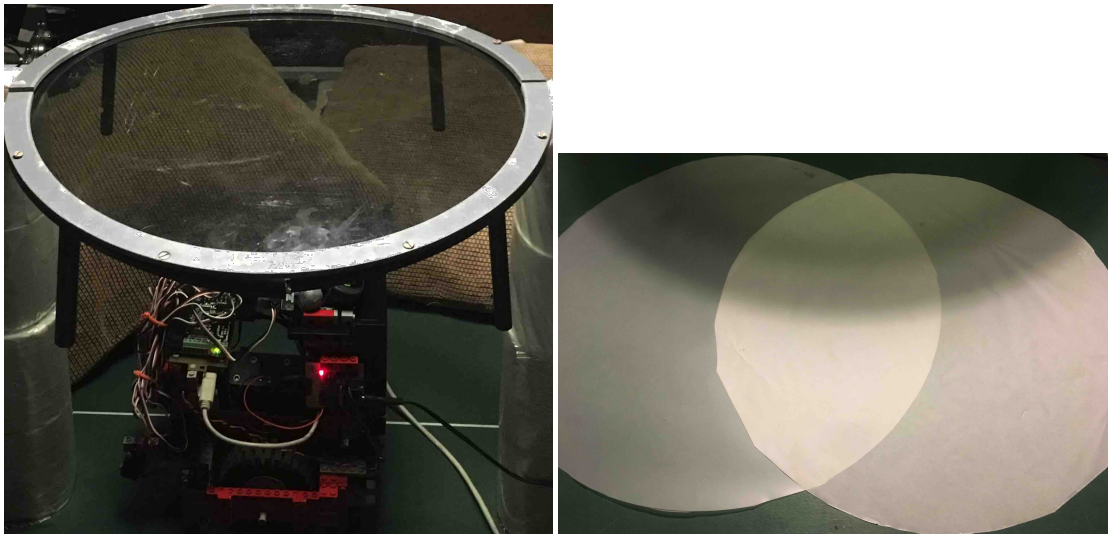

Figure 3.2: Experimental setup for robot experiments investigating possible usefulness predictors for the polarisation cue – setup (left); 'depolarising' filters (right)

In the first year of the project, p was identified as a potential metric for the polarisation cue, as p naturally seems to indicate the quality of the polarisation information available. With the revised definition of usefulness predictors however, it became prudent to experimentally verify that there is actually a correlation between usefulness weightings using this predictor and accuracy of bearing measurements as reported by granularity (mean non-zero bearing change during rotation – see Section 3.1).

Figure 3.2 depicts the experimental setup used during this investigation – the robot was situated under an artificial light source (a ceiling lamp with a halogen bulb) with a linearly polarising filter in between. Paper 'depolarising' filters (which partially elliptically polarise transmitted light, supplied by James Foster of Lund University) were then inserted in different arrangements between the light source and the linearly polarising filter (non-blocking) or between the linearly polarising filter and the robot (blocking).

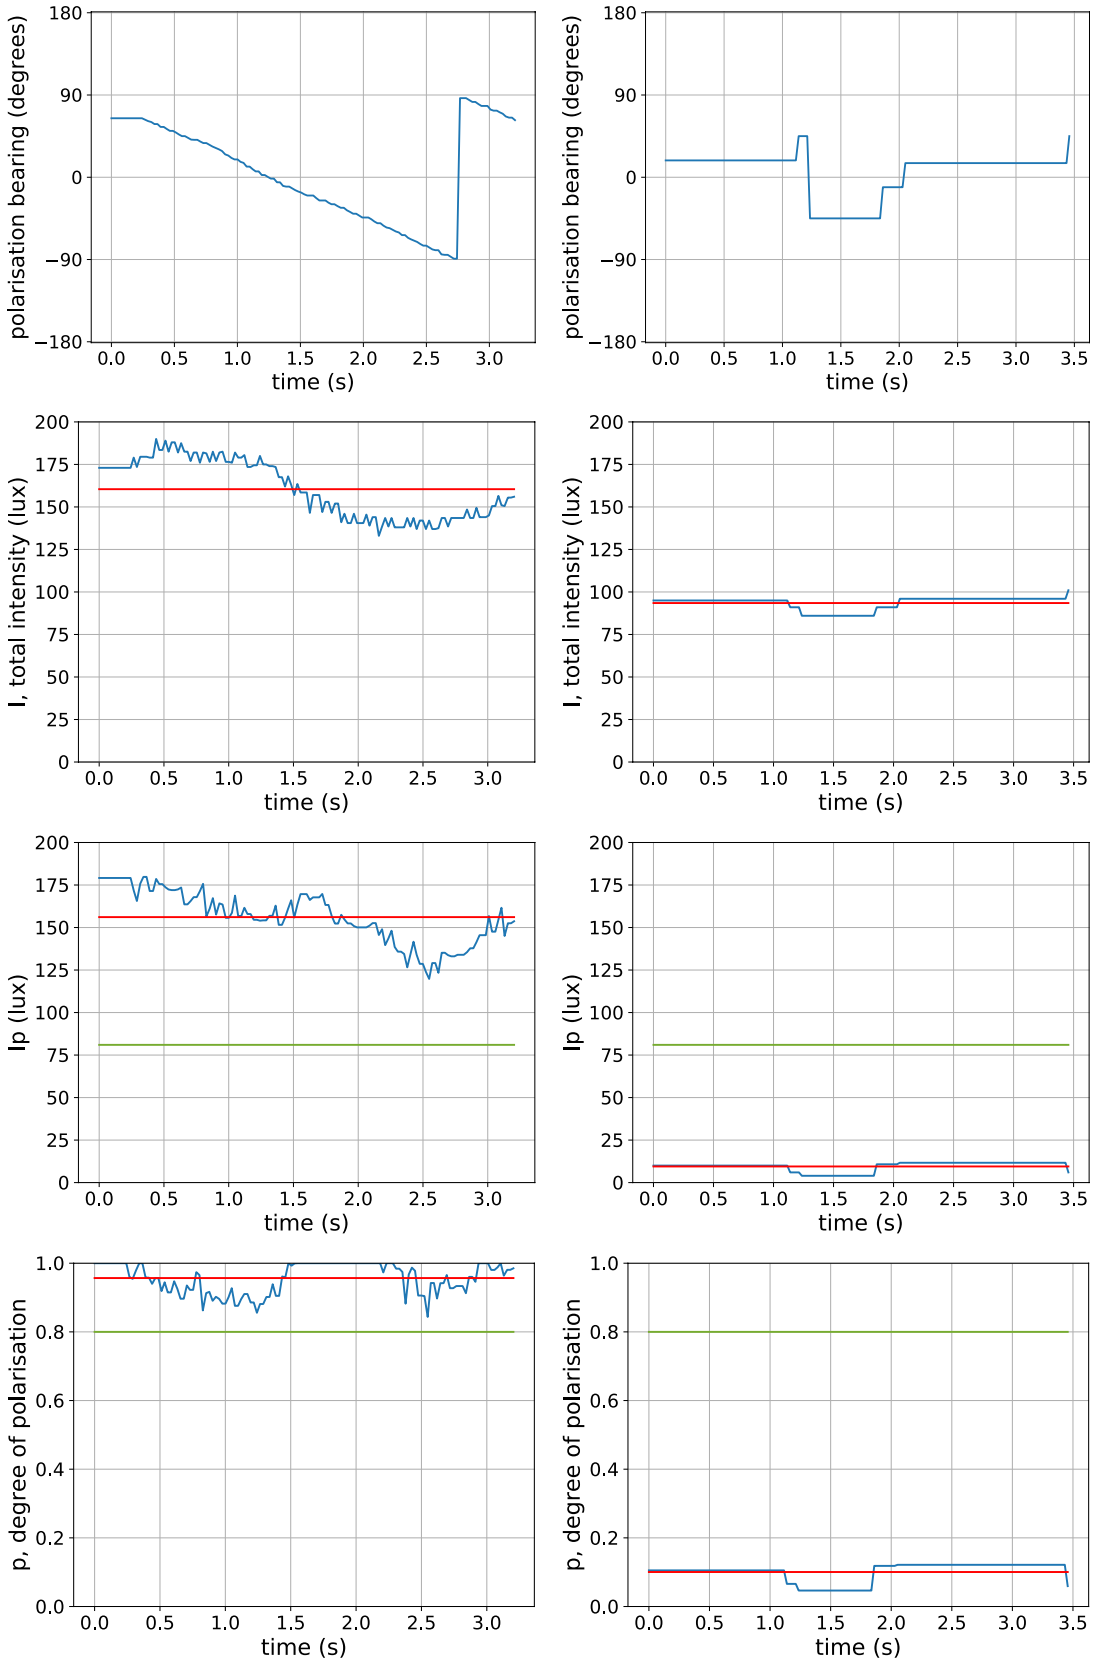

Figure 3.3: Polarisation-related values logged during 180° rotations under different arrangements of depolarising filters – no filters (left); three blocking filters (right); red lines indicate mean values over the whole dance; green lines for the usefulness predictors indicate the chosen availability thresholds

total	blocking	$\mu_{\lvert\delta_+\rvert}(^\circ)$	μ_{Ip} (lux)	$\mu_p \in [0,1]$
0	0	2.09	156.1	0.957
1	0	2.28	149.7	0.973
1	1	2.51	125.6	0.917
2	0	2.63	124.4	0.968
2	1	2.71	111.6	0.906
2	2	18.77	14.6	0.119
3	0	3.25	96.5	0.957
3	1	3.14	93.8	0.894
3	2	23.56	13.0	0.125
3	3	44.26	9.5	0.101
t	0	2.42	116.9	0.969
t	t	*	12.0	0.124

Table 3.1: Numerical results for 180° rotations under different combinations of linearly polarised and 'depolarising' filters – showing the relationship between usefulness, as reported by mean absolute non-zero angle change over the rotation (granularity), $\mu_{\lvert\delta_+\rvert}$, and the mean values of different potential usefulness predictors over the rotation, μ_{Ip} and μ_p, under different combinations of filters; total number of filters = number of blocking filters + number of non-blocking filters, corresponding to the total intensity of light allowed reaching the robot; 't' indicates that a 'totally depolarising' filter was used – one that totally elliptically polarises the transmitted light

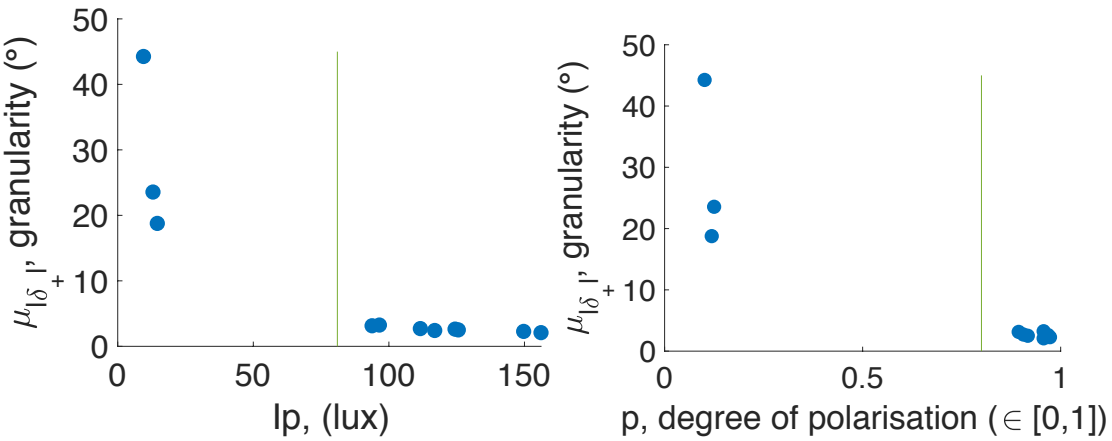

Figure 3.4: Scatter plots of the data in Table 3.1, showing the relationship between usefulness as reported by granularity, and mean values of potential raw usefulness predictors for polarisation; plots produced from data gathered during 180° rotations under different conditions as enumerated in Table 3.1; green lines indicate suggested availability thresholds, a for each raw usefulness predictor; a_{Ip} = 81; $a_p = 0.8$

An experiment consisted of performing a 180° rotation under the filters, logging relevant values Ψ, I, Ip and p. Full results for individual rotations under two different conditions (no filters vs three blocking filters) are given in Figure 3.3. As can be seen, blocking depolarising filters significantly reduce the granularity (looking at the plot of Ψ), as well as Ip and p measurements, but only slightly reduce the overall intensity of light.

Numerical results across all experiments are given in Table 3.1. Results are arranged such that experiments under different total intensities of light are grouped together (described by the total number of filters in between the light source and the robot). Comparing results across each intensity-grouping, it can be seen that the number of blocking filters correlates with both the mean value of the Ip predictor, and the mean value of the p predictor. Figure 3.4 plots this data more clearly, showing the results for all intensity-groupings. There is not a clear linear correlation between granularity and either of the predictions, but for both predictors there is a clear separation between measurements which are useful, and measurements which are not. In other words, suitable thresholding of both usefulness predictors will give good predictions of the availability of the polarisation cue. As such, suitable availability thresholds are suggested, and depicted in Figures 3.3 and 3.4. 'Safer' thresholds were chosen to account for the lack of data in the mid-range of each predictor (a consequence of limited materials available to exactly tune the degree of polarisation presented to the robot).

As there was no clear difference between the performance of Ip and p as usefulness predictors, both could potentially be used in Model CX. However, for theoretical reasons Ip seemed more appropriate. This can be understood by considering a hypothetical example where the intensity of the light source is too low to be well detected, but every single photon is linearly polarised in the same direction – although if p were calculable it would be equal to 1, the near 0 value of I, and therefore near 0 value of Ip, means the angle of polarisation cannot be calculated, making the cue effectively useless. As this hypothetical case is successfully predicted by Ip but not p, ***Ip* was chosen** as the usefulness predictor for the polarisation cue in Model CX.

A similar set of experiments was conceived for the CV intensity-gradient cue (described in Section 1.1.4), with the magnitude of the CV, τ_{CV}, the only potential usefulness predictor identified. Experiments were not carried out as the results seemed obvious – the smaller the value of τ_{CV}, the closer the centroid to the center, and therefore the smaller the number of positions the centroid can take as the robot rotates about the center (due to the finite and uniform resolution of the camera), i.e. the lower the granularity of the CV intensity-cue measurements over the rotation. Therefore τ_{CV} does correlate (but not linearly) with granularity, and so is a valid usefulness predictor. It was also clear that so long as the centroid is sufficiently far from the center so that the image is not near-uniform, any value of τ_{CV} will be sufficiently useful for menotaxis. An availability threshold was arbitrarily chosen to be 2 pixels, corresponding to a 'granularity' of 16 possible positions of the centroid if the magnitude stays constant during the rotation.

Noting from Section 3.2.1 that the range of the usefulness predictor determines the range of cell values ('firing rate'), it seemed appropriate to transform usefulness predictors to be in the range $[0,1]$, allowing 0 to denote minimal firing rate and 1 to denote maximal firing rate of neurons. It also seemed appropriate to mark the midpoint of this range, 0.5, as the 'availability threshold'. Figure 3.5 shows the exact transformation used to map the chosen 'raw' usefulness predictors to usefulness weightings in this range. A steep sigmoidal transformation was used for both as the only relationship very clearly identified for both predictors was binary correlation with granularity (i.e. availability prediction).

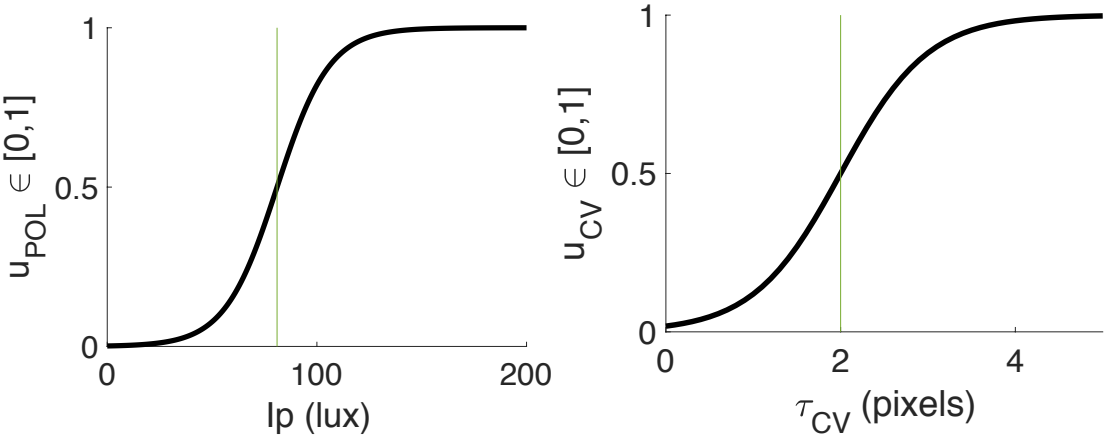

Figure 3.5: Sigmoidal transformation of raw usefulness predictors Ip (left) and τ_{CV} (right) to usefulness weightings u_{POL} and u_{CV} respectively; green lines indicate the chosen availability thresholds $a_{Ip} = 81$ and $a_{\tau_{CV}} = 2$

Together, Section 3.2.1 and this section allow a full specification of the visual inputs to Model CX:

for the polarisation cue: a population of eight neurons, denoted X_{POL}, encoding the angle of linear polarisation $\Psi \in [0^{\circ}, 180^{\circ})$ (also denoted `current_bearing`), weighted by u_{POL}, the sigmoidal transformation of the Ip raw usefulness predictor (also denoted `raw_usefulness`)

for the CV intensity-gradient cue: a population of eight neurons, denoted X_{CV}, encoding the bearing of the CV (also denoted `current_bearing`), weighted by u_{CV}, the sigmoidal transformation of the magnitude τ_{CV} of the CV (also denoted `raw_usefulness`)

The process of how these encodings and weightings are achieved from physical visual input is outside the scope of this project. In the implementation, the neuron values are simply calculated manually, using the empirical calculations described in Sections 1.1.3 and 1.1.4, the sigmoidal transformations depicted in Figure 3.5 and the neuronal encoding described in Section 3.2.1. A convenient consequence of this input format is that for the purposes of behavioural experiments (described in Chapter 4), cues can be easily 'turned off' by setting the weightings to zero for all relevant measurements using that cue.

