# Using reconstructed visual reality in ant navigation research

Zoltán Kócsi

Submitted for the degree of Doctor of Philosophy of the Australian National University.

September, 2021

## Declaration

I declare that the content of this thesis is my original work. Any and all contributions by others are clearly marked and acknowledged.

This work has not been submitted, in whole or in part, for a degree in any other university. Its content has not been previously published except where explicitly stated. I was the first author or a principal contributor to the co-authored papers included in the thesis.

_____

Zoltán Kócsi

Canberra, 2021

*Dedicated to the memory of those to whom I would have been proud to show this work, but no longer can.*

## Abstract

Visual insect navigation is an active research topic. Insects have low resolution eyes and a tiny brain, yet they continuously solve very complex navigational problems; an ability that underpins fundamental biological processes such as pollination and parental care. Understanding the methods they employ would have profound impact on the fields of machine vision and robotics.

As our knowledge on insect navigation grows, our physical, physiological and neural models get more complex and detailed. To test these models we need to perform increasingly sophisticated experiments.

Evolution has optimised the animals to operate in their natural environment. To probe the fine details of the methods they utilise we need to use natural visual scenery which, for experimental purposes, we must be able to manipulate arbitrarily.

Performing physiological experiments on insects outside the laboratory is not practical and our ability to modify the natural scenery for outdoor behavioural experiments is very limited. The solution is reconstructed visual reality, a projector that can present the visual aspect of the natural environment to the animal with high fidelity, taking the peculiarities of insect vision into account. While projectors have been used in insect research before, during my candidature I designed and built a projector specifically tuned to insect vision; the first of its kind to my knowledge.

To allow the animal to experience a full panoramic view, the projector completely surrounds her. The device (dubbed “the Antarium”) is a polyhedral approximation of a sphere. It contains 20 thousand pixels, which is sufficient resolution for *Myrmecia* eyes. The pixels are made out of light emitting diodes (LEDs) that match the spectral sensitivity of *Myrmecia*. Furthermore, the device can also cater for polarisation, a feature of light that humans cannot see but many insects, including *Myrmecia* ants, can. In addition, insects have a much higher fusion frequency limit than humans, therefore the device has a very high flicker frequency (9kHz) and also a high frame rate (190fps).

In the Antarium the animal is placed in the centre of the projector on a trackball. To test the trackball and to collect reference data, outdoor experiments were performed where ants were captured, tethered and placed on the trackball. The apparatus with the ant on it was then placed at certain locations relative to the nest and the foraging tree and the movements of the animal on the ball were recorded and analysed. The outdoor experiments proved that the trackball was well suited for our ants, and also provided the baseline behaviour reference for the subsequent Antarium experiments. In addition, these experiments were also designed to test a hypothesis about the animals not only memorising where to go but also where *not* to go (attractive and repellent views).

To assess the Antarium, the natural habitat of the experimental animals was recreated as a 3-dimensional model. That model was then projected for the ants and their movements on the trackball was recorded, just like in the outdoor experiments. Initial feasibility tests were performed by projecting a static image (without feedback from the trackball to the projected image), which matches what the animals experienced during the outdoor experiments. To assess whether the ant was orienting herself relative to the scene, we rotated the projected scene around her and monitored whether she re-aligned herself with the rotated image or not. Various statistical methods were used to compare the outdoor and in-Antarium behaviour.

The results proved that the concept was solid, but they also uncovered several shortcomings of the Antarium. Most importantly, due to insufficient information available at the time of the LED selection, the actual LED wavelengths are a poor match to the spectral sensitivity of the ants, especially in the UV region.

Nevertheless, even with its limitations the Antarium was used to perform experiments that would be very hard to do in a real environment.

In one experiment the foraging tree was repositioned in or deleted from the scene to see whether the animals go to where the tree *is* or where by their knowledge it *should* be. The results suggest the latter but the absence or altered location of the foraging tree certainly had a significant effect on the animals, including some of them aborting the foraging trip and returning to the nest.

In another experiment the scene, including the sky, was re-coloured to see whether colour plays a significant role in navigation. Due to the spectral mismatch, that experiment had no usable result, except that even the very small amount of UV information that the device provides was able to statistically significantly improve the navigation of the animals.

To rectify the device limitations discovered during the experiments a new, improved projector was designed. Custom made LEDs that do match the spectral sensitivity of the animals were employed. The geometry of the device as well as the LED arrangement has been improved to provide a better visual representation of the environment than the existing model. The Antarium Mk-II device also uses custom optics to increase light intensity and minimise spurious light reaching the animal. It has higher pixel density, increased frame rate and better colour resolution than its predecessor. The new projector is currently in the implementation phase, many of its major components (including the custom LEDs) have already been purchased.

## Acknowledgement

I owe a great deal to my supervisory panel for their help during my journey: Dr. Allen Cheung for his mathematical insights; Dr. Ajay Narendra for all he taught me about ants, their biology and behaviour; and my chair, Prof. Jochen Zeil who is the most patient supervisor a candidate could wish for. Not only he taught me an enormous amount about insects in general and ants in particular, but his infectious respect for and interest in any living creature has really changed the way I look at the world around me.

I would like to thank the members and visitors of our laboratory: Prof. Robert Taylor, the living lexicon of ants; Dr. Trevor Murray with whom I spent many a day (and night) building the Antarium and running experiments in it; Dr. Fiorella Ramirez-Esquivel for all she taught me about ant sensory systems and for our chats about biology; Mr. Jesse Wallace for his help in building/fixing the Antarium; Prof. Willi Ribi, Ms. Ladina Ribi, Prof. Eric Warrant, Prof. Hansjürgen Dahmen, Dr. Antoine Wystrach, Dr. Sebastian Schwarz, Dr. Moosarreza Zahedi for their help, advise and inspiration; Ms. Chloé Raderschall and Dr. Lisa Vlahos for their support.

The most kind and helpful Ms. Karen Scholte from the HDR office guided me through a lot of the administrative hurdles of my PhD. Thank you.

I had the privilege to meet many very bright and helpful people at various conferences and to have thought provoking and eye opening conversations with them. I had several particularly interesting and inspiring chats with Prof. Barbara Webb, Prof. Bruno van Swinderen, Prof. Ken Cheng, Prof. Andrew Barron, Dr. Sylvain Foret (RIP) and Prof. Jan Hemmi.

I am also indebted to my immediate family: my children, Bence and Dorottya for their ‘Go, Dad, go!’ attitude towards my starting a PhD and their interest in my research; and my wife A/Prof. Krisztina Valter who, apart from being my wonderful and supportive partner for the most part of my life, very patiently showed the errors in my ways of applying engineering criteria to biological experiments and taught me a lot about vertebrate vision and neurological

systems. At the other end of the ancestral tree, I must thank my parents and grandparents who instilled the love of science and culture in me, encouraged me to study, to be curious, to ask questions and to seek answers.

Finally, I was lucky to grow up at a place and time where free high-quality education, easy access to arts, science and books were considered fundamental human rights; where knowledge was respected and studying was encouraged. I wholeheartedly thank all my teachers, formal or otherwise, who taught me everything I know and who inspired me to keep learning. I will never be able to pay down my debt to you.

## Abbreviations

| | |
|---|---|
| ANN | Artificial neural network |
| ASOT | Anterior-superior optic trackt |
| DRA | Dorsal rim area |
| CX | Central complex |
| FPGA | Field programmable gate array |
| FV | Full-vector |
| GPS | Global positioning system |
| IDF | Image difference function |
| IP | Internet protocol |
| LED | Light emitting diode |
| MB | Mushroom body |
| PCB | Printed circuit board |
| PI | Path integrator |
| RGB | Red, green, blue |
| RMS | Root-mean-square |
| SD | Standard deviation |
| TCP | Transmission control protocol |
| UDP | User datagram protocol |
| USB | Universal Serial Bus |
| UV | Ultraviolet |
| ZV | Zero-vector |

# Table of Contents

**Introduction....17**
Ants....19
Ant vision....22
The eyes....22
Geometry....24
Anatomy....26
Dorsal Rim Area....35
The ocelli....35
The optic lobes....37
Retina....39
Lamina....39
Medulla....40
Lobula complex....41
Higher level brain centres....42
Central complex....44
Mushroom bodies....47
Summary....52
Navigation....55
Path integration....57
Speed measurement....59
Direction measurement....60
Summing vectors with neurones....64
The reach of path integration....65
Navigation using the visual panorama....68
Cognitive maps....68
Landmark based navigation....69
Panorama based navigation....71
Learning walks....78
Where does that leave us?....81
**The design of the Antarium....85**
Abstract....86
Abbrevations....87
Introduction....87
The Antarium design....94
Geometry....94
Pixel arrangement....96
Polarisation....99
LED driving....100
Flicker considerations....100
Video delivery and frame rate....101
Architecture....103
Power distribution....105
Thermal considerations....106
Distributor board....108
Design tools....110

3D rendering and driver software....111
The trackball system....112
Antarium camera....113
Proof of concept....113
Outlook....116
Conflict of Interest....117
Author Contributions....118
Funding....118
Acknowledgments....118
**Outdoor trackball experiments....119**
Introduction....119
Published work with the trackball....120
Abstract....121
Introduction....121
Materials and Methods....123
Ants and experimental site....123
Data analysis....124
Agent-based modelling....125
Results....131
Myrmecia ants released on the tread-mill....131
Agent-based modelling....135
Testing model-predictions with Myrmecia....136
Discussion....137
Alignment matching and visual memories....141
Continuously integrating attractive and repellent views....143
Testing the model's prediction....145
Integration with path integration....146
Outlook....146
Acknowledgements....148
Supplementary figures....149
Repeated release experiment....152
Method....152
Results....153
Conclusion....156
**Proof of concept Antarium experiments....157**
Abstract....157
Introduction....157
Methods....157
Animal handling....158
Magnetic mounting....158
Data Analysis....159
Nomenclature....160
Preprocessing....162
Ball to scene coordinate transformation....163
Path to vector sequence transformation....164
Statistical methods....167
Experiment 1....170
Aim....170