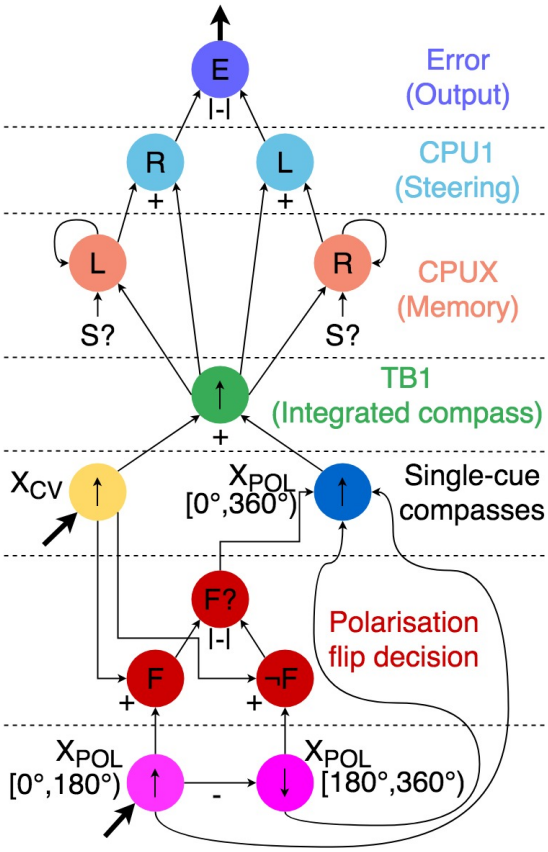

Figure 3.6: CX menotaxis model – network depicting the full model in terms of neuronal connections and operations – inputs $X_{POL} \in [0^{\circ}, 180^{\circ})$ and X_{CV} are defined in Section 3.2.2; the polarisation flip decision is described in Section 3.2.3; the 'cue-integration' and calculation of the final error output is explained in Section 3.2.4

3.2.3 Extending the range of the polarisation cue

The most immediate problem that had to be addressed before successfully combining cues is that the polarisation cue is in the limited range $[0^{\circ}, 180^{\circ})$ (due to the physical nature of linearly polarised light), whereas the required range for successful menotaxis is obviously $[0^{\circ}, 360^{\circ})$ – if using only the twice-repeated range of the polarisation cue, sufficient disturbances (e.g. a fall) could cause the beetle to end up totally reversing its ball rolling direction, bringing it back into the danger of competition.

In order to disambiguate between the two possible opposing bearings indicated by the polarisation cue, it was assumed that available information is utilised from full-ranged cues such as the intensity-gradient cue. The specific implementation proposed is that the original polarisation bearing is shifted by 180° only if the resulting 180°-shifted bearing is closer than the 0°-shifted bearing to the intensity-gradient bearing (in angular distance). This cut off point (the point at which the 0°-shifted and 180°-shifted bearings are equidistant to the intensity-gradient bearing) is appropriate when the assumption is made that both cues are in the same frame of reference (as in the 'cue-integration' mechanism described in Section 3.2.4), and arbitrary otherwise.

A proposed neuronal circuit implementing this transformation can be seen under the 'polarisation flip decision' section of Figure 3.6. This uses the vector comparison method already seen in the steering mechanism (described in Section 2.3), and a multiplexer-style circuit to choose between the two possibilities:

1. Population $X_{POL} \in [0^\circ, 180^\circ)$ is 'flipped' (rotated by 180°) to give $X_{POL} \in [180^\circ, 360^\circ)$ – this could be accomplished by a simple mapping of the eight neurons to shift the values by four ($45^\circ \times 4 = 180^\circ$).
2. Population $X_{POL} \in [0^\circ, 180^\circ)$ is pointwise summed with X_{CV} with the resulting circular weighted mean vector encoded in population F; similarly, population $X_{POL} \in [180^\circ, 360^\circ)$ is pointwise summed with X_{CV} with the resulting circular weighted mean vector encoded in population ¬F.
3. The magnitudes (as indicated by the sum of all neurons in the population) of populations F and ¬F are compared, with the resulting binary signal indicating which of the bearings encoded by $X_{POL} \in [0^\circ, 180^\circ)$ and $X_{POL} \in [180^\circ, 360^\circ)$ are closer to X_{CV}.
4. This binary signal, F?, is used in a multiplexer-style circuit to assign X_{POL} ($\in [0^\circ, 360^\circ)$) to the appropriate value – either 0°-shifted $X_{POL} \in [0^\circ, 180^\circ)$ or 180°-shifted $X_{POL} \in [180^\circ, 360^\circ)$.

3.2.4 Implementing the snapshot mechanism

Once the range of the polarisation cue has been extended, the mechanism for combining cues depends on the assumptions made about the relationship between bearings given by different cues. Behavioural experiments on dung beetles show explicitly that dung beetles do **not** rely on innate knowledge of the natural relationship between the angle of polarisation and the position of celestial bodies (the relationship depicted in Figure 2.2), and can successfully combine polarisation and intensity-gradient cues when they are in 'conflict' according to this natural relationship [3]. The idea then is that menotaxis should be achievable regardless of the angular distance between cues at the time of the snapshot.

For this reason, cues cannot be combined until they are brought into the same 'frame of reference' relevant to the snapshot i.e. they cannot be combined until the angular difference between them at the time of snapshot can be ignored. Two possible implementations were considered, 'cue integration' and 'error integration' (only the cue integration method was actually implemented, and is depicted in Figure 3.6):

Cue integration The 'cue integration' method explicitly brings all cues into the same frame of reference **before** they are combined. This is achieved by rotating the bearings encoded by all cues to match a particular 'reference' cue e.g. the intensity-gradient cue. This transformation is remembered, stored in the snapshot for each cue, and then applied to all subsequent bearing measurements using that cue (until the next snapshot). It should be noted that a neural mechanism to achieve this transformation

could not be conceived. Instead, Figure 3.6 assumes all cues are provided in the same frame of reference. In the actual implementation, the transformations were achieved outwith the neural model. This is a major flaw with this method of cue combination, discussed further in Section 5.2.

Once all cues are in the same frame of reference however, it was easy to combine all cues via a pointwise sum, where cues are weighted in this sum by their species-dependent reliability weighting (see Section 3.1). It was also straightforward to modify the memory component of the circuit to implement the snapshot mechanism – all that was required was summing with 'CPUX' cells that encode the weighted snapshot bearing (instead of 'CPU4' cells that encode the home vector in the CX path integration model). CPUX cells are differentiated from CPU4 cells because it is not clear that the their different functionalities can be switched between via neuromodulation – a question which is outside the scope of this project.

In full, the remainder of Figure 3.6 can be understood by considering the following 'cue integration' implementation:

5. Population TB1 is calculated with the weighted pointwise sum $0.75 \times \mathrm{X_{CV}} + 0.25 \times \mathrm{X_{POL}}$.
6. If the binary input signal S? (also called the `snapshot_trigger`) is high, then the multiplexer-style CPUX memory circuit sets the values of CPUX-L and CPUX-R to be the current TB1 value shifted 45° to the left and to the right respectively i.e. a snapshot is taken.
7. Otherwise, the binary input signal S? is low, and so the multiplexer-style CPUX memory circuit sets the values of CPUX-L and CPUX-R to their previous values (via the hypothetical recurrent connection) i.e. the snapshot is maintained.
8. The steering circuit then works exactly as described in Section 2.3.
9. The output of Model CX is an error signal E (also called `error`) which when positive indicates that a right turn should be taken and when negative indicates that a left turn should be taken.

Note that the origin of the snapshot signal S? was not considered within the scope of this project (in behavioural experiments, this is determined by user input). However, a simple explanation might be that this signal is a simple or weighted sum of multiple different factors that together determine the decision to end a dance.

Error integration The alternative 'error integration' method instead brings cues into the same frame of reference **after** they are combined. This is achieved by storing all weighted cues in separate snapshots (requiring more CPUX-style cells). The steering 'error' signals CPU1-L and CPU1-R for each respective cue can then be combined in a weighted sum before computing the final error signal. The behaviour produced by the cue integration and error integration methods is equivalent (proof omitted). Because of the implementation difficulties noted above, the error integration method was considered more plausible on reflection. The pros and cons for each method are discussed in more detail in Section 5.2.

3.3 The regions vector (RV) intensity-gradient cue

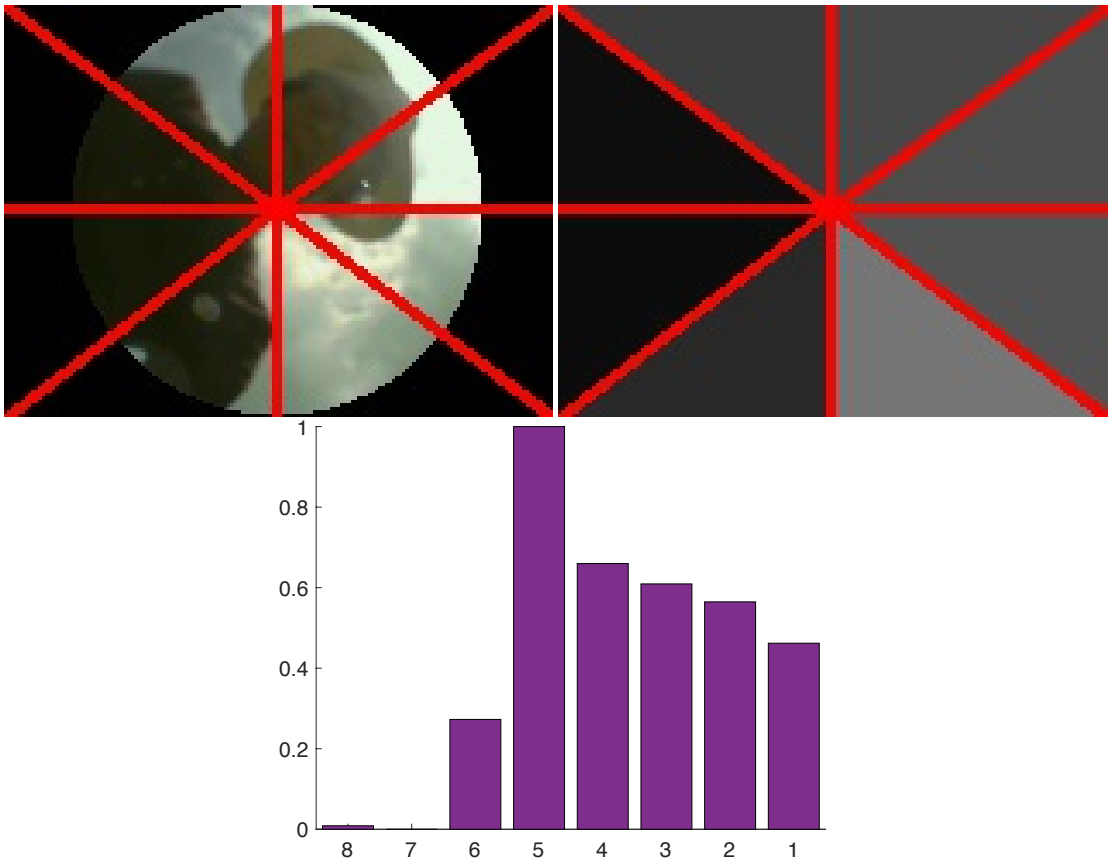

Figure 3.7: Visualisation of the RV cue – for a particular camera frame, regions are indicated by red lines (top left) and then greyscale brightness values summed over each region (top right) resulting in the final RV (bottom, values scaled to be in the range $[0, 1]$)

The eight-valued neural encoding of cues presented in Section 3.2.1 inspired a novel mechanism for utilising directional information within the intensity-gradient cue, called the regions vector (RV) intensity-gradient cue. To calculate this cue for a given camera frame, the visual field is split into eight regions (shown in Figure 3.7), and the total intensity of each region is calculated by summing over every greyscale pixel contained within the region. The RV cue is then represented simply by coupling together these eight intensity values into a single vector, ordered clockwise according to the position of the respective regions (as can be seen in Figure 3.7). In terms of the snapshot mechanism described in Section 3.2.4, it is not clear how exactly this cue could produce menotaxis behaviour. However, this was a simple implementation to attempt, and actually turned out to work in natural conditions (menotaxis using the RV cue is successfully demonstrated in Section 4.4.7). The pros and cons of the CV and RV cues are discussed in Section 5.1.2 with respect to dung beetle behavioural experiments.

A simple potential raw usefulness predictor, referred to as the magnitude of the RV, and denoted τ_{RV}, was devised: taking the difference between the maximal and minimal values of the RV (before the RV is scaled as in Figure 3.7). This potential predictor was not considered in great detail – it is not clear that predictions actually correlate with granularity. As an availability predictor however, this metric seems sound – a uniform image will result in a magnitude approaching zero, appropriately indicating that bearing information is not available from the intensity-cue.

From a neural perspective, the RV cue could also be represented by a population of eight neurons, the activity of each corresponding to an element of the RV. This representation of the RV cue shares some of the properties of the neural bearing-encoding described in Section 3.2.1: eight neurons are required to encode the cue; and rotating the image by 45° corresponds to a shift of cell values by one place. However the crucial ability to meaningfully encode a usefulness weighting could not be conceived of within this framework. Similarly, it was unclear how the bearing indicated by the RV could be combined with bearing-encoded cues using either of the cue integration or error integration methods. Overall then, the RV could not be combined with the polarisation cue within the Model CX framework, and could only be experimentally tested in isolation.

3.4 A more realistic control mechanism

In the first year of the project, the control mechanism `control_straighten` was used, as described in Section 1.1.5. Reviewing the results of the first year of the project, it seemed that the accuracy of menotaxis was particularly high compared to dung beetle results. At the time, this was thought due to the simplicity and deterministic nature of the robot's programming compared to the complexity of dung beetle behaviour (with respect to individual differences and conflicting behavioural priorities).

On reflection however, it was realised that there is also a fundamental difference between the `control_straighten` control mechanism and that used by dung beetles, unrelated to the wheeled vs legged design, presence of the dung ball, or complexity of behaviour. The difference is that in the `control_straighten` control mechanism the robot only ever corrects its bearing by rotating on the spot, whereas the dung beetle will **sometimes** rotate on the spot in the form of performing a 'redance' to match the chosen snapshot, but most of the time will correct its bearing by **turning** i.e. continuing to roll its ball forward but curving its path. The continuous forward movement during this forward-curve will in general result in less accurate menotaxis, simply because corrections are less effective at reducing the error and there is less 'time' to correct when this forward movement constantly brings the beetle closer to exiting the arena. Reflecting this separation between a 'redance' and forward-curve movement, two separate control mechanisms were implemented: `control_straighten` as before, and `control_forward`. The idea is that the robot will perform `control_straighten` at the start of an experiment, mimicking the dung beetle dance, and then will perform only `control_forward`, mimicking ball rolling, for the rest of the experiment.

The `control_forward` control mechanism was as follows (repeated every iteration):

- Inputs: `error`; the current speed s.
- Sample random Gaussian noise to add to the `error` signal: $\texttt{noise} \sim \mathcal{N}(\mu = 0, \sigma = 0.05)$.
- Depending on whether $|\texttt{error} + \texttt{noise}| > 0$ or $|\texttt{error} + \texttt{noise}| < 0$, determine the forward-curve required (clockwise or anticlockwise) to reduce the error signal.
- Set the motor speeds to $2 \times s$ for the outer wheel of the curve and 0% for the inner wheel.

The fact that during `control_forward` the robot is always moving in a forward-curve, and never fully forward, means there is no longer any need for an undesirable threshold as in `control_straighten`. It should be noted however that due to the camera delay, this will result in exaggeratedly sinusoidal paths in the event where perfect menotaxis would otherwise be possible. It is unclear how many 'curvature' settings the beetle takes, or if there is some element of proportional control in the control mechanism used by the dung beetle, but the single curvature setting of `control_forward` was found to be sufficient for menotaxis. The addition of Guassian noise to the error signal was inspired by the CX path integration model [2], and resulted in more realistic behaviour under control conditions (discussed further in Section 5.3).

3.5 Avoiding terrestrial landmarks

Chapter 4 describes robot experiments carried out in order to test the ability of Model A and Model CX to produce menotaxis behaviour using different cues and under different conditions. In the process of performing these experiments, it was always necessary to move out of view of the robot while experiments were underway. For the polarisation cue, it was necessary simply not to block the vertically directed polarised light sensors from their view of the sky.