Methods....170
Experimental design....170
Results....171
Discussion....172
Antarium deficiencies....172
Animal behaviour....173
Population variation....174
Antarium Improvements....174
Experiment 2....176
Aim....176
Methods....176
Experimental design....176
Scene preparation....178
Experimental paradigm....181
Trackball data preprocessing....182
Questions and Predictions....182
Results....184
Q1: Do the animals show orientation? Q3: Do the animals differentiate between scenes?....184
Q2. Do the animals react to rotations? Q3 Do the animals differentiate between scenes?....189
Q4: Do the animals behave at the various scenes as expected? Q5: Are there significant differences between animals?....195
Evening vs. morning experiments....208
Discussion....210
Experiment 3....213
Aim....213
Method....213
Results....214
Discussion....219
Conclusion....220
**Altered reality experiments....221**
Abstract....221
Introduction....221
Experiment 4....223
Introduction....223
Methods....225
3D model preparation....225
Animal handling....226
Results....227
Rotational image difference function analysis....233
Discussion....248
Experiment 5....251
Method....251
Results....255
Discussion....265
Summary....266

**The design of Antarium Mk-II.....269**
Colour matching and LED selection.....269
The geometry of the new Antarium.....285
LED arrangement.....289
Custom optics.....298
Changes to the electronics.....300
Polarisation.....301
Other considerations.....305
Manufacturing cost.....305
Other species.....306
Electrophysiology in the Antarium.....307
**Outlook.....309**
**Bibliography.....313**

# Introduction

The subject of this thesis is the use of virtual reality in investigating the visual navigation of ants. That raises the questions why study insect navigation and why do that with ants?

Visual navigation of insects is an active research topic. Insects have small brains and the acuity of their compound eyes is low compared to vertebrate eyes. To put it in perspective, a honeybee's brain contains about 1 million neurones (Menzel and Giurfa, 2001) whereas humans have somewhere between 85 billion (Williams and Herrup, 1988) and 120 billion (Walløe et al., 2014). The honeybee eye has an angular resolution of about 1.85° (Ryan et al., 2020) while for a human healthy vision is defined as the discrimination ability of one minute of arc (0.017°). Yet insects perform visual navigational tasks that would challenge any vertebrate, including humans. Unveiling the mechanisms insects employ to reliably extract navigational information from a low-resolution image using a seriously constrained computational machinery would no doubt heavily influence the fields of machine vision and robotics. Furthermore, insects show a certain level of intelligence, far beyond simple reflexes, the existence of which seems obvious when one monitors them navigating. Thus, neuro-ethological examination of their navigation might aid us understanding the emergence of intelligence from neural circuits and indeed move us a bit closer to answering the philosophical question “what is intelligence”.

Working with ants is beneficial for several reasons: Ants are abundant even in urban centres. They are easy to capture and to work with, both in their natural habitat and in the laboratory. Some species rely almost exclusively on their vision for navigation. As social insects, they have a fixed home that they return to, a navigational task that they regularly perform. They are pedestrian which eliminates all the complexity associated with a flying insect’s need to stabilise the flight, to control altitude and to compensate for wind. A walking insect’s task is less demanding, although, arguably, proximity to the ground introduces new challenges due to the visual and physical clutter that a flying insect can avoid. Nevertheless, in the context of navigation walking insects probably face less challenging problems than flying ones and, above all, are easier to study.

This chapter gives a brief introduction to the ants and the basic neurobiology of their visual system. A summary of our current knowledge about their visual navigation is also presented.

## Ants

Ants are hexapod insects that emerged from primitive wasps about 150 million years ago (Johnson et al., 2013). Ants form the large family of Formicidae within the order Hymenoptera. Around ten thousand species have been identified and it is believed that that figure represents only about 50 to 70% of all ant species. They are estimated to form around 15-20% of the biomass of all terrestrial animals (Schultz, 2000).

On the insect socialisation scale (Michener, 1969) ants are eusocial, the highest level of the scale. They live in colonies; the average colony size varies by species from a few tens of individuals to over 10 million (Beckers et al., 1989).

The colony has a queen who lays eggs. Species can be *monogynous*, having one queen. *Polygynous* species can have multiple queens, one or more of which lays eggs. There are a low number of males, hatching from unfertilised eggs, who have no role other than inseminating a future queen. Everything else is done by the workers, who are all sterile females. Their tasks include foraging for the colony; tending the brood; building, maintaining and cleaning the nest; as well as fighting intruders. Some species have a caste system where the workers are specialised for the various roles with obvious morphological differences, while in other species all workers look identical.

The experiments described in this thesis were performed using three ant species from the genus *Myrmecia*, in particular, *M. croslandi*, *M. pyriformis* and *M. midas*. These species were chosen because they are solitary, visual navigators; they cover all temporal niches; they had already identified nests and forging habitat at the Campus Field Station at the University; and there were already collected reference data about their behaviour.

All but one *Myrmecia* species are endemic only to Australia. The genus appeared relatively early in ant evolution (Crozier et al., 1995).

Identifying *Myrmecia* species is not an easy task. Several species look almost indistinguishable and show very similar, though not identical, behaviour. Ogata and Taylor (1991) describe the unique features of the various *Myrmecia*

species. Based on that article the Antwiki website published a step-by-step identification guide:

http://www.antwiki.org/wiki/Key_to_Australian_Myrmecia_Species

All three species that were used in the experiments live in underground nests. The adults feed on sugary substances but the brood is carnivorous. Thus, foragers collect sugar and they must also catch other insects to feed the larvae. The foragers collect the food on nearby *Eucalyptus* trees (5-15m from the nest). They also hunt for prey on the ground. Like all species in the genus, they have large, strong mandibles and a potent sting.

Images of several *Myrmecia* ants, including the ones used in the experiments, are shown in *Fig. 1*.

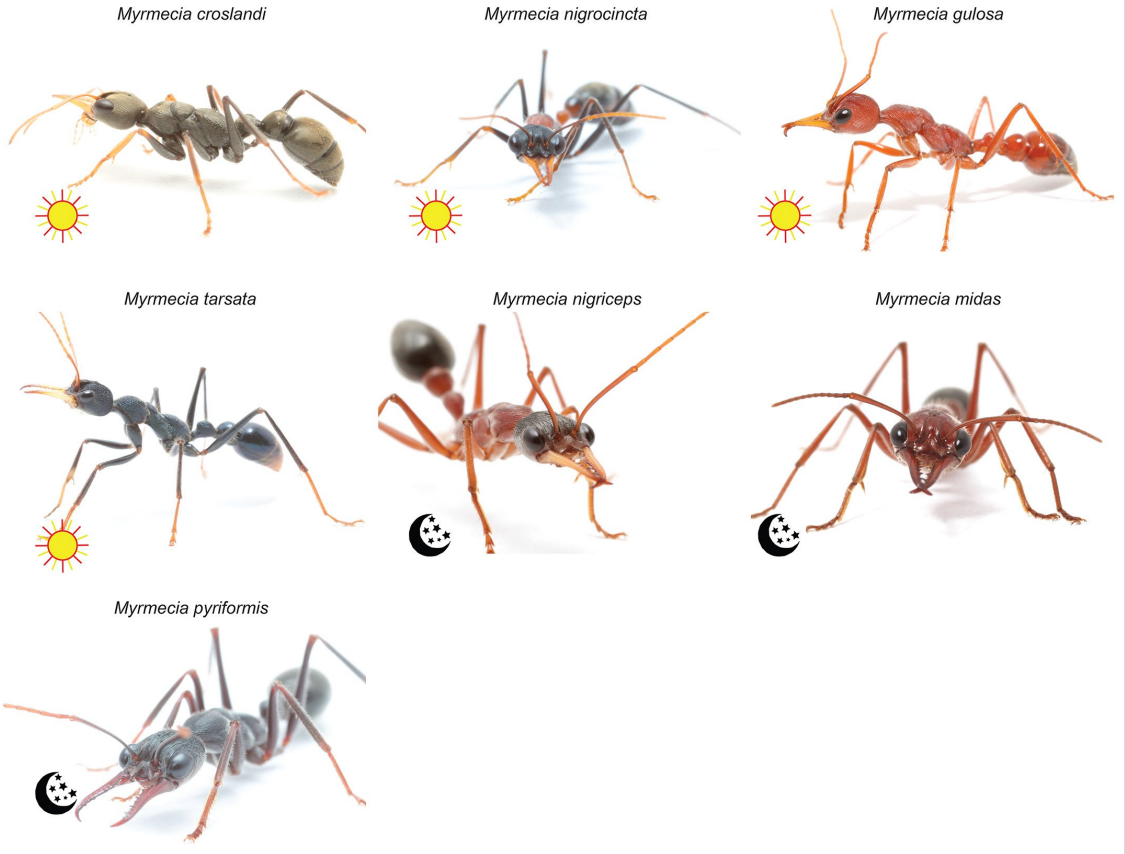


*Fig. 1: Pictures of Myrmecia ants*

*The little symbol bottom-left of the animal indicates whether the species is diurnal (sun) or crepuscular / nocturnal (moon and stars). Image adapted from (Sheehan et al., 2019)*

*M. croslandi* are diurnal. They start to forage in the late morning and return to the nest around mid-afternoon or earlier, if it is a really hot day (Jayatilaka et al., 2014). Workers are 10 to 15 mm long. These ants are active in the warm period of the year. In Canberra, where the experiments were performed, that is roughly from late October to early May. During the cold period the nest entrance is sealed and the ants are hibernating (Jayatilaka et al., 2011).

*M. pyriformis* are nocturnal. They usually start to forage just after sunset, when the ambient light level drops below a threshold (Narendra et al., 2010) and either return to the nest late in the evening or they spend the night on the foraging tree and return at dawn. They forage all year around, staying in the nest only if the ground temperature drops below a certain level (Reid et al., 2013). The major workers are 20-25 mm long.

*M. midas* are very similar to *M. pyriformis* both in appearance and behaviour. *M. midas* are slightly smaller and deep dark brown instead of pitch black. They are also somewhat less aggressive than *M. pyriformis*. However, *M. midas*, like *M. croslandi*, seal the nest for winter (personal observation over several years). In addition, they start to forage about half an hour before sunset (Freas et al., 2017).

## Ant vision

There are three major processing steps involved in forming behaviour based on the visual environment. First, light from the panorama is transformed to neuronal signals by two compound eyes. Then specialised neural structures process those signals and extract the useful information content of the scene. Finally, that data are then passed to the decision making and memory related parts of the brain, from where behaviour emerges.

This section provides an overview of these components. It barely scratches the surface, describing only what is necessary to understand how vision can affect behaviour. For a more in-depth look, the book "Arthropod Brains" (Strausfeld, 2012) provides a detailed (and entertaining) account of the neuroanatomy and function of arthropod brains, together with the history of their research. The insect neuroanatomy nomenclature is described in (Richter et al., 2010) and (Ito et al., 2014).

### The eyes

The two basic eye types are the *camera eye* and *compound eye*.

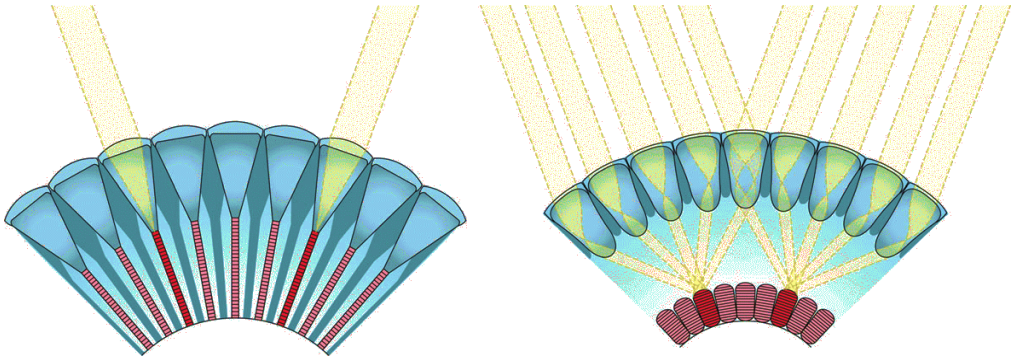

*Fig. 2: Fundamental compound eye designs*

*In the apposition compound eye (left) each lens collects the light from a single spatial direction and delivers it to a single photoreceptor unit. In superposition eyes (right) each lens collects light from several spatial directions; beams from multiple lenses, but from the same direction, converge on a single photoreceptor unit. Image adapted from (Warrant, 2017)*

Camera eyes have a single lens (or a pinhole) that projects an image to the retina covered by photoreceptors. All vertebrates and many invertebrates have camera eyes.

Compound eyes have an array of small lenses, usually arranged as a hexagonal lattice bent to form a convex surface. There are two fundamentally different designs, the *apposition* eye, often found on diurnal and crepuscular insects and the *superposition* eye, mostly used by nocturnal insects. These types are shown in *Fig. 2*. The *neural superposition* eye, not demonstrated in the figure, is a hybrid design, combining principles from the other two; it can be found in, for example, flies.

Evolution produced some extremely intriguing optical arrangements using curved mirrors, wave retardants, variable refraction index lenses and so on. The book "Animal Eyes" (Land and Nilsson, 2002) gives a comprehensive account of those solutions. Details of the compound eye designs and the inherent limitations to their visual acuity are examined in (Land, 1997).

Ants have apposition eyes so only that eye design will be described below.

In apposition compound eyes each lens collects light from a single spatial direction and focuses it on photoreceptors assigned to that lens. That is, each lens and photoreceptor unit provides a single pixel information. This eye type is very old. All the basic structures found in modern apposition compound eyes were already present in *Trilobite* eyes, over 500My ago (Schoenemann et al., 2017; Scholtz et al., 2019).

*Myrmecia* have a pair of apposition compound eyes, one each at the two sides of the head. In addition, they have three small camera eyes, called *ocelli*, on the dorsal-anterior aspect of their head, arranged in a triangular formation. The function of the ocelli is not yet fully understood; what is known about their role in navigation will be discussed later.

A brief description of the apposition compound eye can be found below. It should be noted that there are significant variations between the eyes of different orders and even families of insects. These differences will not be detailed.

***Geometry***

The compound eyes of *Myrmecia* are slightly distorted hemispheres protruding from the lateral aspects of the head. The surface of the hemisphere is covered with a hexagonal matrix, with each facet being a few tens of micrometres in diameter. Each such facet is the distal surface of an *ommatidium*, a unit delivering one pixel of visual information to the animal. *M. croslandi* workers and *M. pyriformis* minor workers have about 2,350 facets in each eye while *M. pyriformis* major workers have around 3,600 (Greiner et al., 2007; Narendra et al., 2011). No data are available for *M. midas*, but they are generally very similar to *M. pyriformis*, so it is a reasonable assumption that their eyes are also similar.

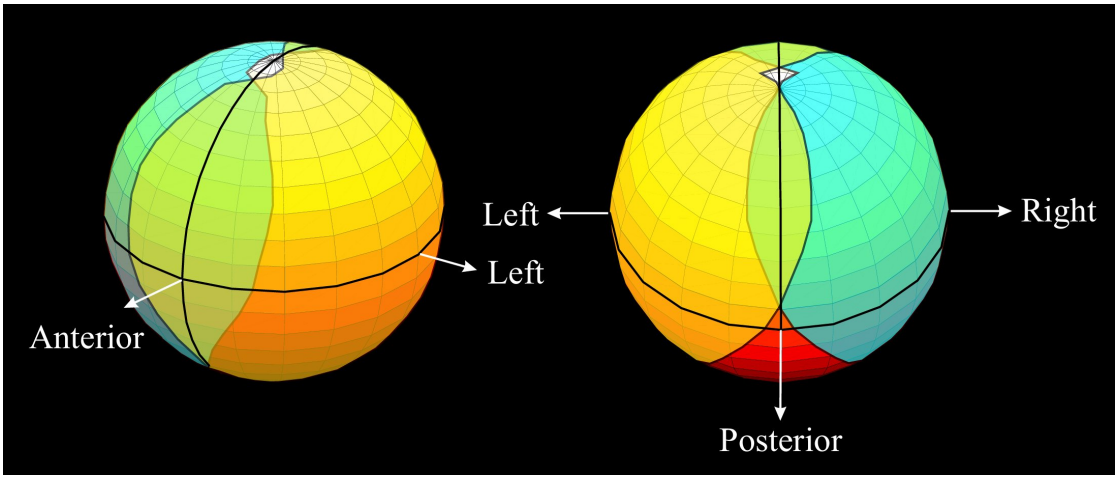


*Fig. 3: Visual field of M. pyriformis*

*The yellow and cyan colours represent the fields of the left and right eyes, respectively. The anterior and posterior green areas indicate the binocular fields. The posterior-ventral area (marked red) and a small area at the anterior-ventral aspect (not visible on the image) are blind spots, occluded by the body. The white area at the dorsal aspect is also a blind spot, there are no ommatidia that look exactly at the zenith. Courtesy of Jochen Zeil, based on in vivo optical measurements by Eric Warrant.*

Due to the geometry of the eyes, the animal has a nearly spherical visual field, with only a small ventral section being occluded by the body. The anterior visual fields of the two eyes overlap, giving the animal binocular vision within that area. That is important for capturing prey (Via, 1977). The visual field for *M. pyriformis* has been mapped, as shown in *Fig 3*.

The metric of the resolving capability of the compound eye is the *inter-ommatidial angle*, the angular difference between the optical axes of adjacent ommatidia. For a perfect hemisphere and identical sized hexagonal facets the inter-ommatidial angle can easily be derived from the number of facets. The sphere has a radius of R and the hexagons have a 'diameter' (the distance between two parallel sides) of D. The surface area of the sphere is $4\pi R^2$ while that of the hexagon is $\frac{\sqrt{3}}{2}D^2$ and so the number of hexagons that tile the sphere will be

$$n=\frac{4\pi R^2}{\frac{\sqrt{3}}{2}D^2}=\frac{8\pi}{\sqrt{3}}\cdot\frac{R^2}{D^2}$$

Solving for D we get:

$$D=\sqrt{\frac{8\pi R^2}{\sqrt{3}n}}=\sqrt{\frac{8\pi}{\sqrt{3}}}\cdot\frac{R}{\sqrt{n}}\approx 3.81\cdot\frac{R}{\sqrt{n}}$$

The equator of the sphere is $2\pi R$ long, spanning 360°. The D diameter hexagons therefore are

$$\Delta\Phi=\frac{360\,D}{2\pi R}=\frac{360}{2\pi}\cdot\frac{3.81\frac{R}{\sqrt{n}}}{R}=\frac{218.3}{\sqrt{n}}$$

degrees apart. The number *n* in the above equation refers to the number of facets over the full sphere; normally the number of ommatidia is given for one eye, that is, a hemisphere. Correcting for that, the final equation for the average inter-ommatidial angle is

$$\Delta\Phi=\frac{154.36}{\sqrt{n}}$$

where *n* is the number of facets per eye and the result is in degrees.