For the intensity-gradient cues however, it was more important not to accidentally provide the robot with directional information in the form of terrestrial landmarks (e.g. the human experimenter). The idea is that for the intensity-gradient cues, any visible landmark that remains fixed while the robot rotates will skew the directional information available from the sky, in some cases completely overriding it. For example, a large dark terrestrial landmark will likely skew the CV to point away from the landmark, ignoring the intensity-gradient visible from the sky above. This effect can be seen in Figure 3.8. A similar effect will likely occur for the RV, with a large dark terrestrial landmark lowering one or more of the regional mean-intensity values to the point where this terrestrial information overrides the effect of the celestial intensity-gradient cue. It should be noted that when the sun is directly visible, this extremely concentrated area of brightness is likely to influence the CV and RV more than any dark terrestrial cue could.

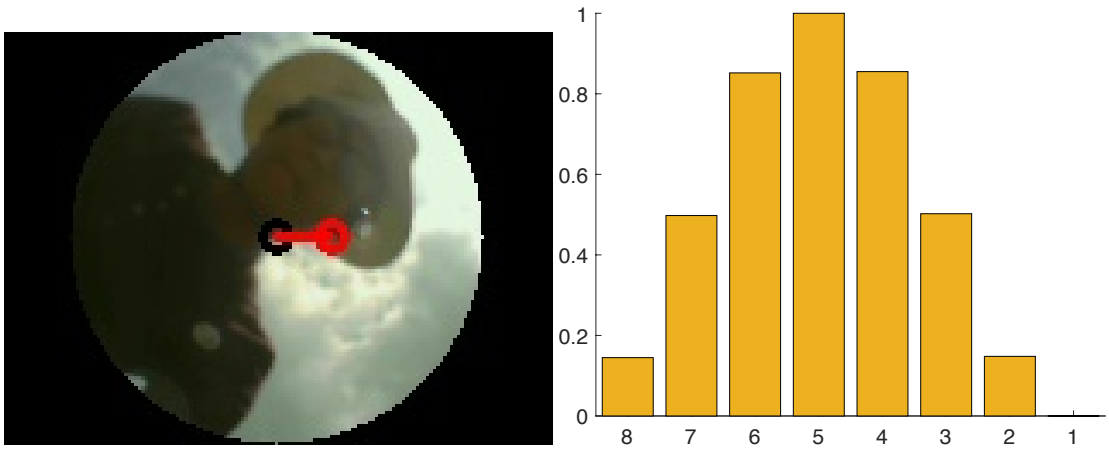

Figure 3.8: Accidental selfie – left: the effect of a human experimenter on the CV cue; the radius of the visual field is set to 100% of the height; black circle = visual center; red circle = centroid of greyscale image; red line = CV vector; the CV vector points directly away from the dark body of the human experimenter; right: 8-point sinusoidal encoding of the CV vector

As dung beetles have not been shown to use terrestrial information explicitly ([39], discussed further in Section 5.1.3), it was deemed necessary to avoid the use of this information where it would otherwise be available. For one attempted experiment, it was not manually possible to avoid the view of surrounding trees, so it was decided to implement a mechanism by which the visual field of the robot can be restricted.

The shape of the restricted visual field was arbitrarily chosen to be circular. The desired functionality from the point of view of the experimenter was to be able to input a radius value (a percentage $\in [1\%, 100\%]$ of the height of the original visual field) which would cause the visual field to become limited to a circle with the given radius and centered on the center of the visual field. This was implemented with the use of a boolean mask checking that that the Euclidean distance from the center to each pixel was less than the radius in pixels. For efficiency reasons, a new mask is created only when the user inputs a new radius value, and the mask is applied to every camera frame by setting the value of every excluded pixel to 0 (black) before the frame is processed further. The implementation was successful, and Figure 3.8 shows the effect on the field of view for a single frame when the radius is set to 100%.

3.6 System architecture

Figure 3.9 gives a high-level conceptual view of the system as a whole, showing how the various components developed this year and last year work together to produce experimentally testable menotaxis behaviour in the robot. This section repeats and updates the description given in the first year of the project, quotation marks indicating direct quotes (section numbers are updated) from [1, Section 3.5].

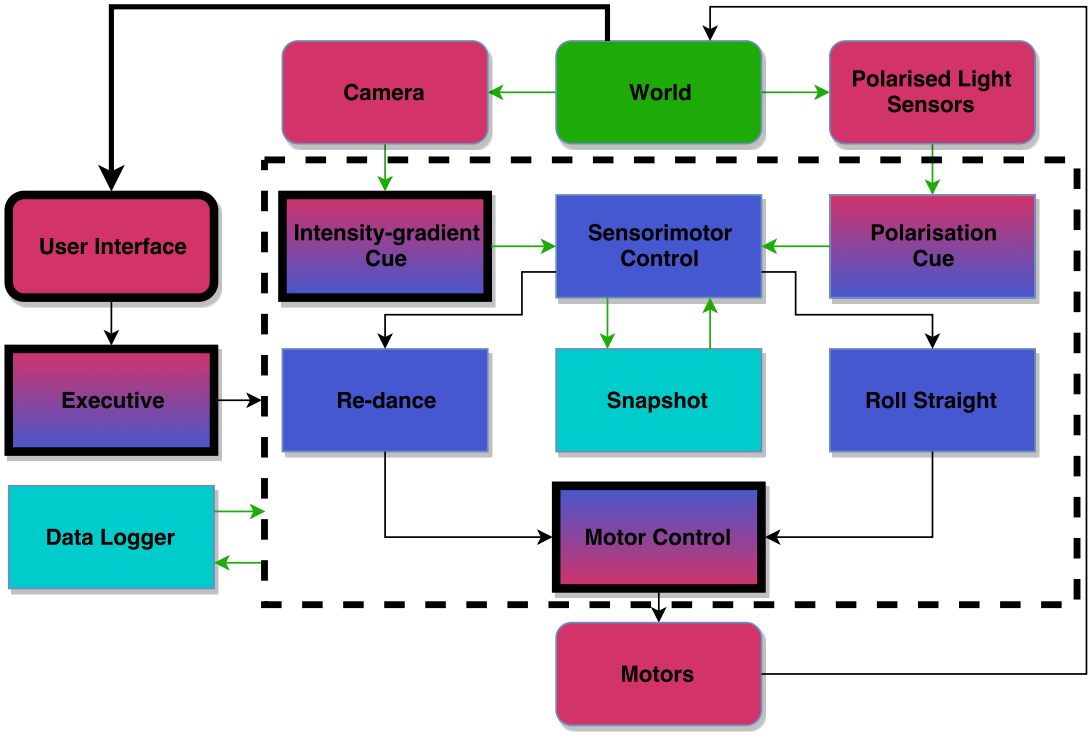

Figure 3.9: System architecture diagram; modified from [1, Figure 3.12]

"Arrows are used to indicate the flow of sensory information (green) from the world into the system, and the flow of control (black), from the user (thick black arrow) to actuators and the world. Rounded boxes (pink) represent hardware interfaces between the world and 'entry-points' to the system (pink-blue):"

- "The motor control block executes pulse-width modulation as described in Section 1.1.2, directly controlling the motor actuators which move the robot, thus changing the state of the world."
- The polarised light sensors capture visual information from the world and feed this into the polarised cue block, which calculates the angle of linear polarisation and the intensity of the linearly polarised component, as described in Section 1.1.3.
- The camera captures visual information from the world and feeds this into the intensity-gradient cue block, which processes the camera image by applying the visual field limitation described in Section 3.5, calculating the CV as described in Section 1.1.4, and calculating the RV and accompanying raw usefulness predictor as described in Section 3.3.
- "The user interface captures user input from the world and feeds this into the executive block (described below)."

"Square boxes represent control blocks (blue) and data blocks (cyan) which store persistent information:"

- The sensorimotor control block takes input `snapshot_trigger` from the executive block and `current_bearing` and `raw_usefulness` from each of the polarisation and intensity-gradient cues, executes either Model A as described in Section 1.1.5 or Model CX as described in Section 3.2, and outputs an `error` signal.
- The redance block takes the `error` signal from the sensorimotor control block, and the current speed s and threshold value from the executive block, executes the `control_straighten` control mechanism as described in Section 3.4, and outputs appropriate motor commands to the motor control block.
- The roll straight block takes the `error` signal from the sensorimotor control block, and the current speed s from the executive block, executes the `control_forward` control mechanism as described in Section 1.1.5, and outputs appropriate motor commands to the motor control block.

"The thick dashed line around these software blocks represents the 'menotaxis' package – a collection of visual processing and behavioural algorithms which can be used for menotaxis. The executive block however, controls the execution of these algorithms, and also has direct control of the motors, and many other parameters, settings and minor functionality (for example, saving camera screen-shots). Executive control can come from either:"

- "user commands sent from a connected laptop to the fit-PC; or"
- "a preprogrammed procedure (making use of the flick switch inputs to the robot) that can be run by double-pressing the power button on the fit-PC."

"The preprogrammed procedure used for experiments is presented in Section 4.3. Similarly to the control block, the data logger can communicate with any block in the menotaxis package, and is used to load and save data to the fit-PC file system in order to e.g. keep settings across multiple executions or save experimental data for analysis."

"Thick lines around the user interface, executive, intensity-gradient cue and motor control blocks indicate that these blocks run on separate threads." Thread timing issues are described fully in [1], with only minor changes implemented for this year of the project.

Chapter 4

Results

This chapter reports the scientific results of the main experiments performed during this year of the project. Results were collected during a field trip to Johannesburg, described in Section 4.1; Sections 4.2 and 4.3 give the details necessary to reproduce the approach to the main experiments; the results including interpretation and analysis are given in Section 4.4. Overall analysis and discussion with respect to the research questions and aims is left until Chapter 5.

4.1 Field trip to Johannesburg

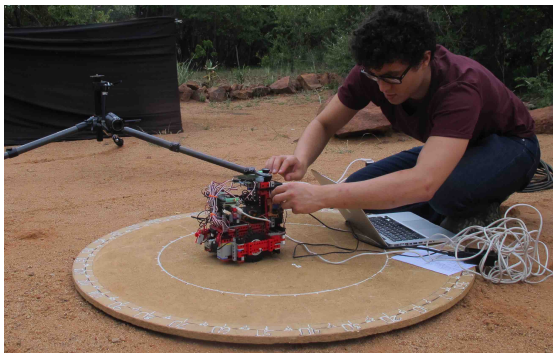

Figure 4.1: Hard at work!

Following aim 4(a), it was preferred to travel with the robot to a location where dung beetles can be sourced. Following up on contact established with Professor Marie Dacke of Lund University during the first year of the project, options considered included visiting Lund University, where dung beetles are routinely imported from near Johannesburg, South Africa, or visiting Johannesburg directly. The advantage of visiting Johannesburg was that experiments could be performed under the natural environment (the South African sky) of the dung beetles involved. The backup option would have been to visit Lund University after fresh dung beetles had been imported, but luckily Professor Dacke and others from the Lund Vision Group were planning a field trip to Johannesburg at a suitable time, and so I was able to accompany them. It should be noted that this trip operated with the help of Professor Marcus Byrne and others of the University of Witwatersrand.

A significant amount of the time allocated to this year of the project was dedicated to organising, planning and participating in this field trip. Most notably, a 1000-word grant application, including a proposed budget, was written for an Innovation Initiative Grant to finance the field trip. The full text of the application is presented in Appendix

B. The application was successful with the full requested amount granted. Flights were researched and booked, and two days were allocated for travel in each direction, leaving seven full days of work. In response to restrictions on the transport of batteries on aircraft, the power source for the motor board was switched from a large 12V battery to a battery pack of 10 rechargeable AA batteries. Additionally, protective packaging for the robot had to be sourced, with damage to the robot the biggest risk of travelling.

The field trip was largely a success, and the robot was robust enough to suffer no damage as the result of travelling. A major unquantifiable benefit of the trip was exposure to discussions, conversations and the fantastic work environment created by Professor Dacke and her team on the field trip. This led to many ideas for how the dung beetle biorobot and CX model could be further developed and utilised to aid research in the ongoing neuroethological investigation of dung beetle navigational behaviours. Unfortunately, due to the placement of the field trip so late in the academic year, many of these ideas could not be implemented, and are left as suggestions for future work in Chapter 5. Some ideas discussed on the field trip were developed to the point of implementation however: the use of Ip as a usefulness predictor for the polarisation cue (Section 3.2.2); and the final approach to the integration of cues (Section 3.2.4).

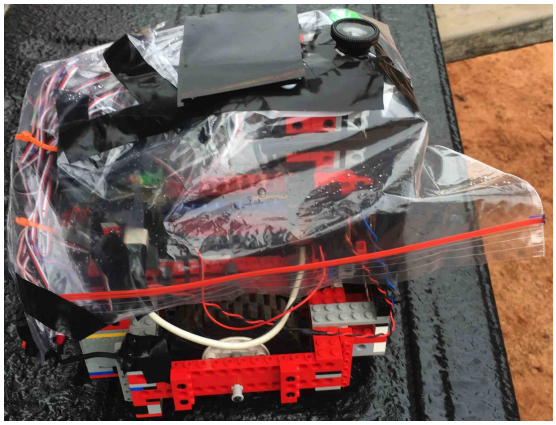

Figure 4.2: Robot doesn't like rain – the robot was successfully made water-resistant to mild rain with a plastic bag and electrical tape.

The amount of data collected on the field trip was quite small compared to what might have been hoped for. This was due to a combination of choosing to implement the new ideas mentioned above which took time to debug, debugging in general due to unexpected problems (notably the need for avoiding landmarks addressed in Section 3.5 and the problems with detecting polarisation discussed in Section 4.4.5), and bad weather – it was unfortunately totally overcast for all but one of the seven days, and rained for the last three. While experiments could be performed in overcast conditions, it was preferred to wait for sunny conditions (particularly for the sake of the dung beetles, which are less amiable to ball rolling under clouds). To allow experiments to be performed in the rain, the robot was made water resistant (see Figure 4.2), but unfortunately experiments in the rain could not be completed due to the risk of water damage to the arenas used.

For these reasons, all experiments reported were actually taken on a single day – the day I flew from Johannesburg. This left limited time and conditions for experiments. The most basic experiments were performed (except using the polarisation cue), but there was not enough time to collect a large amount of data for each experiment, and the more interesting and complex experiments requiring clear skies and specific sun altitudes could not be performed. Furthermore, during the rainy days leading up to 'experiment day', a non-trivial amount of effort was directed towards making the system easily configurable so as to make the process of running experiments as quick and

efficient as possible. While not technically interesting enough to report, a system of scripts was developed to speed up the process of saving and organising results (screenshots and logged values), generating and organising graphs (especially to sanity-check experimental conditions before running full experiments), and setting motor and model parameters. The time taken for a full experiment was successfully reduced from 2+ hours to 30 minutes, with most data and graphs presented in Section 4.4 generated automatically at the time of the experiment. Without these improvements, it would not have been possible to perform all of the reported experiments in the time that was remaining.

4.2 Experimental setup

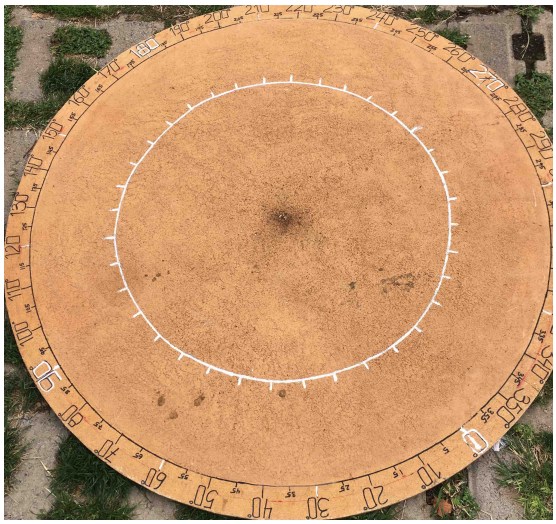

Figure 4.3: Wooden arena used in all full experiments

This section describes the experimental setup for all experiments described in Section 4.4. All experiments were carried out outdoors, approximate coordinates (-26.1907, 28.0313), in the car park for the Biology building in the main campus of the University of Witwatersrand, Johannesburg, South Africa, between 13:30 (solar altitude 67°) and 16:00 (solar altitude 36°) on Saturday 24th February 2018.

Figure 4.3 shows the arena used, which was levelled to provide a completely flat surface before commencing experiments. The radius of the inner white circle is 33cm, and the radius of the outer black circle is 50cm. In order to end robot experiments in roughly the same manner, experiments were halted by running in and manually activating the distance and/or bump sensors of the robot as soon as the front of the robot appeared to cross the outer black circle (before bearing markers). Exit bearings for robot experiments were reported by eye, using the bearing markers ion the arena, to the nearest 10°. Dung beetle experiments were halted once the beetle crossed the inner white circle. Exit bearings for dung beetles were reported by eye, using the bearing markers on the arena, to the nearest 5°. The arena is made of wood, and is slightly rough. It was found that the robot could turn smoothly at speed setting 45%, taking approximately $\texttt{rotation\_time} = 7s$ to rotate the full 360° (making the angular speed $\approx 51^\circ/s$). With all terrestrial landmarks safely out of the sight, the radius of the visual field was set to 100%.

4.3 Experimental procedure

This section describes the experimental procedure followed for all experiments described in Section 4.4. Following the paradigm for dung beetle behavioural experiments described in Section 2.1, two behavioural algorithms, `choose` and `follow`, were designed and programmed onto the robot. Note that the dung beetle dance is outside the scope of this year of the project, and so the choose and follow algorithms in this year of the project no longer involve the `dance` procedure used in the first year of the project. Instead, to allow a random bearing to be chosen uniformly from $[0^\circ, 360^\circ)$, the time taken for the robot to rotate 360° must first be recorded empirically, denoted the `rotation_time` (in milliseconds).