For the 3,600 ommatidia per eye of *M. pyriformis* the equation gives 2.57°. But since the eye is not a perfect hemisphere, the angle varies across the visual field. Where the eye has larger curvature, the inter-ommatidial angle is larger, while a locally flattened region offers better resolution. For *M. pyriformis* the

inter-ommatidial angle varies between 1.1° (lateral field) and 3.0° (anterior field) (Reid, 2010). For a more in-depth analysis of the visual acuity of the compound eye, see (Kirschfeld, 1976) and (Land, 1997).

Another important metric is the *acceptance angle*, the half-width of the cone of directions from where light can reach the photoreceptors. Land (1997) published a paper where he collated all available measurements to-date and concluded that diurnal insects tend to under-sample the image, which results in higher contrast (higher than human vision) while nocturnal insects are more likely to over-sample. Unfortunately, the only Hymenoptera for which data were available for this review was the honeybee. More recently Schwarz *et al* (2011a) reported that desert ants have an acceptance angle smaller than the inter-ommatidial angle. No data are available for *Myrmecia.*

A resolution in the degrees range might seem inadequate compared to the 0.017° resolution of human vision. However, it will be shown later that ants use methods of navigation where the low resolution of the eye is not detrimental and, within certain limits, is even advantageous.

***Anatomy***

The distal end of the ommatidia (i.e. the surface of the eye) are covered by a hard, thin, transparent protective layer called the *cornea*. The cornea is part of the optical assembly of the ommatidium. Directly below it is the *cuticular lens*, a transparent bi-convex structure. The cornea and the cuticular lens together are usually just referred to as the lens. The lens is not biologically active. Below the lens is the *crystalline cone*, a transparent material secreted by surrounding cells called the *Semper cells*. Light incident to the cornea is focused by that assembly and leaves the crystalline cone as a narrow (few microns diameter) beam at the centre of the ommatidium.

A detailed examination and mathematical analysis of the light collection properties of the optical assembly of the honeybee, *Apis mellifera*, can be found in (Varela, 1970a). The optical assembly is shown on *Fig. 4*. The paper's scope is limited to refractive optics. Since its publication it has been shown that, due to the small size of the internal structures of the ommatidium, diffraction and waveguide phenomena also must be taken account. Those calculations were carried out for flies (Stavenga and van Hateren, 1991; Stavenga, 2003). No similar analysis has been published for *Myrmecia* ommatidia.

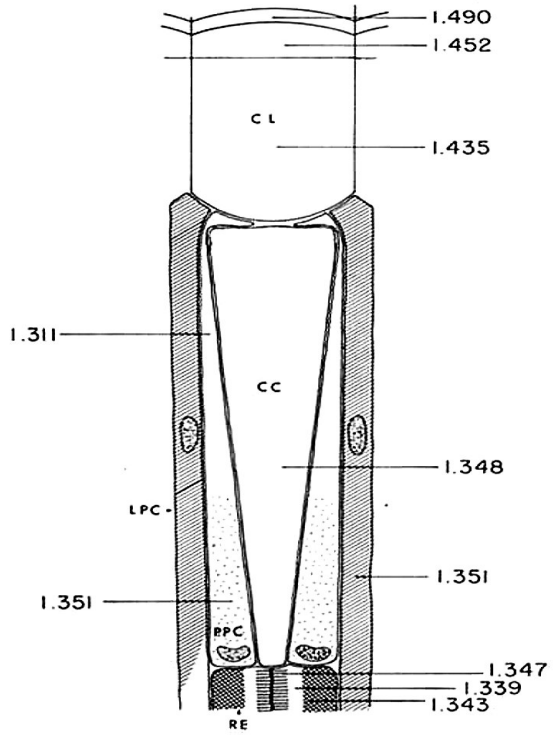


*Fig. 4: Optical assembly of an ommatidium*

*The numbers are the refractive indices of the indicated structures, data are for the honeybee. CL:cuticular lens, CC:crystalline cone, LPC:long pigment cells (separating ommatidia), PPC:primary pigment cells (controlling light intake), RE:retinula cells (photoreceptors). Adapted from (Varela, 1970a)*

The amount of light exiting the optical assembly is controlled by the primary pigment cells. They can narrow the diameter of the light beam leaving the optical assembly, thus limit the amount of light entering the rhabdom.

The degree of intervention that the pigment cells can perform depends on the temporal habitat of the animal. Diurnal species show markedly less control over the light amount than nocturnal species (Narendra et al., 2016). Nocturnal species have larger rhabdoms (Narendra et al., 2017), needed to collect the feeble night light. Those rhabdoms would be overwhelmed by daytime light levels without an effective pupillary mechanism.

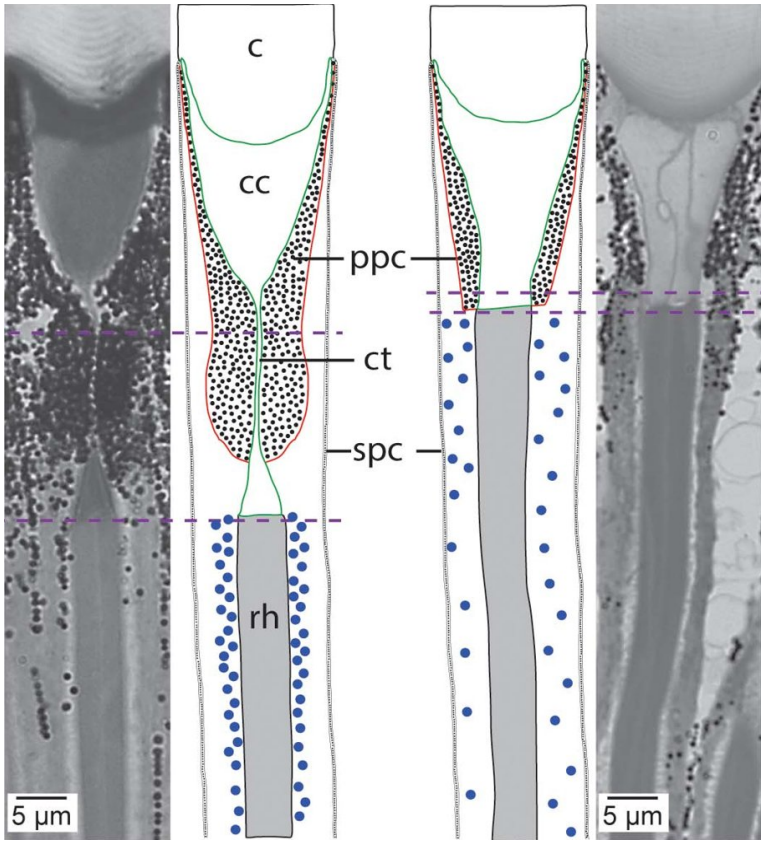


*Fig. 5: Pupillary mechanism*

*The left hand side shows the microscopic and schematic arrangement under bright light conditions while the right hand side shows a dark-adapted ommatidium. C:cuticular lens, cc:crystalline cone, ct:crystalline cone tract, ppc:primary pigment cells, spc:secondary pigment cells, rh: rhabdom. Adapted from (Narendra et al., 2013a).*

As shown in *Fig. 5*, under bright conditions the primary pigments cells invade the crystalline cone tract and pigment migrates to the intrusion, constricting the light beam. In a dark adapted ommatidium the pigment cells move away from the crystalline cone tract and allow all the light collected by the optical apparatus to exit (Narendra et al., 2013a, 2016).

It should be noted that since the animal has spherical vision, one ommatidium always looks directly at the sun during the day or the moon during the night (both the Sun and the Moon have an apparent diameter of about 0.5°). Other ommatidia look at other parts of the scenery, possibly dark shadows. Therefore the qualifiers of ‘bright conditions’ and ‘dark conditions’ apply not only to the overall scene but also on a per-ommatidium basis. It is not known whether the adaptation is strictly local (each ommatidium works independently) or ommatidia can affect each other. In addition, the effect of dark/light adaptation to visual information extraction has not been studied.

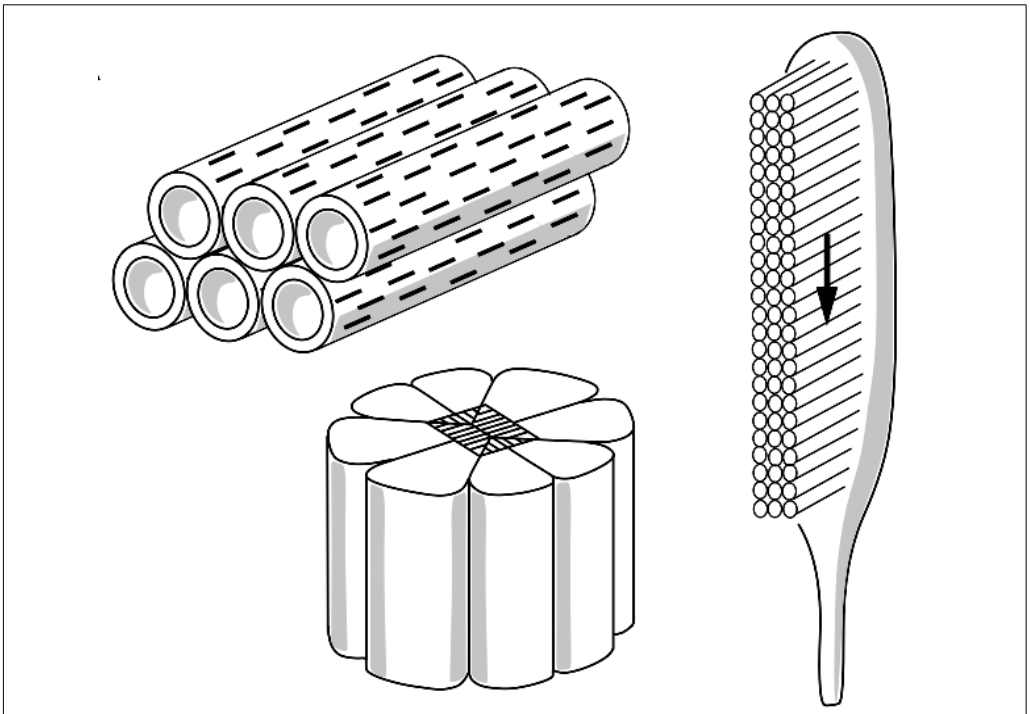

*Fig. 6: Photosensory apparatus*

*The centre image shows the retinular cells surrounding the rhabdom. On the right is a single photoreceptor with the microvilli that it sends into the rhabdom; the microvilli from a single cell are the rhabdomere. Top left shows that the elongated opsin molecules are aligned with the longitudinal axis of the microvillus the membrane of which they are embedded in. Adapted from (Labhart, 2016).*

Below the crystalline cone is the photosensory apparatus.

Photoreceptor or *retinular* cells are assembled in a ring formation; the number of retinular cells is usually between 6 and 9, depending on species. These cells

are elongated, thus forming a long, thin tube around the longitudinal axis of the ommatidium. The inside cavity of the tube is called the *rhabdom*. The light exiting the crystalline cone enters the rhabdom, which can function either as a light guide or as a waveguide.

The rhabdom is not empty space. The retinular cells send many thin (around 50nm in diameter) finger-like parallel processes, called *microvilli*, into the light beam cavity. The collection of all microvilli from a single retinular cell is called a *rhabdomere.*

The surface of every microvillus is packed with photosensitive units. A complex transmembrane protein, *opsin*, is bound to a photosensitive molecule, the *retinal*[1] (together they are called *rhodopsin*). When the retinal absorbs a photon with the correct wavelength (which depends on minute differences in the opsin molecule that contains the retinal), it will undergo a configuration change. That causes a configuration change in the opsin, which, in turn, triggers a complex chemical cascade leading to the opening of ion channels. The consequent depolarisation ultimately results in neurotransmitter release at the base (proximal end) of the cell (Fain et al., 2010; Sterling and Laughlin, 2015; Honkanen et al., 2017). That process is called *phototransduction*.

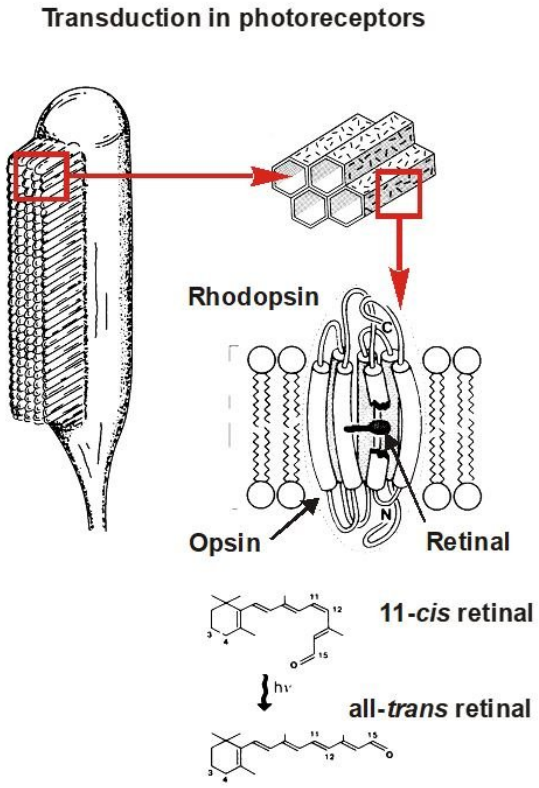


*Fig. 7: Phototransduction*

*Only the first step of the cascade, the configuration change of the retinal molecule is shown. Adapted from (Stavenga and Hardie, 2011)*

1 *In invertebrates small chemical variations of the molecule occur, for example* Drosophila *uses the 3-hydroxy-retinal variant (Arshavsky, 2010).*

The components of the photosensory apparatus are shown in *Fig. 6* while *Fig. 7* illustrates phototransduction.

The rhabdom can be *fused* or *open*. In the former, the rhabdomeres completely fill the space inside the rhabdom and all rhabdomeres have the same optical axis as that of the optical assembly of the ommatidium. In the latter, rhabdomeres remain separate, forming several light guides, each with a slightly different optical axis. Hymenoptera have fused rhabdoms. Their properties are examined in (Snyder et al., 1973).

Both the optical and the photodetector assembly are surrounded by (separate) pigment cells. Those cells optically isolate adjacent ommatidia from each other.

The quantum efficiency, noise minimisation and energetic optimisation of insect photoreceptors and a comparison with vertebrate ones is discussed in detail in the book "Principles of Neural Deign" (Sterling and Laughlin, 2015).

The retinal molecules in the rhodopsin are elongated and they are sensitive to the polarisation direction of light. Polarisation is a property of transverse waves, including electromagnetic radiation. The details of the phenomenon are far beyond the scope of this thesis, involving a lot of mathematics. However, for our purposes it is sufficient to know that electromagnetic radiation, which includes light, is an oscillating electric field and an oscillating magnetic field which are orthogonal to each other (and both are orthogonal to the direction of energy propagation). If the polarisation is linear, then the electrical field (E) and the magnetic field (B) are contained in two orthogonal planes, shown in *Fig. 8.*

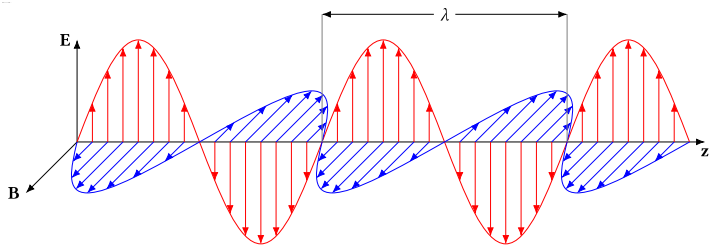


*Fig. 8: Linear polarisation of electromagnetic waves*

*The radiation, with wavelength λ, propagates along the z axis. The oscillating electrical field is parallel to the E axis while the magnetic field with the B axis.*



The actual direction of the E axis is not fixed. If we rotate both the E and B axes around the z axis, the rule that all three axes should be orthogonal to each other will not be violated. Mathematically, oscillating fields can be generated by rotating vectors, thus the electrical component of the propagating electromagnetic wave is called the *E-vector.* Its direction is that of the E axis. In a beam of light each individual photon can have its E-vector point to an arbitrary direction: the light is not polarised. However, if all E-vectors point in the same direction, the light is said to be polarised.

A photon with its E-vector parallel with the longitudinal axis of the retinal molecule has a significantly higher probability of triggering the configuration change than a photon with an orthogonal E-vector.

As shown in *Fig. 6*, the retinal molecules are likely to be oriented parallel to the microvillus they are embedded in. Therefore, if all microvilli of a rhabdomere are parallel to each other, the entire rhabdomere is polarisation sensitive. In some cases (discussed later) that is advantageous, but in the general case polarisation sensitivity should be avoided. That is achieved by twisting the retinular cells around the rhabdom and thus rotating the microvillar orientation of the rhabdomere along the optical axis of the rhabdom (Menzel and Blakers, 1975; Meyer and Domanico, 1999).

The photoreceptor cells use multiple opsins which are sensitive to different parts of the spectrum. In most cases a single retinular cell expresses only one opsin. In the case of *Myrmecia* the retinular cells are sensitive to green, blue and ultraviolet (Ogawa et al., 2015), shown in *Fig. 9*.