Behaviour `choose`:

1. Rotate in a randomly chosen direction (equal probability of rotating clockwise or anticlockwise) for a random number of milliseconds, chosen uniformly from $[0, \texttt{rotation\_time})$.
2. Fully halt (pause for a second to be sure) and then set `snapshot_trigger` to be true for one iteration of the selected model.
3. Using this snapshot and the selected model, attempt to go forward in a straight line by executing `control_forward` until the arena boundary is detected (via bump or IR sensors), then halt.

Behaviour `follow`:

1. Load the snapshot saved from the most recent execution of `choose`.
2. Using this snapshot and the selected model, attempt to rotate to match the snapshot angle by executing `control_straighten` until the error signal is less than some generous threshold value (chosen to be 10 for each model).
3. Using the snapshot and the selected model, attempt to go forward in a straight line by executing `control_forward` until the arena boundary is detected (via bump or IR sensors), then halt.

To record CIB results, the following procedure was adopted, with all exit bearings recorded to the nearest 5° for dung beetle experiments, and the nearest 10° for robot experiments:

1. Place the robot/beetle in the center of the arena, facing a chosen initial bearing (alternating between 0° and 180°).
2. Execute the `choose` behaviour / allow the dung beetle to roll its dung ball, and record the exit bearing, denoted the c bearing.
3. Re-place the robot/beetle in the center of the arena, facing 180° from the initial bearing.
4. Execute the `follow` behaviour / allow the dung beetle to roll its dung ball and record the exit bearing, denoted the f_1 bearing.

5. Re-place the robot/beetle in the center of the arena, again facing the initial bearing.
6. Execute the `follow` behaviour / allow the dung beetle to roll its dung ball and record the exit bearing, denoted the f_2 bearing.
7. Two distinct CIB values can be determined by taking the shortest angular distance between f_1 and c, and between f_2 and c.

When a dung beetle is placed in the center of the arena, it is first physically separated from its dung ball, and then the dung ball placed in the center of the arena and the dung beetle placed so that its front two legs are touching the ball and the beetle's anterior is facing the chosen bearing.

The procedure for a full robot experiment can now be fully described:

1. If not already done for a particular arena, manually choose a speed setting, empirically record the angular velocity, empirically record the `rotation_time`, manually set the `rotation_time` for subsequent executions of the `choose` behaviour, and set the radius of the visual field such that no terrestrial landmarks can be seen from any location within the arena.
2. Replace the batteries of the robot with a fully charged battery pack.
3. Note the time and experimental conditions (for later calculating solar altitudes), and take a screenshot of the robot's visual field.
4. Select a model from Model A or Model CX and set relevant model parameters for this particular experiment.
5. Execute a 360° rotation from the center of the arena, logging the values of relevant internal variables during each frame of the rotation (to be presented as graphs characterising the experimental environment at the time of this rotation).
6. Repeat the procedure to record CIB results until there are 10 distinct CIB values for this experiment.

4.4 Experimental results

4.4.1 Dung beetle experiments

Dung beetle experiments were performed between 15:00 and 15:30, with solar altitude varying between approximately 49° and 42° over this time period. Three different dung beetles of the diurnal species *Scarabaeus lamarcki* and two of the diurnal species *Kheper nigroaeneus* were used for dung beetle experiments, but the results are grouped together for both kinds of beetle (the beetle species is distinguished in Table A.1). Beetles were collected from within the game farm 'Stonehenge', 70 km northwest of Vryburg, North West province, South Africa, approximate coordinates (-26.48,24.38), during February 2018.

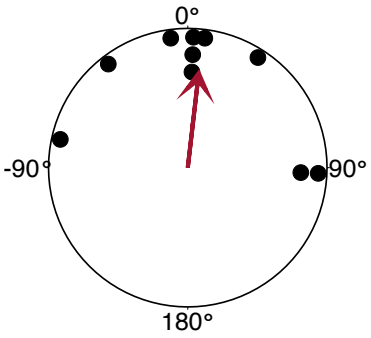

Figure 4.5: CIB results for dung beetle experiments; circular mean vector shown (red arrow)

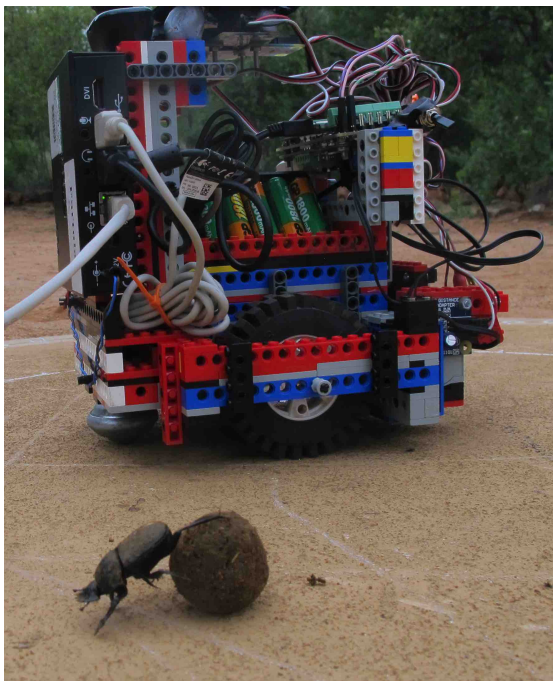

Figure 4.4: Relative size of the robot vs a diurnal dung beetle of the species *Scarabaeus lamarcki*

The most active beetles (out of a container of about 15 of each species) were selected for experiments. Dung beetle experiments were then conducted according to the dung beetle behavioural experiment paradigm described in Section 2.1, specifically step 6 of the procedure for a full robot experiment described in Section 4.3.

One selected *S. lamarcki* beetle could not be encouraged to complete its f_2 roll and so results using this beetle were discarded. The ($n = 10$) remaining CIB results for all dung beetle experiments are given in Figure 4.5. The circular mean vector for CIB results had bearing $\mu_{CIB} = 6.6^{\circ}$, magnitude $\tau_{CIB} = 0.68$, and circular standard deviation $\sigma_{CIB} = 45.6^{\circ}$. The circular mean of the absolute CIB results was $\mu_{|CIB|} = 32.9^{\circ}$.

4.4.2 Control condition

The control condition gives a baseline for how well the control mechanism performs menotaxis in the absence of any visual mechanism. As such, the robot's camera and light sensors were covered (see Figure 4.7), and it was verified that little to no visual information was available to the robot – the screenshot was near total darkness and polarised light sensor values were near

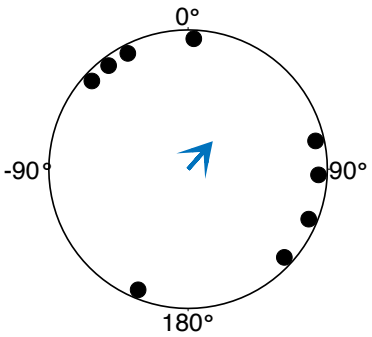

Figure 4.6: CIB results for the control condition; circular mean vector shown (blue arrow)

zero and did not vary as the robot rotated. Model CX was used, with both cues turned on, but note that the results would not be expected to depend on the model used, rather the control mechanism. The expected result then is that due to the stochastic nature of the control mechanism `control_forward`, forward-curved motor commands will be selected randomly. Due to the design of the experiment, this 'random walk' behaviour should result in very curved paths and CIB results distributed approximately uniformly randomly (depending on the size of the arena and the turning circle of the robot while forward-curving), indicating an inability to menotax.

Figure 4.7: Visual sensors obscured in the control condition

One f_2 bearing for the control experiment is not given as the robot's path was curved to the point where it took more than 30 seconds to exit the arena. The ($n = 9$) CIB results for the remaining control condition experiments are given in Figure 4.6. The circular mean vector for CIB results had bearing $\mu_{CIB} = 41.4^{\circ}$, magnitude $\tau_{CIB} = 0.24$, and circular standard deviation $\sigma_{CIB} = 70.7^{\circ}$. The circular mean of the absolute CIB results was $\mu_{|CIB|} = 74.8^{\circ}$. As expected, all paths were significantly more curved than in subsequent experiments (as judged by eye).

The original plan was to compare these control results for the robot to similarly gathered control results for the dung beetles, whereby a 'cap' is glued to the head of each beetle to obscure its view of the sky, as in [17] and [39]. Unfortunately, there was not enough time to perform these dung beetle control experiments.

4.4.3 CV experiment (Model A)

For these experiments, Model A was used, with only the CV intensity-gradient cue turned on. Experiments were performed between 14:30 and 15:00, with solar altitude varying between approximately 55° and 49° over this time period. Figure 4.9 shows typical celestial conditions during these experiments (screenshots were selected from five that were taken over the stated time period). As can be seen, cloud cover was patchy and inconsistent as the clouds move, but the sun was always within view, and bright enough that the CV vector points directly to the sun.

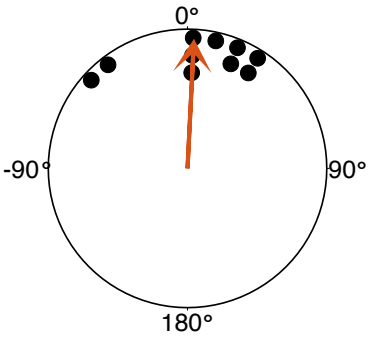

Figure 4.8: CIB results for the Model A using the CV cue; circular mean vector shown (orange arrow)

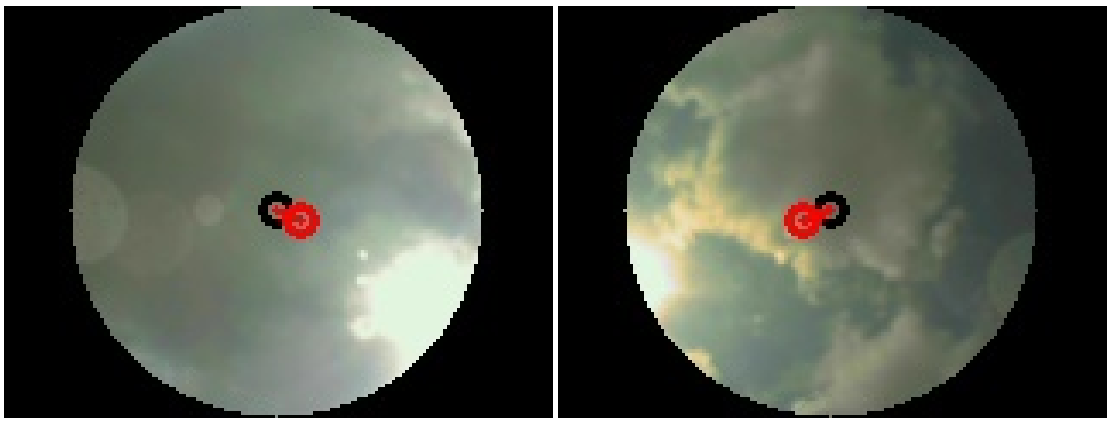

Figure 4.9: Selected celestial screenshots during CV, Model A experiments; black circle = visual center; red circle = image centroid; red line = CV vector

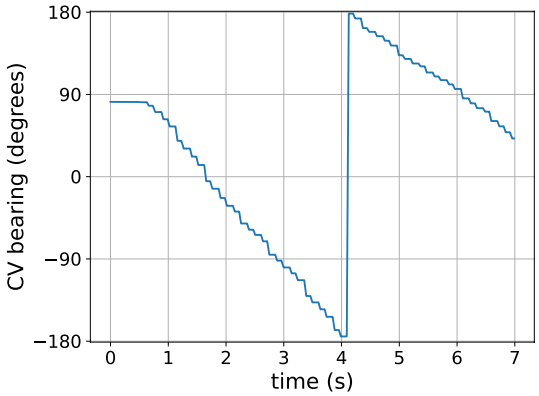

Figure 4.10: CV bearing values logged during a 360° rotation

Figure 4.10 shows CV bearing values logged internally by the system as the robot rotated 360° in the center of the arena. Bearings logged seem accurate and relatively smooth, matching what would be expected for such a rotation assuming the CV vector rotates as the robot does. The ($n = 10$) CIB results are given in Figure 4.8. The circular mean vector for CIB results had bearing $\mu_{CIB} = 2.9^{\circ}$, magnitude $\tau_{CIB} = 0.90$, and circular standard deviation $\sigma_{CIB} = 25.4^{\circ}$. The circular mean of the absolute CIB results was $\mu_{|CIB|} = 19.9^{\circ}$.

For an additional partial experiment using the `choose` behaviour, after the snapshot was taken (as indicated by programmed LEDs on the sensor board) the camera was partially obscured by a hand as the robot was moving in a forward-curve. Judged by eye, it was clear that on obstruction the robot immediately veered off course, and when the obstruction was removed, the robot immediately attempted to recover the original direction. This confirms that the camera must continually be viewing the sky for Model A to menotax successfully.

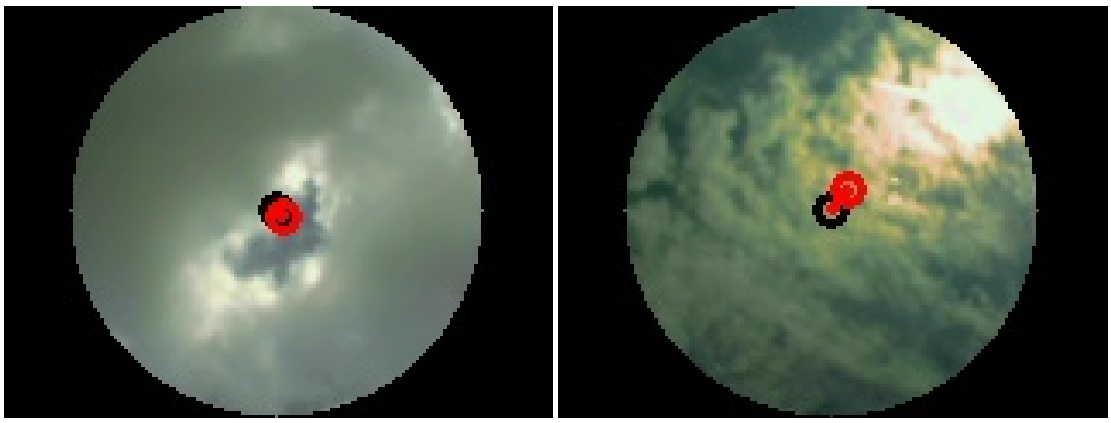

Figure 4.11: Selected celestial screenshots during CV, Model CX experiments; black circle = visual center; red circle = image centroid; red line = CV vector

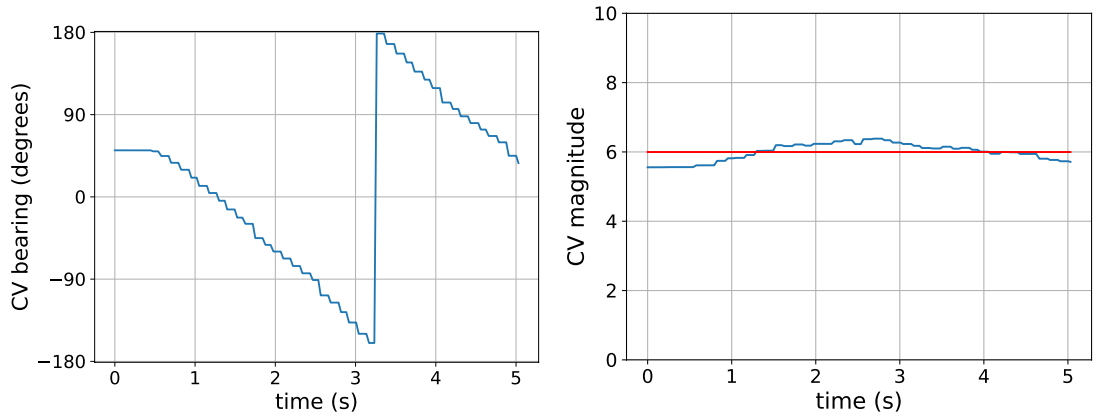

Figure 4.12: CV bearing and raw usefulness predictor values, τ_{CV}, logged during a 360° rotation; red line (right) indicates the mean, $\mu_{\tau_{CV}}$

4.4.4 CV experiment (Model CX)

For these experiments, Model CX was used, with only the CV intensity-gradient cue turned on. Experiments were performed between 13:30 and 14:00, with solar altitude varying between approximately 67° and 61° over this time period. Figure 4.11 shows typical celestial conditions during these experiments. As can be seen, the sun was once completely obscured by very thick clouds (left), but most of the time was clearly visible (right) and able to strongly influence the CV cue.

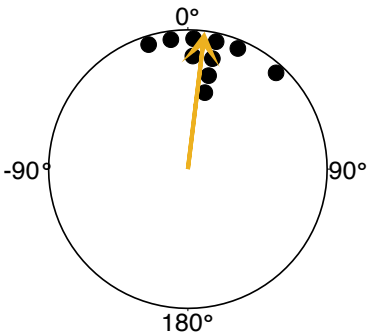

Figure 4.13: CIB results for Model CX using the CV cue; circular mean vector shown (yellow arrow)

Figure 4.12 shows CV bearing and magnitude values logged internally by the system as the robot rotated 360° in the center of the arena. Again, values logged matched what would be expected for facilitating successful menotaxis. The ($n = 10$) CIB results are given in Figure 4.13. The circular mean vector for CIB results had bearing $\mu_{CIB} = 6.9°$, magnitude $\tau_{CIB} = 0.96$, and circular standard deviation $\sigma_{CIB} = 15.4°$. The circu-

lar mean of the absolute CIB results was $\mu_{|CIB|} = 12.9^{\circ}$. Again, a 'hand-over-camera' experiment was performed to confirm that the camera must continually be viewing the sky for Model CX to menotax successfully.

4.4.5 Polarisation experiments

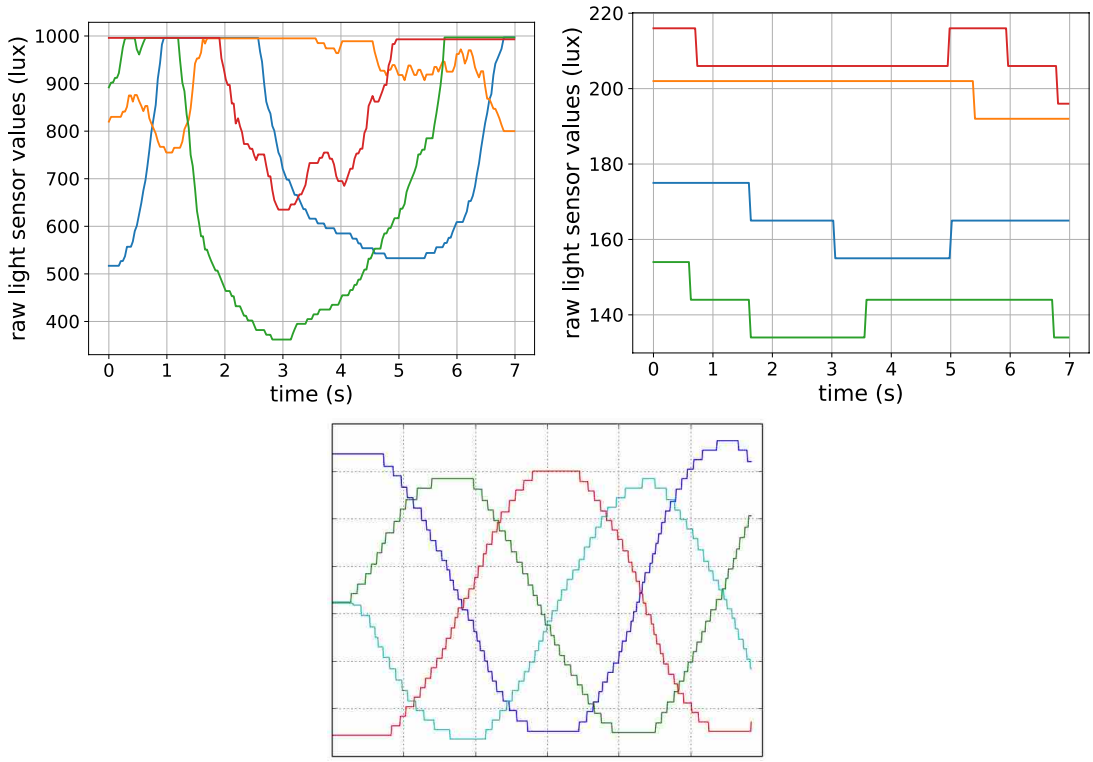

Figure 4.14: Comparison of raw light values logged during a 360° rotation under different conditions – solar altitude approximately 55°-29° in direct sunlight (top left); solar altitude approximately 42° under a partially cloudy sky (top right); under bright indoor lights covered with a linearly polarising filter (bottom, modified from [1])

Similarly to experiments with the CV intensity-gradient cue described in Sections 4.4.3 and 4.4.4, experiments were planned using the polarisation cue for both Models A and CX. Experiments were attempted at around 15:30, with solar altitude approximately 42°. Unfortunately, it was not found possible to utilise the polarisation cue on the day of the experiment.

The main problem, discovered early on in the field trip was that the very high brightness of the South African sky (compared to indoor experiments, or experiments performed in [1] under the Scottish sky) under most conditions would cause the polarised light sensors to saturate (light intensity after passing through the linearly polarised light filters would exceed the light sensors maximum value of 1000 lux). This is problematic, as in the case of saturation, calculating the angle or intensity of the linearly polarised component is no longer possible.

Attempting to overcome this problem, the polarised light sensors were covered with a neutral density filter (visible in Figure 4.2), which partially blocks the transmission of light but seemingly does not affect the polarisation of transmitted light. Different optical densities (strengths of the filter) were tried, and it was found that the need varies with celestial conditions. On exposure to direct sunlight, an optical density of 0.9 (transmitting 12.6% of light [40], achieved by stacking three filters each of optical density 0.3) was insufficient to prevent saturation (Figure 4.14, top left). For comparison, Figure 4.14 (bottom) shows polarised light sensor readings under conditions perfect for extracting polarisation information.

On the day of the experiment, under a partially cloudy sky with no direct sunlight, it was found that this optical density of 0.9 was sufficient to prevent saturation and reduce the light sensor values to a reasonable range (Figure 4.14, top right). Unfortunately, even though light sensor values were reduced to an appropriate range, they did not vary as expected during rotation, indicating that polarisation information was not detectable (verified by calculating and logging polarisation bearing values over the same rotation). It is unclear why this effect was seen, as the celestial conditions indicate that polarisation information should be available from the sky, and the high light sensor values seem to indicate that the optical density of 0.9 was not too high to block this information. It seems possible that by manipulating the optical density a solution might have been found, but unfortunately there was insufficient time on the day of the experiment to investigate this further.

4.4.6 Combination experiment (Model CX)

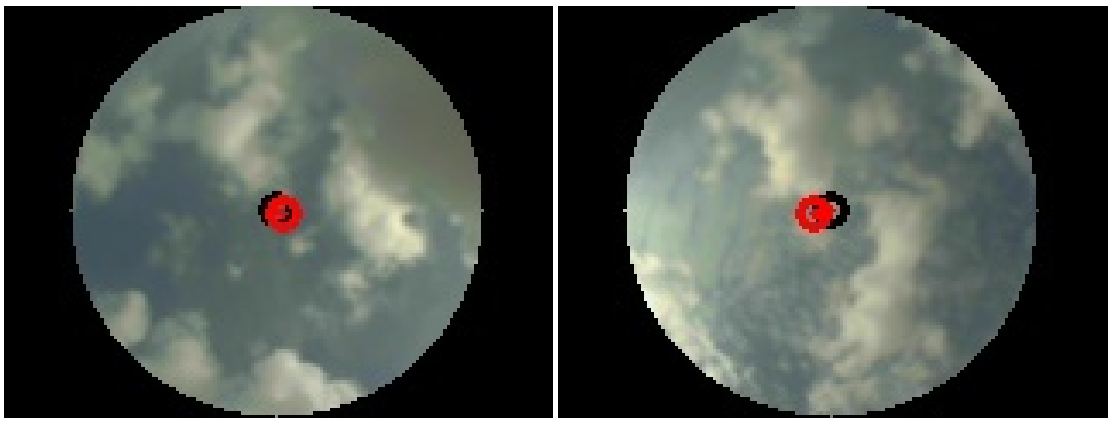

Figure 4.15: Selected celestial screenshots during Model CX combination experiments

Although the polarisation cue was not usable on the day of the experiment, combination experiments were performed in order to demonstrate that in the natural case where the polarisation cue is not useful but is still providing erroneous readings, and the intensity-gradient cue is useful, Model CX is able to ignore the polarisation cue and use only the intensity-gradient cue.

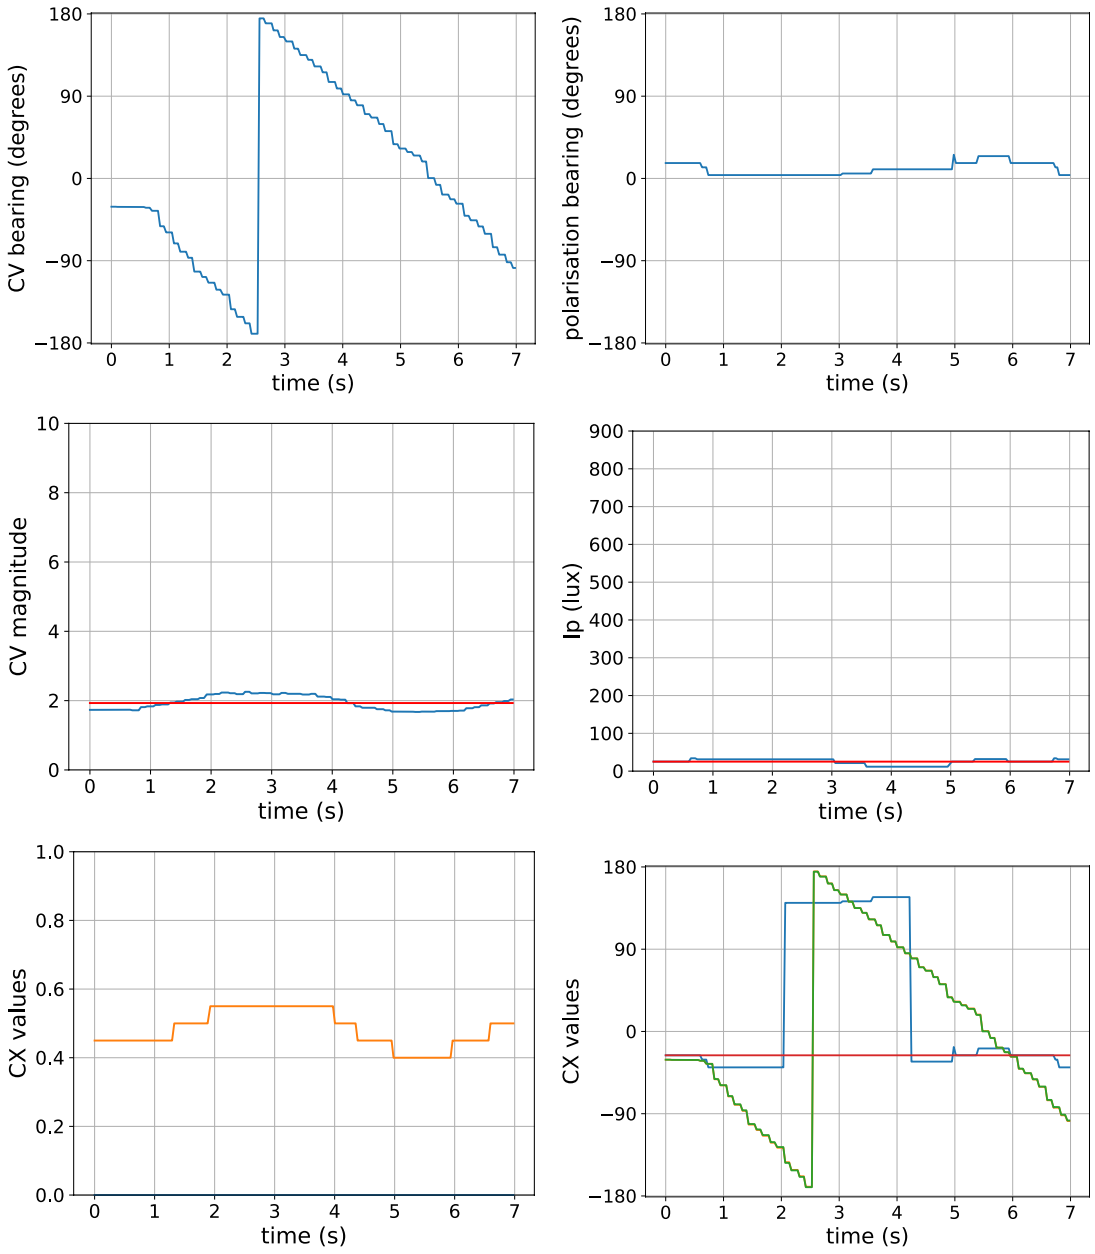

Figure 4.16: CX internal values logged during a 360° rotation; CV bearing (top left), polarisation bearing (top right); CV raw usefulness weightings (middle left, red line indicates mean); polarisation raw usefulness weightings (middle right, red line indicates mean); actual usefulness weightings of intensity-gradient (yellow) and polarisation (value 0) cues (bottom left); bearing encoded by X_{CV} intensity-gradient cue (yellow, identical to green combination values), X_{POL} polarisation cue (blue, after 180°-360° range extension), TB1 combination (green) and CPUX snapshot (red) CX cells

For these experiments, Model CX was used, with both cues turned on. Experiments were performed between 15:30 and 16:00, with solar altitude varying between approximately 42° and 36° over this time period. Figure 4.15 shows typical celestial conditions during these experiments. As can be seen, the clouds remained patchy but the sun was no longer directly visible. Although still sufficient for menotaxis, the effect of the sun on the CV cue was weaker than in experiments reported in Sections 4.4.3 and 4.4.4.

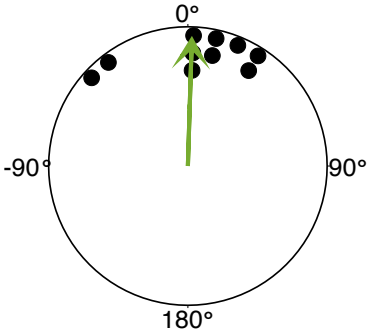

Figure 4.17: CIB results for Model CX combining the polarisation cue and the CV cue; circular mean vector shown (green arrow)

Figure 4.16 (bottom row) shows the magnitude values (left) and bearing values (right) encoded by CX cells logged internally by the system as the robot rotated 360° in the center of the arena. As can be seen, the magnitude values for the intensity-gradient cue are around 0.5 but for the polarisation cue are consistently zero, matching the fact that the CV magnitude is near its availability threshold of 2 pixels whereas the Ip is well below its availability threshold of 81 lux. As expected, the misleading polarisation bearing is thus ignored, and the integrated bearing equal to the CV bearing at all times. The ($n = 10$) CIB results are given in Figure 4.17. The circular mean vector for CIB results had bearing $\mu_{CIB} = 1.8^{\circ}$, magnitude $\tau_{CIB} = 0.91$, and circular standard deviation $\sigma_{CIB} = 24.9^{\circ}$. The circular mean of the absolute CIB results was $\mu_{|CIB|} = 18.9^{\circ}$.

4.4.7 RV experiment (Model CX)

For these experiments, Model CX was used, with only the RV intensity-gradient cue turned on. Experiments were performed between 14:00 and 14:30, with solar altitude varying between approximately 61° and 55° over this time period. Figure 4.19 shows RV bearing magnitude values logged internally by the system as the robot rotated 360° in the center of the arena.

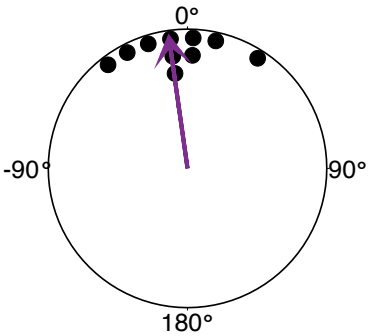

Figure 4.18: CIB results for Model CX using the RV cue; circular mean vector shown (purple arrow)

The ($n = 10$) CIB results are given in Figure 4.18. The circular mean vector for CIB results had bearing $\mu_{CIB} = -8.1^{\circ}$, magnitude $\tau_{CIB} = 0.95$, and circular standard deviation $\sigma_{CIB} = 18.6^{\circ}$. The circular mean of the absolute CIB results was $\mu_{|CIB|} = 16.0^{\circ}$. Again, a 'hand-over-camera' experiment was performed to confirm that the camera must continually be viewing the sky for Model CX to menotax successfully.

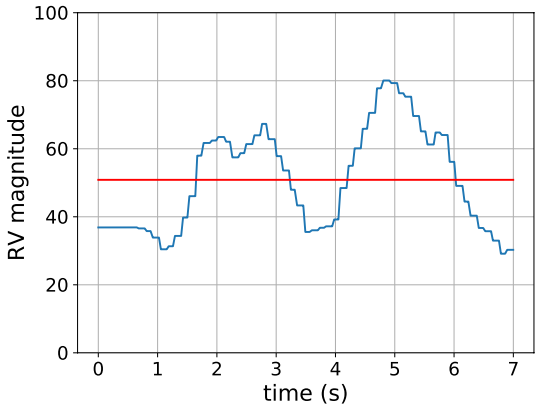

Figure 4.19: RV raw usefulness predictor values, τ_{RV}, logged during a 360° rotation; red line indicates the mean $\mu_{\tau_{RV}}$

4.4.8 Summary of results

Raw exit bearing results for all full CIB experiments[1], organised by experiment, are presented in Table A.1 of Appendix A.