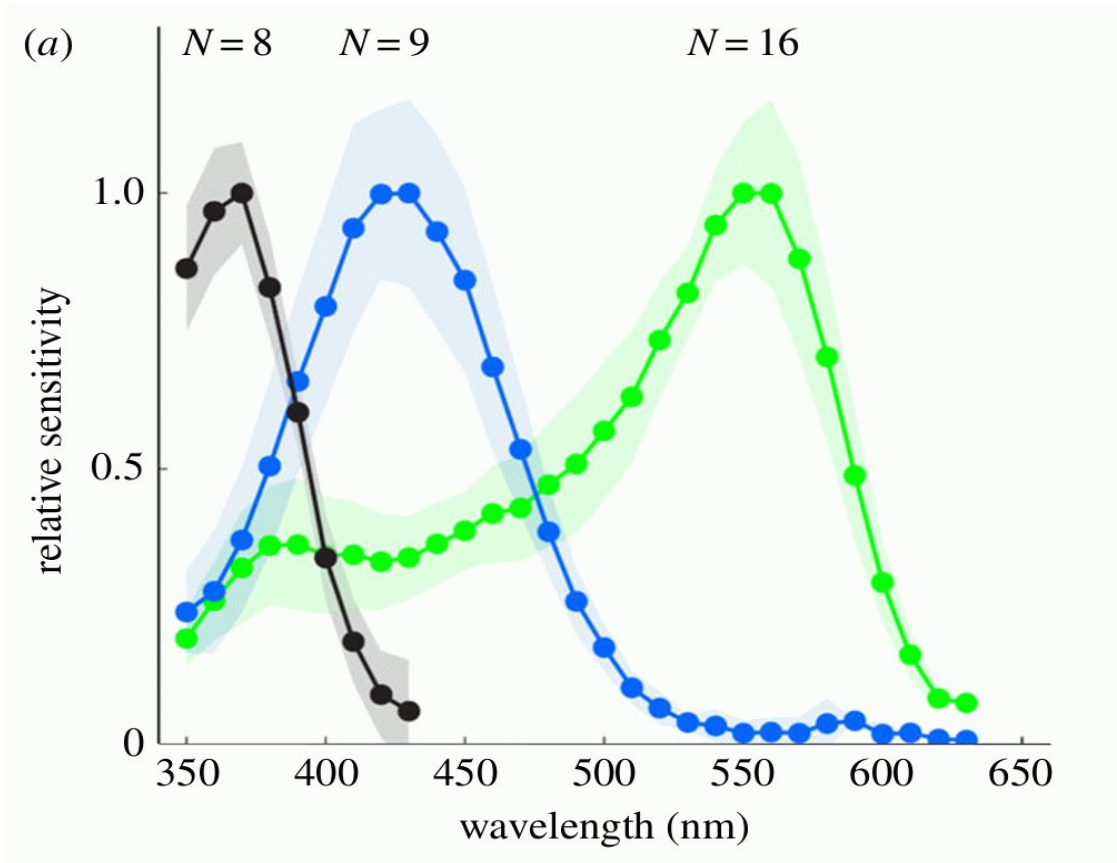


*Fig. 9: Spectral sensitivities of M. vindex photoreceptors*

*Electrophysiologically obtained spectral sensitivity curves. N indicates the number of animals used to collect the data for the particular curve. The peak sensitivities are 370nm for the UV (black curve), 450nm for the blue (blue curve) and 550nm for the green (green curve) sensitive receptors. The shaded areas around the curves indicate the standard deviation. Adapted from (Ogawa et al., 2015).*

It is not clear how the potential ability for colour vision contributes to visual navigation. It is known, however, and will be discussed later, that UV contrast is crucial to obtain both visual compass and landmark panorama information (Barta and Horváth, 2004; Möller, 2002; Schultheiss et al., 2016a; Stone et al., 2014). The number of photoreceptor cells in an ommatidium depends on the order, possibly even on a lower taxonomical level. In Hymenoptera the sweat

bee (Greiner et al., 2004), honeybee (Ribi, 1975), the desert ant (Labhart and Meyer, 1999) and *Myrmecia* (Menzel and Blakers, 1975) are all known to have nine retinular cells.

For *Myrmecia*, Greiner e*t al* (2007) examined the eyes of several species and concluded that the optical sensitivity of the eyes and the number of ommatidia are not proportional to body size but clearly related to the temporal niche the animal lives in: nocturnal animals tend to have larger and more sensitive eyes. The same temporal niche dependent differences were detected by electron microscope examination of the rhabdoms of several *Myrmecia* species (Narendra et al., 2011, 2017).

Sheehan *et al* (2019) examined the volume of brain centres of seven *Myrmecia* species and concluded that diurnal species invested more in their optic lobes (the primary visual processing centres) than nocturnal species, on the other hand the latter invested more into the higher level processing centres. These centres will be discussed shortly.

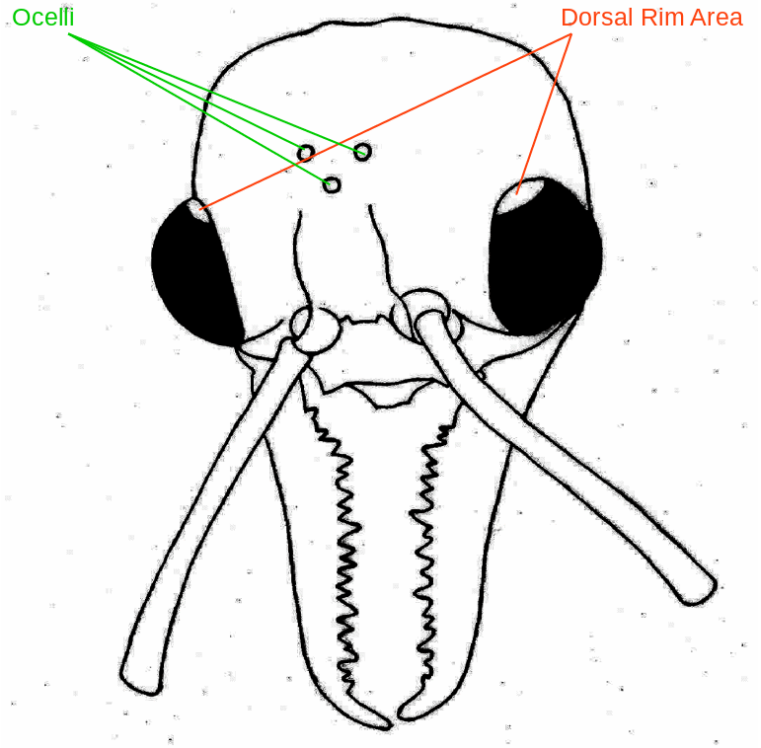


*Fig. 10: The dorsal rim area and the ocelli of M. pyriformis*

*Drawing courtesy of Ladina Ribi*

***Dorsal Rim Area***

For some insects, including *Myrmecia*, in the dorsal region of the eye, in the so-called *dorsal rim area* (DRA) the microvilli from different rhabdomeres are all either parallel or orthogonal to each other and they all carry the same opsin (UV sensitive, in case of Hymenoptera). Furthermore, the retinular cells do not twist around the rhabdom, therefore the microvilli are aligned the same way all through the rhabdomere. Since photoreceptors with microvilli aligned in parallel are sensitive to the direction of polarisation, with two orthogonal sets the animal can measure the polarisation angle of the incident UV light (Labhart, 2016).

These are specialised ommatidia. They often have a wide optical acceptance angle (basically, they are slightly out of focus). The shape of the rhabdom might be different from that of 'normal' ommatidia and there can be other morphological differences of the retinular cells as well (Aepli et al., 1985; Labhart, 1980; Labhart and Meyer, 1999; Meyer and Domanico, 1999), reviewed in (Zeil et al., 2014a). The DRA is shown in *Fig. 10*.

Although many insects measure the polarisation of the UV, some use the blue (Kinoshita et al., 2007) or even the green (Labhart and Meyer, 1999) part of the spectrum.

***The ocelli***

Like many insects, *Myrmecia* have not only compound eyes, but a set of camera eyes as well. These are called the *ocelli*. In the case of *Myrmecia* there are three of them, arranged as a triangle on the dorsal-anterior aspect of the head (see *Fig. 10*).

The role of the ocelli is not clearly understood, but they are believed to be involved in navigation and, for flying insects, in flight stabilisation. In this introduction only the navigational context will be examined briefly and even that will be limited to ants. It should be noted that not all ant species have ocelli.

The basic optical, anatomical and neural blueprint variations of the ocelli are discussed in (Mizunami, 1995). It seems that the general organisation of the ocelli is well preserved across *Hymenoptera* (Ribi and Zeil, 2018). An important aspect of that preserved structure is that the ocellar retina is bipartite, with well

defined ventral and dorsal regions. For certain species the ventral region, corresponding to the dorsal visual field, is composed of rhabdoms with aligned microvilli, which are therefore sensitive to the polarisation of light. For other species there is no or limited alignment of the microvilli.

The first experimental evidence for ocelli supplying compass information[2] for an ant was published almost four decades ago by Fent and Wehner (1985). A more recent experiment by Schwarz *et al* (2011b) confirmed those findings. It should be noted that both experiments used desert ants (*Cataglyphis bicolor* and *Melaphorus bagoti*, respectively) which live in visually parsimonious environments (especially *C. bicolor*) and heavily rely on the sky compass information. A review of polarisation vision by Zeil *et al* (2014a) points out that while the above mentioned experiments prove that the ocelli can supply compass information, the ocelli of the animals involved have not been studied in detail, so we don't know whether those species possess polarisation sensitive, oriented microvilli in their ocellar retina. Narendra and Ribi (2017) examined the ocelli of several *Myrmecia* species (including diurnal, crepuscular and nocturnal) and concluded that their retinal organisations were not particularly suited to extract polarisation information from skylight.

Little data are available on the spectral sensitivity of the ocelli in the literature. However, Ogawa *et al* (2017) performed anatomical, physiological and behavioural experiments with honeybee ocelli. They concluded that the bee ocelli photoreceptors are sensitive to UV (360nm) and blue-green (500nm). The UV receptors have shown higher polarisation sensitivity than the blue-green ones. Furthermore, the ventral visual field of the ocelli has been shown to be most sensitive to vertically polarised light while the dorsal field did not show a preferential E-vector orientation.