Table 4.1 summarises the CIB results for each condition. The circular mean, μ_{CIB} indicates the 'average' exit bearing relative to the corresponding chosen exit bearing – values close to zero indicate that on average the chosen exit bearing is taken. Note that e.g. two CIB values of x° and $-x^\circ$ would produce a zero μ_{CIB}. To account for this, the circular mean of the absolute CIB values, $\mu_{|CIB|}$, is also reported, which would produce a non-zero $\mu_{|CIB|}$ value in such cases. The magnitude of the circular mean vector, $\tau_{CIB} \in [0,1]$, indicates the spread of the corresponding circular mean – a value of $\tau_{CIB} = 1$ indicates that all CIB results were equal to μ_{CIB}, whereas a value of $\tau_{CIB} = 0$ indicates that the CIB results are uniformly spread throughout the range $[0^\circ, 360^\circ)$. Similarly, the circular standard deviation, σ_{CIB}, indicates the spread of the corresponding circular mean – the higher the value of σ_{CIB} the more the spread deviates from μ_{CIB}.

experiment	solar altitude (°)	μ_{CIB} (°)	τ_{CIB}	σ_{CIB} (°)	$\mu_{\lvert CIB\rvert}$ (°)
control	control	41.4	0.24	70.7	74.8
Model CX, CV	67-61	6.9	0.96	15.4	12.9
Model CX, RV	61-55	-8.1	0.95	18.63	16.0
Model A, CV	55-49	2.9	0.9	25.4	19.9
dung beetle	49-42	6.6	0.68	45.6	32.9
Model CX, both	42-36	1.8	0.91	24.9	18.9

Table 4.1: Summary of numerical CIB results – $n = 9$ for the control experiment, $n = 10$ for all other experiments; circular statistics were calculated using [41]

Figure 4.20 visualises the μ_{CIB} and τ_{CIB} results of Table 4.1 by depicting the circular mean vectors (the bearing of a vector is the μ_{CIB} bearing for that experiment, and the length of the vector is the τ_{CIB} value).

[1] A video of an example CIB experiment can be found at `https://youtu.be/3-HJhZamev4`

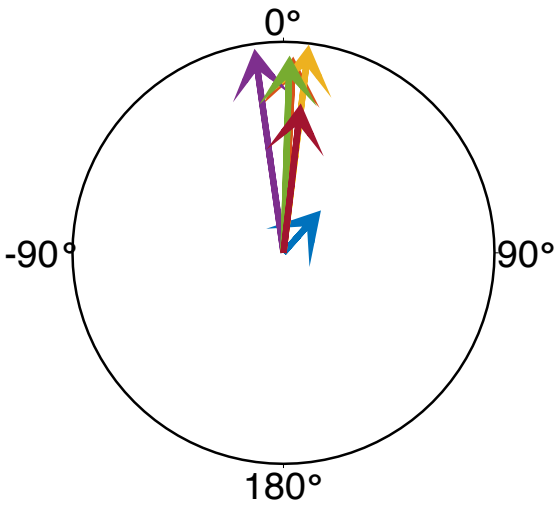

Figure 4.20: Comparison of CIB circular mean vectors for dung beetle (red), control (blue), Model A, CV (orange), Model CX, CV (yellow), Model CX, both (green), and Model CX, RV (purple) experiments; a unit circle (black) is shown to indicate how close the magnitude of circular mean vectors is to the maximal magnitude of 1

For perfect menotaxis we would expect to see values of μ_{CIB} close to 0° and values of τ_{CIB} close to $\tau_{CIB} = 1$, indicating that the average exit bearing taken is close to the chosen exit bearing and that most results were actually close to this average exit bearing. This would correspond to circular mean vectors pointing close to 0°, with arrowhead close to touching the unit circle. We would also expect relatively low values of σ_{CIB}, and values of $\mu_{|CIB|}$ relatively close to 0°.

We can see then from Table 4.1 and Figure 4.20 that menotaxis was clearly successful for all robot experiments except for the control experiment. Especially given the relatively small number of experiments performed, all values are acceptably close to those expected for perfect menotaxis. For the control experiment on the other hand, the spread was very large, with τ_{CIB} close to the zero value that would indicate a totally uniform distribution of exit bearings. For the dung beetle experiment, μ_{CIB} is acceptably close to 0° but the spread is significantly higher than for robot experiments. This is to be expected for dung beetle vs robot experiments however, as results from the literature show large differences in the behaviour of individual dung beetles, whereas the behaviour of the robot is deterministically programmed (the only stochasticity is in the random walk performed when menotaxis is not possible e.g. in the control condition).

Table 4.1 also lists solar altitudes for experiments performed. The idea here is to compare the success of menotaxis across different solar altitudes, especially for the non-control robot experiments. The most notable result is that in terms of spread (τ_{CIB} and σ_{CIB} values), menotaxis seems more successful for the higher altitudes compared to the lower altitudes. This doesn't match what we would expect if considering the altitude of the sun in isolation (the idea that the further from the zenith, the more of a contrast between the solar and anti-solar celestial hemispheres), but makes more sense when considering that in later experiments overall conditions became less bright and the sun stopped being directly visible to the robot.

Table 4.2 offers a comparison of the raw usefulness predictors by solar altitude, as recorded during the 360° rotation at the start of each experiment.

solar altitude (°)	$\mu_{\tau_{CV}}$	$\sigma_{\tau_{CV}}$	$\mu_{\tau_{RV}}$	$\sigma_{\tau_{RV}}$	μ_{Ip}	σ_{Ip}
control	0.09	0.2	3.6	0.04	10.3	0
67-61	1.0	0.7	38.3	5.8	25.6	11.0
61-55	3.6	1.0	50.9	14.9	182.5	72.5
55-49	8.7	0.7	79.4	11.5	337.3	149.9
42-36	1.9	0.2	44.4	5.6	25.1	7.5

Table 4.2: Comparison of raw usefulness predictors during a 360° rotation at different solar altitudes; dances were performed closer to the start than the end of each time period, meaning the higher solar altitude is more likely

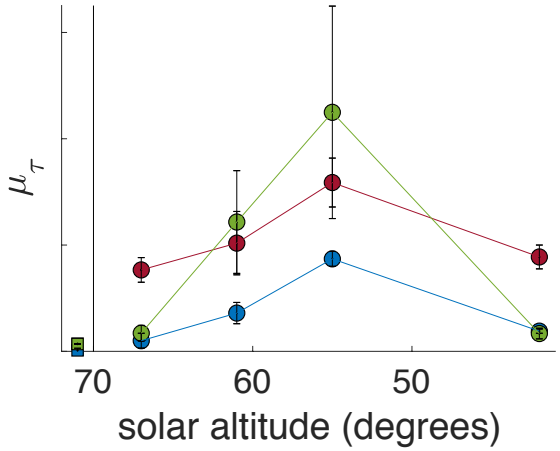

Figure 4.21: Comparison of raw usefulness predictors during a 360° rotation at different solar altitudes – mean values, μ_τ, of CV (blue), RV (red) and polarisation (green) raw usefulness predictors; for visibility, the CV predictor is scaled by 5 and the polarisation predictor is scaled by 1/3; error bars denote sample standard deviation, σ_τ; black line separates control condition (squares) from outdoor conditions (circles)

The most salient result from Table 4.2 and Figure 4.21 is the clear difference between the values of the raw usefulness predictors for the control vs non-control conditions – it is clear that by suitably thresholding each raw usefulness predictor, the lack of availability of all cues in the control condition could have been predicted. This is a positive result, showing that for all cues the raw usefulness predictors can correctly predict availability.

Considering now the intensity-gradient cues CV and RV, we can see that the predicted usefulness of the cues increased as the solar altitude decreased, up until the last experiment when the sun was no longer visible. Matching the analysis of the CIB results above, it seems that in practice, the exact solar altitude doesn't matter as much as whether or not the sun is actually visible. Cloud conditions then do not seem to affect predicted usefulness – although intuitively it makes sense that partially cloudy conditions can be problematic due to the constantly changing picture of the sky, it seems that in practice this doesn't matter for menotaxis so long as the sun is visible.

Considering now the polarisation cue, it is interesting to note that despite the problems with saturation described in Section 4.4.5, the raw usefulness predictor also increases with decreasing solar altitude up until the last experiment. Bearing in mind that the dances were performed with the neutral density filter installed, it seems that the result for the last experiment matches the inability to determine polarisation angle for this experiment. While the explanation for this failure remains elusive, it is a positive result that the lack of availability of the polarisation cue could have been predicted by suitable thresholding of the raw usefulness predictor.

In summary, it first should be noted that especially considering the low sample size, comparisons made may not be significant, and so the data can only weakly support and conclusions made. Despite this limitation, the results seem to demonstrate quite clearly that menotaxis can be successfully produced using the CV cue with Model A or the CV or RV cues with Model CX. Furthermore, for the limited conditions tested (one cue unavailable), integration of CV and polarisation cues in Model CX was successful. Finally, the usefulness predictors for all cues (CV, RV and polarisation) correlate with the availability of the respective cues, and even with the solar altitude which is thought to predict the usefulness of cues. For all experiments carried out, it was clear that under the cloudy conditions on the day of the experiment, the visibility of the sun affects the ability to menotax using intensity-gradient cues.

Chapter 5

Discussion

This chapter aims to discuss the implications of model building from Chapter 3 and experimental results from Chapter 4 on possible answers to the three research questions, restated here:

Section 5.1:
How are celestial visual cues, particularly the intensity-gradient and polarisation cues, encoded in the dung beetle CX for menotaxis purposes?

Section 5.2:
How are different cues combined in the CX during menotaxis?

Section 5.3:
By what neurologically plausible mechanism does the dung beetle CX produce menotax behaviour?

Section 5.4 then goes on to evaluate the project as a whole, in particular noting the limitations of the work; and finally Section 5.5 draws overall conclusions with respect to the three research questions.

5.1 How are menotaxis cues encoded in the CX?

5.1.1 The weighting of cues

Dung beetle behavioural experiments show that when the intensity-gradient cue is mirrored whilst the polarisation cue stays the same, dung beetles turn to roll in a bearing that lies between the bearings suggested by the individual cues. This shows that within the context of a single snapshot, these cues are **combined**, and it seems likely that all cues used for menotaxis are similarly combined.

Crucially, the intermediate bearing chosen in these experiments was dependent on the altitude of the sun at the time of the experiment, suggesting that in the combination of intensity-gradient and polarisation cues, at least one of the cues is **weighted** according to the altitude of the sun. The altitude of the sun is known to affect the quality of these cues with respect to the accuracy of menotaxis, and so it was assumed that the weighting of cues reflects their potential usefulness for menotaxis. This **usefulness** weighting of cues is conceptually separate from **reliability** weighting of cues, which reflects the a priori weighting of some cues (e.g. the intensity-gradient cue) over others (e.g. the polarisation cue). Within this framework, the preference ordering of cues referred to in the literature (such as in [15]) reflects this reliability weighting, and is likely determined by evolution. This conceptual clarification between cue combination and cue preference is a novel contribution of this project.

While reliability weightings are likely innate, usefulness weightings are most likely coupled with each measurement. Applying this hypothesis to the sinusoidal neural encoding of cues suggested within the CX path integration model led to a novel interpretation of the encoding – the amplitude of the sinusoidal representation encodes the usefulness weighting of the bearing measurement encoded. This interpretation held when incorporated into the full CX model for menotaxis (which includes cue combination). Usefulness predictors (measurable quantities related to the bearing measurement for each cue that correlate with menotaxis accuracy) were suggested for both the intensity-gradient and polarisation cues, although experimental results from Section 3.2.2 and Chapter 4 suggest that these metrics may be better described as availability predictors (i.e. to turn cues 'on' or 'off'). While conceptually successful, the cue encoding and combination using these usefulness predictors was not properly tested experimentally. This marks one of the major limitations of this year of the project, discussed further in Section 5.4.

Dung beetle behavioural experiments could be performed to experimentally test the hypotheses developed in a wider range of conditions than the behavioural experiments considered. For example, indoor experiments could be carried out to verify the cue combination behaviour observed, and quantify the relationship between altitude of a point-source and weighting. Similarly, the degree of polarisation could be varied to see how dung beetles respond, and in particular when combined with a fixed intensity-gradient cue.

5.1.2 Intensity-gradient cue: CV vs RV

The regions vector (RV) intensity-gradient cue was newly developed for this year of the project, inspired by the encoding of bearings within the CX path integration model – instead of a sinusoidal encoding, average brightness values for non-overlapping regions of the visual field are simply stored directly.

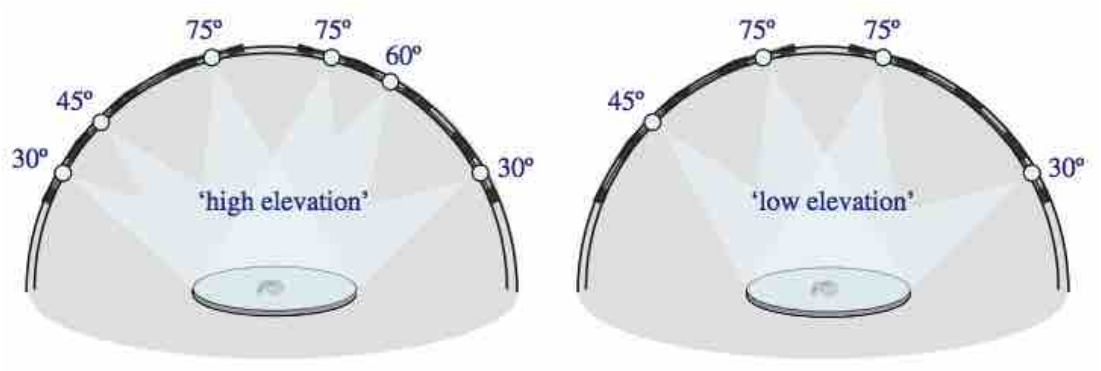

Figure 5.1: Visualisation of the 'high elevation' (left) and 'low elevation' (right) conditions; modified from [29]

Although the RV cue was shown to be sufficient to produce menotaxis within the CX model developed, it is unclear why this is the case. Specifically, it has not been shown that the CX steering mechanism will always produce the correct error signal when comparing RV vectors. Although the experimental results seem convincing, given that the experiments were quite rushed and only 5 full experiments were performed, it is not inconceivable that the data is unreliable. In other words, the results would be a lot more convincing if the mechanism behind RV comparison was well understood.

Nevertheless, assuming that the RV cue is sufficient for menotaxis for some conditions, it is worth comparing the RV and CV cue under more complex 'edge case' conditions, to see if there are scenarios where one cue accounts for dung beetle behavioural results but the other does not.

In [29], experiments were performed involving more complex patterns of intensity than the usual single green light spot. The conditions shown in Figure 5.1 were presented, where the intensity patterns lie within a single narrow band across the dung beetle's visual field, but are non-symmetrical about the center of the field. As such, the centroid will lie on the band but not in the center, and will rotate as the field of view rotates, resulting in successful menotaxis if the CV cue were being used. This does not match dung beetle results however, which show a uniform distribution of bearings (inability to menotax). These are the only known experimental results that do not support the hypothesis that the CV cue is used by dung beetles.

Under the same conditions, the values of the RV cue are expected to be identical when the field of view is rotated 180°. This would be expected to produce a bimodal distribution of bearings, similar to when using the polarisation cue in isolation. Again, this does not match dung beetle results. From these conditions then, it is clear that neither the RV nor the CV cue is actually being used by dung beetles.

On the other hand, in the first year of the project, it was shown that under totally symmetrical conditions about the center (e.g. two lights at either side of the center), the CV cue does account for dung beetle results (inability to menotaxis). Under the same conditions, the RV cue would again be expected to produce a bimodal distribution however. For this edge case then, the CV cue predicts dung beetle behaviour better than the RV cue.

Overall, the CV cue comes off more favourably compared to the RV cue as implemented. However, neither cue is totally supported by the results, and so the idea of summing over regions should not be ruled out. The division of the visual field into the eight regions as implemented was chosen arbitrarily – it may be that a more thoughtful division (possibly involving a different number of regions) will produce an intensity-gradient cue that does account for the edge cases described in this section. Developing such a division is left for future work.

5.1.3 Terrestrial landmarks

In the literature, it is usually assumed that dung beetles do not use 'terrestrial landmarks' for menotaxis, mostly based on the results of [39]. However, following the discussion of Section 3.5 it became clear that the robot using the CV cue (and probably the RV cue) will very successfully use terrestrial landmarks for menotaxis **if these landmarks are visible**. Given the discrepancy between the literature assumptions and the robot's behaviour, this section re-examines the experiments and results from [39] with the aim to explain these results within the context of intensity-gradient cues such as the CV and RV cues:

- Experiments 1 and 2 of [39] demonstrate that dung beetles ignore terrestrial cues (a dung pile, bushes, trees, a water-tower, and the skyline) when celestial cues (a clear sky) are available.
- Experiment 3 demonstrates that dung beetles wearing a cap obscuring their dorsal eyes cannot menotax at all.
- Experiment 4 demonstrates that under totally overcast conditions, menotax performance is greatly reduced despite the availability of terrestrial cues.

From these experiments, the authors of [39] conclude that dung beetles 'ignore landmarks for straight-line orientation'. Arguably however, the approach of separating landmarks from celestial cues is not meaningful. It seems very unlikely that dung beetles can distinguish what is terrestrial from what is celestial, and it is predicted that if terrestrial cues take up a significant portion of the visual field used for the intensity-gradient cue, they will be used for menotaxis by dung beetles in the same way as by the robot.

However, it does seem that celestial cues are generally preferred to terrestrial cues. This can be explained by assuming that the portion of the visual field used for the intensity-gradient cue is mostly skyward facing, and so the terrestrial contribution to any intensity-gradient based cue will usually be negligible. Particularly when the sun is visible, this is likely to dominate any intensity-gradient cue utilised.

This interpretation can explain the results of [39] more clearly:

- For experiments 1 and 2, celestial cues are available and so likely dominate terrestrial cues, especially if the sun is visible (not clear from the report).

- For experiment 3, it is clear that the entire visual field used for the intensity-gradient cue is obscured, terrestrial cues and all.

- For experiment 4, the idea was to remove celestial cues (the overcast sky) meaning that an inability to menotax would demonstrate that terrestrial cues were not being used for menotaxis. Contrary to this expectation, dung beetle paths for experiment 4 were actually significantly straighter (mean path length 116.4 ± 9.4cm) than for experiment 3 (mean path length 217.7 ± 46.2cm), showing that **some** cues were being used. The results could be explained by terrestrial cues in the periphery being used but not being as useful for menotaxis (e.g. the sources of terrestrial cues are closer, and so move a lot while the dung beetle rolls). A simpler explanation might be that the polarisation or intensity-gradient cues were used, but their usefulness was decreased by the overcast conditions.

This analysis gives a different conclusion from the authors of [39] – rather than ignoring landmarks, it is simply the case that the visual field of dung beetles (for the purposes of menotaxis) is vertically facing and narrow enough that most terrestrial landmarks cannot be seen.

This conclusion now raises a methodological issue – for experiments 3 and 4 dung beetle paths were tracked using a camera suspended 3m above the center of the arenas used, but the possibility of the dung beetles using the camera as a cue does not seem to have been considered. The very reason for the poor performance in experiment 4 might have been that the beetle **was** using the camera as a cue, but that this was a very poor cue as it moved so much relative to the beetle while rolling. Under clear skies when the sun is visible or the polarisation cue is preferred, the use of the camera in this way may not adversely affect dung beetle paths as much.

For this reason, it is suggested that a more comprehensive behavioural study be carried out on dung beetles to determine the exact field of view considered for the intensity-gradient cue. Results of such an investigation could have major implications on the way dung beetle behavioural experiments should be carried out in future. For example, using a single green light spot indoors, the light could be lowered until the beetle is no longer able to menotax. Similarly to this hypothetical experiment, it was shown in the experimental results for the robot in Chapter 4 that as the sun naturally lowered, menotaxis immediately became more difficult once the sun disappeared from the robot's field of view – it is likely that the same drop in performance occurs for dung beetles.

The discussion of this section is mostly possible due to the ability to visualise what the robot is seeing during experiments, giving interpretable knowledge of the exact visual inputs at different times. Whilst clearly more difficult for dung beetles, the insights of this section emphasise the importance of visualising exactly what dung beetles can actually see during experiments. This will allow easier control over visual inputs during behavioural experiments, and the confidence to know for sure which aspects of the experimental environment can be utilised and which cannot. As such, the pursuit of a suitable experimental methodology to visualise the dung beetle's field of view is a major recommendation of this project.

5.2 How are menotaxis cues combined in the CX?

The method of pointwise summing the neural encoding of cues to combine bearing information in a weighted fashion is obviously dependent on the neural encoding, and so this research question overlaps with the first. The method of combining cues also turned out to be coupled with the implementation of the snapshot mechanism, and so this research question similarly overlaps with the next. This section chooses to focus on the cue combination and snapshot mechanisms leaving the next section to focus on the steering mechanism and related control issues.

With the aim of combining the polarisation and CV intensity-gradient cues, a necessary preliminary step was to extend the range of the polarisation cue from the natural $[0^{\circ}, 180^{\circ})$ to the full $[0^{\circ}, 360^{\circ})$. During the project, a CX circuit was successfully implemented to do so, which involves utilising directional information from the intensity-gradient cue (any full ranged cue could potentially be used for this purpose).

Talking about 'cue combination' can sometimes be misleading in that it is actually errors that are combined: angular offsets between each cue at the current time vs the time of the snapshot. Two behaviourally equivalent mechanisms were proposed to achieve this error combination. They are behaviourally equivalent in that both implementations accomplish the same mathematical function of finding the circular weighted mean of these errors to give a final error signal. The difference between so called 'cue integration' and 'error integration' is whether or not errors are combined implicitly (achieved by remembering the angular difference **between different cues** at the time of the snapshot) or explicitly. The pros and cons of each method are considered here:

Cue integration:

- Cues are integrated to form an explicit 'current bearing' (frame of reference is snapshot-dependent), an intuitively appealing concept.
- The integrated current bearing is not weighted (weighting of each cue in the snapshot is stored implicitly in the weighted combination at the time of the snapshot), meaning the encoding can be enforced by a 'ring attractor' (see Section 2.3 for details) i.e. there is always a current bearing, again intuitively appealing.
- The killer problem: a neural mechanism with the assumed encoding of the cues could **not** be conceived to achieve the operation of explicitly calculating the angular distance between cues, or of reapplying a stored offset to subsequent measurements – it remains possible that a neural mechanism exists, but the operations involved do not seem natural to the encoding, and a successful implementation would likely involve converting back and forth between some other representation.

Error integration:

- This method is simpler in that no offset need be stored for each separate cue, although more complex in that multiple error signals must be calculated – overall, error integration is simpler within the context of the particular neural encoding of cues.
- The weighting of each cue is stored explicitly in the snapshot (as the snapshot is now a coupling of weighted bearing measurements for each cue) – this easy access to the snapshot weightings seems appealing, although no use for these could be conceived.

The cue integration method was implemented first, and the main flaw quickly became apparent. Before conceiving of the error integration method however, much time was dedicated (unfruitfully) to trying to overcome this problem. The error integration method itself was not devised until the field trip, and there was unfortunately not enough time to implement this fully. Due to the behavioural equivalence however, experimental results from Chapter 4 should hold for both methods. However, the process of implementing the error integration in full would have likely lead to more insights and a model of the neural mechanisms that may direct further neuroethological studies. As such, implementation of the error integration method is recommended as the next task for the dung beetle biorobot project.

5.3 How is menotaxis behaviour produced in the CX?

The CX steering mechanism used in the model was taken directly from the CX path integration model. Within the context of the cue integration method, this steering mechanism was combined with a simple snapshot circuit which stores the integrated cue in a snapshot to be compared every iteration with the current integrated cue (i.e. the snapshot mechanism from [3]). Experimental results from Chapter 4 indicate that this mechanism works as expected, and is sufficient for menotaxis.

While not strictly dealt with by the CX (it is assumed that the error output is a binary signal indicating the direction to turn in order to match the snapshot bearing), another factor in the production of menotaxis behaviour is the control mechanism used. A new control mechanism was devised during this year of the project that produces curved forward movement – this was thought preferable to the unrealistically straight paths produced by the control mechanism used during the first year of the project. The evolutionary rational behind this type of movement is clear – natural selection does not strictly favour dung beetles with very straight paths, rather dung beetles that get quickly away from the competition of the dung heap (without accidentally coming back). The speed gains by not stopping to dance all the time clearly outweigh the more curved path.

A perhaps more interesting discussion is to focus on what happens in the ‘control’ condition (when the visual sensors of the robot or beetle are completely obscured) in relation to the output of the CX model. As noted in Section 3.4, a small amount of

Gaussian noise is added to the error signal before determining which direction to turn. This mimics the noise added to the error signal in the CX path integration model [2] which for this model results in search-like circling behaviour once the home vector has been traversed. This is because at this point the home vector will have zero magnitude (the CPU4 updates continue while the home vector is traversed), making all comparisons equal. The random noise means a random turn direction is taken each iteration, resulting in this 'nest searching behaviour'. Note that in the model error noise is explicitly added, whereas in the dung beetle it is very plausible that noise is simply the result of noisy neuronal properties and interactions.

The benefit of added noise for menotaxis is similar – when the error signal (before noise) is actually zero (e.g. if the visual sensors are obscured both during and after the snapshot, meaning the snapshot bearing always exactly matches the current bearing), there is a uniformly random probability of forward-curving anticlockwise vs clockwise. The result is a 'random walk' of curves, producing randomly curved paths. The alternative without noise in this case is that the robot will move as straight as possible in the absence of visual cues. Arguably, this alternative would be more beneficial for the purposes of getting away from the dung pile as quickly as possible. However, it was noted that in the literature when dung beetles fail at menotaxis due to being totally deprived of visual input (as in [39] and [17]), their paths are randomly curved in this manner.

The situation is less clear when the cues are removed between the snapshot and ball rolling, as in experiments performed in [3]. The result in such cases was a 'uniform distribution of bearings', however it is unclear if the beetles were rolling straight in randomly chosen directions or if they were taking curved 'random-walk' paths. This ambiguity makes it difficult to confidently claim whether or not this control mechanism (including noise) accounts for dung beetle behaviour. A recommendation for future dung beetle behaviour experiments then is to always report the curvature of paths (even if only qualitatively).

5.4 Critical evaluation

As in the first year of the project, the main method of evaluating the model was to perform behavioural tests (implemented on the biorobot) using the same experimental methodology as used for dung beetle behavioural experiments. In this year of the project, the next step was taken by actually performing dung beetle experiments in parallel to robot experiments, ensuring an identical procedure was adopted and allowing results to be directly compared (with the caveat that the celestial conditions were constantly changing throughout all experiments). In the sense that both the robot and dung beetle were able to menotax under these conditions, the project can be evaluated positively – the central task of implementing and testing a CX model of dung beetle menotaxis was completed successfully.

There are of course many limitations to the work carried out. Although the CX menotaxis model was successfully implemented, any success should be qualified by the limitations and limited scope of the implemented model:

- The scope of the model is strictly limited to operations after cues have been fully encoded, and ends with a simple binary signal to turn left or right. There is no attempt to model low-level visual processing, and similarly, there is no attempt to model actuation beyond the decision of which direction to turn. Although limited, this scope reflects the fact that hardware and hardware-software interfaces of the robot do not resemble those of the dung beetle.

- Similarly, there has also been no attempt to model cues other than the polarisation and intensity-gradient cues, as the biorobot is not equipped with sensors to capture these. However, a successful model should be easily extendable to include additional cues – while not demonstrated, it seems clear that any additional (fully encoded) cues could simply be included as additional inputs to the reliability-weighted TB1 sum. In this sense the model does seem suitably extendable.

- Compared to the CX path integration model [2], new neurons and connections are introduced that have not been anatomically derived. As such, the model can only claim to be anatomically plausible, rather than anatomically constrained.

- The level of abstraction within the model is high, similar to the level of explanation of the CX path integration model in Section 2.3 – neuron activity is considered only in terms of the neural encoding described, where 'activity' is a dimensionless numerical quantity, and the physical realisation of operations on these quantities is not considered.

- Even neurological plausibility is not assured – no attempt was made to demonstrate that neurological mechanisms such as the multiplexer-style circuit (used in the snapshot circuit and the polarisation extension circuit) or the recurrent connections (the CPUX cells connect to themselves as inputs to the snapshot multiplexer) are physically possible with neural circuitry. The disclaimer then is that the details of the model remain speculative, and are only plausible as far as there are no operations assumed that obviously cannot be accomplished by physical neurons.

- From [33], we know that although the brain regions responsible for visual processing (and indeed the eyes themselves) are much bigger for nocturnal compared to diurnal dung beetles, the CX brain regions between the species are indistinguishable. It seems safe to assume that this will also be the case for crepuscular ball-rolling dung beetles. The behavioural predictions of any CX model developed should therefore match behavioural experiments involving and comparing diurnal, crepuscular and nocturnal dung beetles (e.g. [15]). This is indeed the case, as changing reliability weightings in the TB1 sum can allow species-specific cue preferences to be captured. This was not experimentally demonstrated however.

- Even within the specified scope, the model could still be improved, in particular to use the more promising error integration method rather than the cue integration method (see Section 5.2).

There are also limitations affecting the quality of the experimental results:

- The major limitation of the results, and perhaps of the project as a whole, was that there was no final experimental demonstration that dung-beetle style cue combination (or cue preference) can be produced by the model. Experiments to verify this were planned, but due to the problem with the polarisation cue (described in Section 4.4.5), the limited time to troubleshoot this, and the unfortunate weather conditions, there simply wasn't enough time during the field trip to perform these planned experiments. Specifically, it experiments replicating those of [14] were planned, whereby the sun is mirrored in order to modify the intensity-gradient cue without modifying the polarisation cue.

- If these planned experiments were performed, it was also planned to use the comparison to dung beetle experiments to tune model parameters (the reliability weightings of cues and the sigmoidal functions used to transform raw usefulness predictors). If successful, this would have produced much stronger evidence in support of the model.

- Along the same lines, the planned control condition (using caps) for the dung beetle experiments could not be performed.

- The other obvious limitation, which has already been discussed in previous sections, is the low amount of data collected for the experiments that were performed.

- It was hoped that results for this year of the project could be compared to results from the first year of the project. However, results related to metrics for cue preference could not be compared directly, due to changes in the metrics used. To avoid this limitation, all results for this year of the project are reported in terms of the raw usefulness predictor (before passing through the sigmoid) in the hope that results become more future-proof.

- As in the first year of the project, the (seemingly irreducible) camera delay for the intensity-gradient cue cause issues with the accuracy of menotaxis – due to the new control mechanism however, this now manifested in sinusoidal paths in the case of successful menotaxis. Fortunately, with appropriate adjustments to the speed of the robot (corresponding to amplitude of the path taken), this factor will not have affected results significantly.

Overall, the project was successful in what it achieved, but perhaps did not achieve as much as it should have. The depth of investigation and quantity of meaningful results could have gone further, and perhaps would have if the weather conditions during the field trip were more fortunate. Weather conditions should should not have been relied upon however, and on reflection, more work should have been carried out prior to the field trip to prepare for the possibility of bad weather, and a contingency plan to carry out indoor (artificial cue) experiments should have been carried out.

On a more positive note however, the scope of the project was significantly narrowed compared to the first year of the project, giving a much more focused feel to the report. I believe that overall, this has resulted in a more comprehensive and meaningful investigation, at least with respect to the research questions chosen. I hope then that the contributions of the project, particularly this year of the project, will have significant scientific value in terms of directing future dung beetle behavioural research, and securing the future of the dung beetle biorobot project.

5.5 Conclusions

Noting the overlap between the three research questions, the following specific conclusions can be drawn with respect to each question:

How are celestial visual cues, particularly the intensity-gradient and polarisation cues, encoded in the dung beetle CX for menotaxis purposes?

- The RV cue implemented or a similar regions-based intensity-gradient cues are potential candidates for how the intensity-gradient cue is encoded, however an encoding of the CV cue remains more likely.
- Every bearing measurement is likely coupled with a weighting reflecting a cue-specific prediction of the accuracy of that measurement with respect to menotaxis.
- The degree of linear polarisation, but more likely the intensity of the largest linearly polarised component of light are potential encoding-appropriate predictors of menotaxis accuracy for the polarisation cue.
- The length of the CV (encoding-appropriate), or the maximal difference between values of the RV, are potential predictors of menotaxis accuracy for the intensity-gradient cue.
- Bearing and weighting information can be coupled together in an eight-neuron population encoding, operations on which are suitable for a neurologically plausible menotaxis mechanism.
- Beetles likely make no distinction between terrestrial and celestial cues, and likely use all information available from their albeit narrow and celestially-directed visual field.

How are different cues combined in the CX during menotaxis?

- Polarisation bearing information can be combined in a neurologically plausible manner with directional information from a full-ranged cue to extend the range of the polarisation cue from $[0^\circ, 180^\circ)$ to $[0^\circ, 360^\circ]$, producing appropriately unimodal bearing distributions utilising the extended polarisation cue.
- Angular distances (errors) between the snapshot and current bearings for different cues are likely combined with respect to the weightings of the relevant measurements.

- Errors are likely stored separately in the snapshot and combined explicitly.
- During cue combination, cues as a whole are likely also weighted by the reliability of the cue within the dung beetle's natural environment (as determined by natural selection).

By what neurologically plausible mechanism does the dung beetle CX produce menotax behaviour?

- Control signals which can be used to produce menotaxis behaviour can themselves be produced by neurologically plausible mechanisms that:
  - Extend the polarisation cue as described above.
  - Combine errors as described above.
  - Implement the snapshot mechanism.
- In the absence of strong control signals from the menotaxis mechanism, control signals in dung beetles are likely noisy to the point of producing 'random walk' curved paths in these conditions.

Finally, based on the discussion in Sections 5.1, 5.2 and 5.3, and qualified by the limitations noted in Section 5.4, the overall conclusion can be stated simply: 'the CX model for dung beetle menotaxis presented, and the underlying assumptions made in the development of this model, represent the best neurologically plausible explanation for cue encoding, cue combination, cue preference and menotaxis behaviour to date'.

Appendix A

Raw Data

control			
	50	125	0
	220	310	180
	310	150	310
	80	210	–
	200	170	310
Model CX, CV	c (°)	f_1 (°)	f_2 (°)
solar altitude 67°			
	340	20	350
	240	260	240
	350	0	340
	240	240	220
	140	150	150
Model CX, RV			
solar altitude 61°			
	170	200	150
	220	190	180
	120	120	130
	160	150	160
	340	330	330

Model A, CV			
solar altitude 55°			
	130	160	150
	330	290	280
	10	40	30
	160	160	160
	40	50	40
dung beetle	c (°)	f_1 (°)	f_2 (°)
solar altitude 49°			
S. lamarcki	50	55	10
K. nigroaeneus	160	160	80
S. lamarcki	230	230	230
S. lamarcki	120	150	110
K. nigroaeneus	270	0	0
Model CX, both			
solar altitude 42°			
	320	270	280
	160	190	180
	210	210	210
	250	280	260
	250	250	260

Table A.1: Raw data from all full CIB experiments; 3 different *Scarabaeus lamarcki* and 2 different *Kheper nigroaeneus* beetles were used for beetle experiments; an additional *S. lamarcki* beetle was discarded as it could not be encouraged to make its f_2 roll; one f_2 bearing for the control experiment is not given as the robot circled for more than 30 seconds; True North (as determined by a magnetic compass) was at approximate bearing 170°; solar altitudes calculated using a free online tool [42] using the experiment times (13:30, 14:00, 14:30, 15:00, 15:30) and the location (University of the Witwatersrand, Johannesburg, South Africa).