It thus remains an open question how input from the ocelli may augment the visual information that the compound eyes provide to the navigational system of the animal. Based on experiments with *M. bagoti*, Schwarz *et al* (2011c) suggested that it was possible that the ocelli are part of a separate navigation

2 *Obtaining compass information from the polarisation pattern of the sky will be discussed on page 60.*

mechanism that runs in parallel with, and independently of, the one mediated by the compound eyes.

Although the ocelli seem to play a role in visual navigation, that role has not yet been clearly identified. Thus, in the rest of this thesis the ocelli will be mentioned where relevant, but otherwise will not be examined or discussed in detail.

## The optic lobes

Like with the apposition compound eye itself, the grand plan of the neural machinery processing the visual information is the same throughout the insect world. However, minor variations with significantly different performance have evolved, reflecting the unique needs of the individual species. The description below is simplified, omitting much detail. Its aim is only to illustrate the neural machinery at the ant's disposal. *Fig. 11* shows the brain regions that will be discussed.

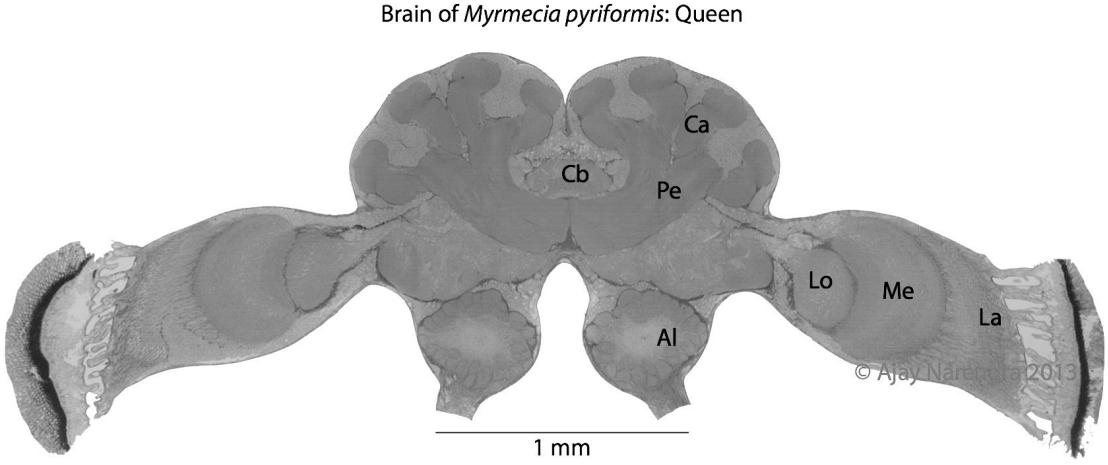


*Fig. 11: Brain regions involved in visual navigation*

*On the sides are the eyes. Further details are in the text. The area markings are: La:lamina, Me:medulla, Lo:lobula complex, Ca:mushroom body calyx, Pe:mushroom body peduncle, Cb:central body. The area marked Al:antennal lobe belongs to olfaction and is not discussed in the text. Image courtesy of Ajay Narendra.*

The collective name for all photoreceptor cells is the *retina*. The raw visual information from the retina is processed by the *optic lobes* of the animal's brain. The optic lobe consists of four *neuropils*[3] (Strausfeld, 2012). Each neuropil contains repeating units which are *retinotopic*, that is, each unit corresponds to exactly one ommatidium in the eye and these units have the same spatial relationship to each other as that of the ommatidia[4]. These units are called *cartridges* or *columns*. They contain several neurones each. The cartridges are separated from each other by glial cells. The columns receive retinotopic input and provide retinotopic output. Their neurones make local interconnections between each other. The neuropils also contain *tangential* neurones. These are not part of the columnar arrangement; rather, they make connections to several columns. There are many different types with a wide variety in the number of columns they contact. Many of them are *amacrine*, with no axon and as such with no clear direction of the flow of information through the cell. Tangential cells connect to each other as well, forming a very complex information processing and mediating network between the columns. The arborisations of the tangential neurones are located in well distinguishable *layers* or *strata* of the neuropil. In addition, output neurones, gathering information from both columnar and tangential neurones, send axons to higher level processing centres of the brain. There are fewer output neurones than ommatidia, but their axonal tract is still representative of the retinal map, only sub-sampled. *Fig. 12* shows a highly simplified connection diagram of the optic lobe (columnar neurones not drawn).

3 *A neuropil is a region which is very densely packed with synapses between neurones. In insects the cell bodies of the neurones often remain outside of this region and only the axonal or dendritic processes are sent into the neuropil. In older literature the neuropils were often called ganglia but it is no longer the case; the word ganglion is now reserved to specific neuropils of the animals.*

4 *With the exception of the lamina, the retinotopic organisation has not been explicitly demonstrated, but it is generally accepted to be true for the neuropils in the optic lobe.*

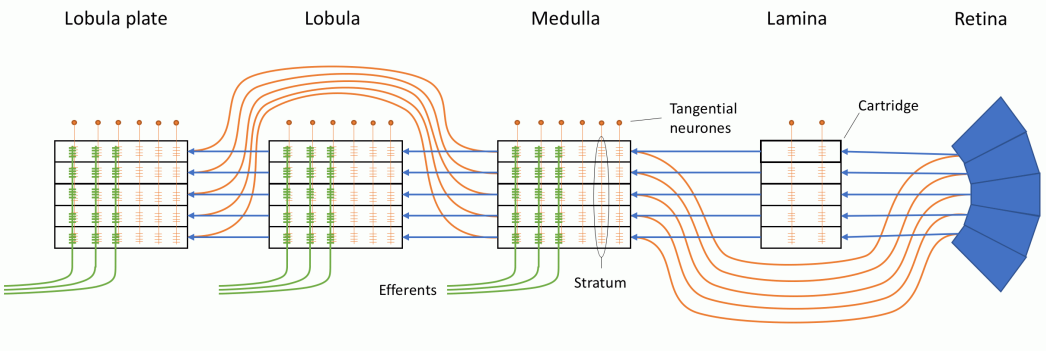


*Fig. 12: Simplified schematics of the optic lobe*

## Retina

The retina is the collective name for all photoreceptor cells in all the ommatidia. Some photoreceptors of an ommatidium synapse to the neurones in the first optic lobe neuropil, the lamina, via *short visual fibres*. Others send axons (the *long visual fibres*) across the lamina and synapse to the neurones in the second neuropil, the medulla. It is known for *Drosophila* that the green sensitive receptors terminate in the lamina while the blue and UV sensitive ones in the medulla. No such information is available for *Myrmecia*.

## Lamina

The sensory data from the retina first enter the *lamina*. It has a relatively simple structure and for some species its organisation is known in detail; for example (Armett-Kibel et al., 1977) describes the dragonfly lamina, (Hamanaka et al., 2013) that of the butterfly, while the bee lamina is examined in (Greiner et al., 2004; Ribi, 1975; Rusanen et al., 2017; Varela, 1970b). The lamina connectome for *Drosophila* has been mapped and is described in, for example (Rivera-Alba et al., 2011). That paper established that there are 5 columnar neurones in each lamina cartridge receiving input from some of the photoreceptors, while (Ribi and Scheel, 1981) identified 4 lamina columnar neurones for the honeybee.

Electrophysiological recordings suggests that the lamina neurones perform roles such as noise filtering and contrast equalisation (Laughlin, 1981; Sterling and Laughlin, 2015). They also might process colour information (Chen et al., 2020) and condition signals for motion detection (Tuthill et al., 2013).

Electrophysiology on hawkmoth lamina monopolar cells indicate that the lamina can perform light level dependent spatial integration, that is, deliver high resolution image in bright conditions and trade in resolution for sensitivity as the light level drops (Stöckl et al., 2020).

## Medulla

The medulla is a complex neuropil with many strata and morphologically different cells e.g. (Ribi and Scheel, 1981). It is a very dense neuropil so performing electrophysiological measurements on it is hard, although not impossible. Recordings have been made in flies, e.g. (DeVoe and Ockleford, 1976), crickets e.g. (Honegger, 1978; Yukizane et al., 2002), locust, e.g. (Osorio, 1986). In fruit flies, functional imaging is also available as a tool to monitor neural activity in the medulla (and other neuropils) e.g. (Hardcastle et al., 2021).

In Hymenoptera, the axons of the medulla output neurones form the *anterior-superior optic tract* (ASOT), one of the visual pathways to the brain.

There are also connections from the medulla to the central complex (discussed shortly).

A century ago, based on his comparative neuroanatomical studies, Ramón y Cajal argued that the organisation of the insect lamina and medulla together is very similar to that of the neural circuitry found in the vertebrate retina (Sanes and Zipursky, 2010). Unfortunately, the operation of the vertebrate retina layers is still far from completely understood (MacNeil and Masland, 1998; Masland and Raviola, 2000; Masland, 2001; Masland and Martin, 2007). As clearly spelt out in a review paper by Marsland (2012), the vertebrate retina does not send raw pixel information to the higher level brain centres. Rather, it uses local nonlinear neural filter circuits which retinotopicly extract at least 20 local features from the raw image. The extraction circuits have feature dependent spatial sampling frequencies and receptive field sizes. About half of the extracted features have been identified; such as local average intensity, local motion, local colour contrast and so on. The other half, at the time of writing this

thesis, remains a mystery. The important point is that the vertebrate optic nerve carries the extracted features and not the raw pixels.