Appendix B

Innovation Initiative Grant Application

Project Title Field-testing a dung beetle biorobot in South Africa

Project Overview Travel to Johannesburg to field-test a dung beetle biorobot in its native environment.

Aims and Objectives In recent years, the navigational abilities of South African dung beetles have been shown to be unique within the animal kingdom. Not only can these beetles see the sun, but their tiny eyes can detect the (invisible to us) polarisation pattern of the day and night sky, and even the stars and the Milky Way galaxy. Perhaps more extraordinarily, these beetles have been shown to use all of this information to navigate whilst rolling their dung balls.

My two-year UG4/UG5 project investigates this navigational ability. The overall aim of the project is to model the neural circuitry responsible for utilising these visual cues, and replicate the same navigational behaviour on a biorobot – if the model is simple, neurologically plausible (based on known neural structures and properties) and produces the correct behaviour, then we can be confident we understand how the tiny dung beetle brain produces this complex behaviour. The crucial point is that the chaotic visual information available in the sky is too complex to simulate realistically, either in a computer simulation or in a controlled laboratory setting. In order to draw confident conclusions from this investigation, it is therefore necessary to demonstrate the robot 'works' under the same South African skies at which these fascinating insects spend their lives gazing.

The aim of this IIG project is to travel with a team of biologists to their field-testing location in South Africa, and there spend 7 days testing the robot within the same natural habitat in which this team have been observing beetles for many years.

Potential benefits or Significant Results In South Africa, carefully controlled experiments on the robot will be carried out, and the results compared to identical experiments performed on dung beetles (either already published or performed on the same trip). If the results match, we may confidently conclude that the model represents our best current understanding of the neurology of this complex behaviour. If the results do not match, this would suggest that our understanding of the behaviour or environment is incomplete. It is very likely in this case that the nature of the difference in behaviour will immediately suggest modifications to the model that must be made and/or features of the environment that have been misunderstood. A successful demonstration may then be possible within the 7 days. If not, the improvement in our understanding is still likely to be significant.

In either case, the results of the investigation as a whole would be of much greater scientific value if experiments in South Africa were included. This would make a publication of the investigation much more likely, benefitting myself, the research teams and universities involved, and the scientific community. A specific target for publication is the 'Living Machines' conference, likely to be running next year with a submission deadline in March.

As well as publishing the work, the robot and this investigation are likely to be used for public engagement and other events in the future e.g. Science Fairs, university open days, project presentation days, and when technology companies visit the university. A published paper with more significant results, as well as videos of the robot's beetle-like behaviour in South Africa, will greatly increase the quality and engagement of these presentations. In this way, an IIG can directly contribute to increasing public engagement with the sciences, and potential funds directed towards the institute, school, and university.

Is there anything innovative or unique about your project? The biorobotics research methodology used in this project is not itself unique – the idea of building robots to develop and test neural and behavioural models of biological behaviour has been around for a few decades – but in the history of science is still a novel and innovative way to advance biological investigations and understand how insects and other animals achieve their impressive behaviours and abilities.

This project methodology is not unique then, but is a unique application of an existing innovative methodology, and connects for the first time the study of beetle behaviour and behavioural robotics. No dung beetle biorobot has ever been constructed, programmed or tested. Additionally, no neural models of dung beetle navigation behaviour have ever been proposed. The neural model itself is based on a very recently published (Oct 2018) neural model of visual navigation in honey-bees. This project is, therefore, an advancement of this cutting-edge research, extending the model in a new direction to account for a different kind of behaviour, and connecting two similar areas of behavioural biological research.

Feasability of your project This is the second year of the project. In the first year, I successfully constructed the biorobot, implemented a behavioural model (not neurologically plausible) capturing the key features of the navigational behaviour, and successfully tested the robot in a range of limited and unnatural conditions, testing within an artificial arena equipped with vaguely sky-like visual cues.

One potential risk is that the neural model is not completed and tested (to the best extent possible without a South African sky) in time for the trip, planned for February 16th-26th 2018. Assuming conceptual problems are solved, there is enough time to develop, implement, and test the model (based on experience developing the behavioural model last year), but there is always the risk that conceptual problems may not be solved. In this scenario, however, field-testing (and videoing) the behavioural model developed would still be a scientifically valuable endeavour for which this IIG would be necessary for.

In terms of the time in South Africa, I have allowed for two days of travel in each direction, giving 7 full days available for testing, which is about three times the time required to actually carry out the preplanned experiments (based on my experience testing the robot last year).

The only other risk identified is damage to the robot, in particular, due to travel, and weather conditions. The risk of damage due to travel will be minimised by storing the robot in a suitable way as hand-luggage. The risk of damage due to weather will be minimised by making the robot as waterproof as possible (to be done this year), and by keeping the robot out of the sun when not running experiments. In general, tools and spare parts will be taken so that it is possible to repair any non-catastrophic damages during the trip.

Bibliography

[1] R. H. El-Naggar, "A biorobotic investigation of the visual mechanisms controlling dung beetle celestial menotaxis," Bachelor's thesis, School of Informatics, The University of Edinburgh, 2017.

[2] T. Stone, B. Webb, A. Adden, N. B. Weddig, A. Honkanen, R. Templin, W. Wcislo, L. Scimeca, E. Warrant, and S. Heinze, "An Anatomically Constrained Model for Path Integration in the Bee Brain," *Current Biology*, vol. 27, no. 20, pp. 3069–3085, 2017. [Online]. Available: https://doi.org/10.1016/j.cub.2017.08.052

[3] B. el Jundi, J. J. Foster, L. Khaldy, M. J. Byrne, M. Dacke, and E. Baird, "A Snapshot-Based Mechanism for Celestial Orientation," *Current Biology*, vol. 26, no. 11, pp. 1456–1462, 2016. [Online]. Available: https://doi.org/10.1016/j.cub.2016.03.030

[4] V. Ivan, "Robotics: Science and Systems – Practicals – Getting Started," [Online]. Available: http://wcms.inf.ed.ac.uk/ipab/rss/practicals/getting-started. [Accessed: 20-Mar-2017].

[5] D. Lambrinos, H. Kobayashi, R. Pfeifer, M. Maris, T. Labhart, and R. Wehner, "An Autonomous Agent Navigating with a Polarized Light Compass," *Adaptive Behaviour*, vol. 6, no. 1, pp. 131–161, 1997. [Online]. Available: https://doi.org/10.1177/105971239700600104

[6] Wikipedia contributors, "Stokes parameters — Wikipedia, The Free Encyclopedia," [Online]. Available: https://en.wikipedia.org/w/index.php?title=Stokes_parameters&oldid=756356856. [Accessed: 26-Mar-2017].

[7] J. B. Tatum, "Polarized light and the stokes parameters," [Online]. Available: http://astrowww.phys.uvic.ca/~tatum/physopt/physopt4.pdf. Accessed: [26-Mar-2017].

[8] Wikipedia contributors, "Rayleigh sky model — Wikipedia, The Free Encyclopedia," [Online]. Available: https://en.wikipedia.org/w/index.php?title=Rayleigh_sky_model&oldid=764458102. [Accessed: 14-Feb-2017].

[9] B. el Jundi, J. Smolka, E. Baird, M. J. Byrne, and M. Dacke, "Diurnal dung beetles use the intensity gradient and the polarization pattern of the sky for orientation," *Journal of Experimental Biology*, vol. 217, no. 13, pp. 2422–2429, 2014. [Online]. Available: http://jeb.biologists.org/content/217/13/2422

[10] A. Rosebrock, "Finding the Brightest Spot in an Image using Python and OpenCV," [Online]. Available: http://www.pyimagesearch.com/2014/09/29/finding-brightest-spot-image-using-python-opencv/. Accessed: [26-Mar-2017].

[11] V. V. Hafner, "Adaptive Homing – Robotic Exploration Tours," *Adaptive Behavior*, vol. 9, no. 3–4, pp. 131–141, 2001. [Online]. Available: http://dx.doi.org/10.1177/10597123010093002

[12] A. Vardy, "Biologically Plausible Methods for Robot Visual Homing," Ph.D. dissertation, School of Computer Science, Carleton University, 2005. [Online]. Available: https://curve.carleton.ca/system/files/etd/5ce96fc8-c740-4a88-84ee-cc3084cb8f26/etd_pdf/ae25b93a3a96a9d003482a206c4e8e89/vardy-biologicallyplausiblemethodsforrobotvisual.pdf

[13] M. Mangan, "Visual homing in field crickets and desert ants: a comparative behavioural and modelling study," Ph.D. dissertation, School of Informatics, University of Edinburgh, 2011. [Online]. Available: https://www.era.lib.ed.ac.uk/handle/1842/5678

[14] M. Dacke, B. el Jundi, J. Smolka, M. Byrne, and E. Baird, "The role of the sun in the celestial compass of dung beetles," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 369, no. 1636, 2014. [Online]. Available: https://dx.doi.org/10.1098%2Frstb.2013.0036

[15] B. el Jundi, E. J. Warrant, M. J. Byrne, L. Khaldy, E. Baird, J. Smolka, and M. Dacke, "Neural coding underlying the cue preference for celestial orientation," *Proceedings of the National Academy of Sciences*, vol. 112, no. 36, pp. 11 395–11 400, 2015. [Online]. Available: https://doi.org/10.1073/pnas.1501272112

[16] E. Baird, M. J. Byrne, J. Smolka, E. J. Warrant, and M. Dacke, "The Dung Beetle Dance: An Orientation Behaviour?" *PLOS ONE*, vol. 7, no. 1, pp. 1–6, 2012. [Online]. Available: http://dx.doi.org/10.1371%2Fjournal.pone.0030211

[17] L. Khaldy, "The Dung Beetle Dance: The role of visual cues in dung beetle orientation behavior," Master's thesis, Department of Biology, Lund University, 2014. [Online]. Available: http://www.lunduniversity.lu.se/lup/publication/4466562

[18] E. Baird, M. J. Byrne, C. H. Scholtz, E. J. Warrant, and M. Dacke, "Bearing selection in ball-rolling dung beetles: is it constant?" *Journal of Comparative Physiology A*, vol. 196, no. 11, pp. 801–806, 2010. [Online]. Available: http://dx.doi.org/10.1007/s00359-010-0559-8

[19] G. Halffter and E. G. Matthews, *The Natural History of Dung Beetles of the Subfamily Scarabaeinae (Coleoptera, Scarabaeidae)*, 1966. [Online]. Available: http://www.socmexent.org/revista/folia/Num%2012-14/1-308.pdf

[20] M. L. Brines and J. L. Gould, "Skylight Polarization patterns and Animal Orientation," *Journal of Experimental Biology*, vol. 96, no. 1, pp. 69–91, 1982. [Online]. Available: http://jeb.biologists.org/content/96/1/69

[21] T. W. Cronin, E. J. Warrant, and B. Greiner, "Celestial polarization patterns during twilight," *Applied Optics*, vol. 45, no. 22, pp. 5582–5589, 2006. [Online]. Available: http://ao.osa.org/abstract.cfm?URI=ao-45-22-5582

[22] J. Gál, G. Horváth, A. Barta, and R. Wehner, "Polarization of the moonlit clear night sky measured by full-sky imaging polarimetry at full Moon: Comparison of the polarization of moonlit and sunlit skies," *Journal of Geophysical Research: Atmospheres*, vol. 106, no. D19, pp. 22 647–22 653, 2001. [Online]. Available: http://dx.doi.org/10.1029/2000JD000085

[23] I. Pomozi, G. Horváth, and R. Wehner, "How the clear-sky angle of polarization pattern continues underneath clouds: full-sky measurements and implications for animal orientation," *Journal of Experimental Biology*, vol. 204, no. 17, pp. 2933–2942, 2001. [Online]. Available: http://jeb.biologists.org/content/204/17/2933

[24] R. Hegedüs, S. Åkesson, and G. Horváth, "Polarization patterns of thick clouds: overcast skies have distribution of the angle of polarization similar to that of clear skies," *Journal of the Optical Society of America A*, vol. 24, no. 8, pp. 2347–2356, 2007. [Online]. Available: http://josaa.osa.org/abstract.cfm?URI=josaa-24-8-2347

[25] M. Dacke, P. Nordström, C. H. Scholtz, and E. J. 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*, vol. 188, no. 3, pp. 211–216, 2002. [Online]. Available: http://dx.doi.org/10.1007/s00359-002-0295-9

[26] M. Dacke, M. J. Byrne, C. H. Scholtz, and E. J. Warrant, "Lunar orientation in a beetle," *Proceedings of the Royal Society of London B: Biological Sciences*, vol. 271, no. 1537, pp. 361–365, 2004. [Online]. Available: http://rspb.royalsocietypublishing.org/content/271/1537/361

[27] 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 of London B: Biological Sciences*, vol. 366, no. 1565, pp. 697–702, 2011. [Online]. Available: http://rstb.royalsocietypublishing.org/content/366/1565/697

[28] M. Dacke, E. Baird, M. Byrne, C. H. Scholtz, and E. J. Warrant, "Dung Beetles Use the Milky Way for Orientation," *Current Biology*, vol. 23, no. 4, pp. 298–300, 2013. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0960982212015072

[29] J. J. Foster, B. el Jundi, J. Smolka, L. Khaldy, D.-E. Nilsson, M. J. Byrne, and M. Dacke, "Stellar performance: mechanisms underlying Milky Way orientation in dung beetles," *Philosophical Transactions of the Royal Society of*

London B: Biological Sciences, vol. 372, no. 1717, 2017. [Online]. Available: http://rstb.royalsocietypublishing.org/content/372/1717/20160079

[30] B. el Jundi, J. J. Foster, M. J. Byrne, E. Baird, and M. Dacke, "Spectral information as an orientation cue in dung beetles," *Biology Letters*, vol. 11, no. 11, 2015. [Online]. Available: http://rsbl.royalsocietypublishing.org/content/11/11/20150656

[31] R. Loesel, D. R. Nässel, and N. J. Strausfeld, "Common design in a unique midline neuropil in the brains of arthropods," *Arthropod Structure & Development*, vol. 31, no. 1, pp. 77–91, 2002. [Online]. Available: https://doi.org/10.1016/S1467-8039(02)00017-8

[32] K. Pfeiffer and U. Homberg, "Organization and Functional Roles of the Central Complex in the Insect Brain," *Keram Pfeiffer and Uwe Homberg*, vol. 59, no. 1, pp. 165–184, 2014. [Online]. Available: https://doi.org/10.1146/annurev-ento-011613-162031

[33] E.-V. Immonen, M. Dacke, S. Heinze, and B. el Jundi, "Anatomical organization of the brain of a diurnal and a nocturnal dung beetle," *Journal of Comparative Neurology*, vol. 525, no. 8, pp. 1879–1908, 2017. [Online]. Available: https://doi.org/10.1002/cne.24169

[34] D. B. Turner-Evans and V. Jayaraman, "The insect central complex," *Current Biology*, vol. 26, no. 11, pp. R453–R457, 2016. [Online]. Available: https://doi.org/10.1016/j.cub.2016.04.006

[35] C. M. Harley and R. E. Ritzmann, "Electrolytic lesions within central complex neuropils of the cockroach brain affect negotiation of barriers," *Journal of Experimental Biology*, vol. 213, no. 16, pp. 2851–2864, 2010. [Online]. Available: https://doi.org/10.1242/jeb.042499

[36] K. Neuser, T. Triphan, M. Mronz, B. Poeck, and R. Strauss, "Analysis of a spatial orientation memory in *Drosophila*," *Nature*, vol. 453, pp. 1244–1247, 2008. [Online]. Available: https://doi.org/10.1038/nature07003

[37] M. V. Srinivasan, "Where paths meet and cross: navigation by path integration in the desert ant and the honeybee," *Journal of Comparative Physiology A*, vol. 201, pp. 533–546, 2015. [Online]. Available: https://doi.org/10.1007/s00359-015-1000-0

[38] Wikipedia contributors, "Optical flow — Wikipedia, The Free Encyclopedia," [Online]. Available: https://en.wikipedia.org/w/index.php?title=Optical_flow&oldid=820418940. [Accessed: 09-April-2018].

[39] M. Dacke, M. Byrne, J. Smolka, E. Warrant, and E. Baird, "Dung beetles ignore landmarks for straight-line orientation," *Journal of Comparative Physiology A*, vol. 199, no. 1, pp. 17–23, 2013. [Online]. Available: http://dx.doi.org/10.1007/s00359-012-0764-8

[40] Edmund Optics Ltd., "Understanding Neutral Density Filters," [Online]. Available: https://www.edmundoptics.eu/resources/application-notes/optics/understanding-neutral-density-filters/. [Accessed: 02-Apr-2018].

[41] P. Berens, "Circular Statistics Toolbox (Directional Statistics) for MATLAB," [Online]. Available: https://uk.mathworks.com/matlabcentral/fileexchange/10676-circular-statistics-toolbox--directional-statistics. [Accessed: 03-Apr-2018].

[42] T. Hoffmann, "SunCalc," [Online]. Available: https://www.suncalc.org/#/-26.18,28.0436,17/2018.02.24/13:30/1/0. [Accessed: 31-Mar-2018].