If Cajal's observation is correct and the initial stages of the image processing mechanisms employed by vertebrates and insects are indeed similar, then we can assume that the medulla, too, extracts local features of the image and these retinotopic feature sets, rather than the actual pixels, are sent to the insect's brain. It is not yet known what those extracted features are, although we do know that colour (Morante and Desplan, 2008) and motion (Spalthoff et al., 2012) related information are among them.

## Lobula complex

The lobula complex is the final processing stage of the optic lobe. For some orders, such as Diptera, it contains two anatomically distinct neuropils, the *lobula* and the *lobular plate*. For other orders, including Hymenoptera, these two neuropils are anatomically not separable (Ribi and Scheel, 1981) and the whole neuropil is simply called the lobula.

The lobula receives retinotopic input from the medulla. It has fewer strata than the medulla. Its exact processing role is not yet known.

The lobular plate receives retinotopic input from both the lobula and the medulla. It is known, mostly from studies performed on flies, to extract motion information, both whole scene (optic flow) and that of objects moving through the visual field, reviewed in (Borst, 2009; Krapp et al., 1998, 2001; Borst and Egelhaaf, 1989; Borst and Haag, 2002; Dan et al., 2018). There is no literature available about the organisation and function of the lobula in ants, but for dragonflies it is known that, among other things, it extracts orientation information (O'Carroll, 1993).

## Higher level brain centres

The optic lobes pre-processed the visual scene and, we can assume, extracted the useful information from it while throwing away unimportant details, noise and general visual clutter. The result then needs to be integrated with other sensory clues, it should be associated with known scenarios, any new information contained in it should be learnt, and, of course, the animal must make decisions based on it.

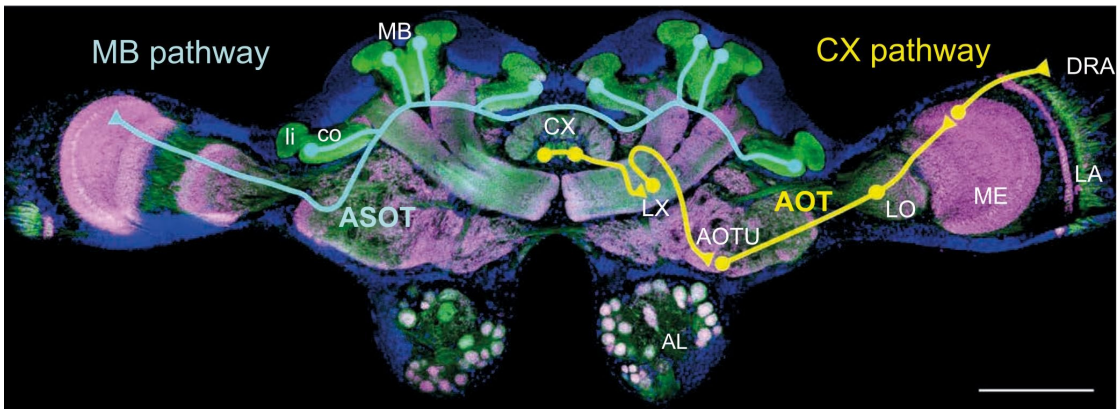


*Fig. 13: The two main visual pathways*

*The cyan path on the left shows that visual information goes from the medulla directly to the mushroom body calyces. It also shows that the visual information from the left and right eyes are combined, and the mushroom bodies of both brain hemispheres receive the visual data from both eyes.*

*The yellow track on the right shows that data collected from the dorsal rim (UV compass information), the medulla (assumed to be colour and motion information), lobula complex (motion) is delivered to the central complex.*

*Brain of the desert ant Cataglyphis fortis. ASOT:anterior superior optic tract, AOT:anterior optic tract, AOTU:anterior optic tubercule, LA:lamina, ME:medulla, LO:lobula complex, DRA:dorsal rim area, MB:mushroom body, li:mushroom body lip, co:mushroom body collar, LX:lateral complex, CX:central complex, AL:antenna lobe. The scale bar at the bottom-right is 0.2mm. Adapted from (Rössler, 2019) .*

All that happens in higher level brain centres. From a visual navigation point of view, the most important two of those are the *mushroom bodies* and the *central complex*. There are two main separate pathways in ants delivering visual information to these centres (Rössler, 2019).

Output neurones of the medulla deliver data to the mushroom bodies. Both mushroom bodies receive the visual information from both eyes. As it was mentioned earlier, we don't exactly know what that information exactly is.

The other pathway is more complex. The medulla processes the dorsal rim area input, which is believed to deliver sky polarisation information. The processed result, together with output from the lobula, connects to neurones in the anterior optic tubercule. Those, in turn, send outputs to the lateral complex and the lateral complex then delivers the information to the central complex. Not shown on the image, but it is also known that the central complex receives information from the ocelli. Furthermore, electrophysiological measurements on bees indicate that the central complex also receives colour information from the medulla (Paulk et al., 2009).

Recently it was discovered that, at least for flies, there is a direct data pathway from the mushroom bodies into the central complex, and an indirect pathway the other way, which will be discussed shortly.

It is my personal belief, and I would like to emphasise that I cannot provide corroborating literature, that the central complex processes directional cues, motion measurements (and control) and in general is responsible to know where the animal is and where she is heading. However, it does it without analysing the actual visual scene. The mushroom bodies, on the other hand, take the data describing the visual surroundings of the animal and try to compare it to known scenes, recognising familiar views and recall information associated with them. Those associated information bits are what the MB sends to the central complex, which then makes the final decision about what the animal should do (i.e. her observable behaviour), including instructing the mushroom bodies to learn a scene and associate certain navigation data with it.

## Central complex

As its name implies, the central complex (CX) is an assembly of interconnected neuropils, located roughly at the centre of the brain. It is an evolutionarily old, well preserved brain region (Homberg, 2008). For completeness, but without going into any detail, in insects its components are the *fan shaped body*, the *ellipsoid body*, the *protocerebral bridge* and the *noduli*. Like the optic lobe neuropils, these components have a columnar arrangement and many layers of tangential neurones, which form an intricate network connecting the components to each other (Hadeln et al., 2020).

The central complex receives input from many sensory systems (Varga et al., 2017), including proprioception. It has long been assumed that the central complex receives information from the mushroom bodies, even though there was no evidence of a direct MB to CX pathway (Collett and Collett, 2018). However, recent neuroanatomical studies did find direct MB→CX and indirect CX→MB connections for *Drosophila* (Li et al., 2020a). The CX is thought to control behaviour; and its outputs, after some more processing, eventually reach the motor control ganglia of the animal (Heinze and Homberg, 2008; Varga et al., 2017).

The neuroanatomy of the central complex throughout the insect world is discussed in (Strausfeld, 2012) and (Heinze and Pfeiffer, 2018). Perhaps the most understood function of the CX is its role in navigation, in particular, the processing of compass and motion information. Connections between columns form a circular neural architecture, which lends itself to represent angular information as a ‘bump’ of neural activity that shifts from column to column as the angle changes (Stone et al., 2017). A review article by Honkanen *et al* (2019) shows how that ring structure can be used to perform different tasks depending on the ecological niche of the animal, from heading maintenance to saccade control, path integration and migratory route following. The authors also suggest that the evolutionary preservation and ubiquity of the CX in insects is perhaps due to its ability to perform a multitude of navigational roles.

Calcium imaging studies suggest that the CX represents direction information as a *ring attractor* (Seelig and Jayaraman, 2015; Turner-Evans and Jayaraman,

2016).[5] Some columnar neurones in the protocorebral bridge, the fan shaped body and the ellipsoid body, together with the interconnecting tangential neurones form a ring-like connection system. Based on known connections in the CX and using a spiking neurone model Kakaria and de Bivort (2017) reported the emergence of ring attractor dynamics in the protocerebral bridge. This dynamics was observed *in vivo*, using optochemical markers in fruit flies (Kim et al., 2017). Collett and Collett (2018) suggests a mechanism where the mushroom bodies (described below) regularly provide general target orientation while the CX is responsible for short-term maintenance of heading and the actual steering of the animal.

There are again parallels between the neural organisations of certain invertebrate and vertebrate brain regions. In case of the central complex its similarity to the vertebrate basal ganglia was reported by Strausfeld and Hirth (2013).

The central complex will be mentioned later when I will discuss path integration. A recent editorial by Heinze and Pfeiffer (2018) summarised what was known about this neuropil at the time. The paper pointed out that the most understood aspect of the central complex is its role in navigation, in particular, in path integration and obtaining a celestial compass direction. Indeed, there is an

---

5 *Attractor dynamics is far beyond the scope of this thesis. In very broad terms, the state of a complex system at any moment can be represented as a point in a multidimensional vector space, usually called the phase space, where each dimension represents some individual parameter of the system. As the state of the system changes, either by itself or due to external disturbance, that point moves through the phase space. If there are preferential points to which the system tends to return after a disturbance, those points are called attractors. If the attractor points form a continuous curve in the phase space, it is then called a line attractor. If a line attractor is a closed curve then it is called a ring attractor.*
*In neurology the system is a collection of neurones and the dimensions of the phase space are the activity levels of the individual neurones. Thus, a point in the phase space represents a particular activity pattern. As the activity of the individual neurones change, the system moves through a trajectory in the phase space. An attractor is formed if the neurones, after some external influence, tend to settle down to a well-defined activity pattern. A ring attractor is formed, for example, if at any time only one or two neurones are active and the external stimulus moves that activity from one neurone to the next and to the next and so on, in a well-defined order. Such an arrangement can be used to represent angles (or any periodic abstract construct) cheaply, as only one or two neurones need to be active at any time.*

increasing body of research about that topic, for example (Fiore et al., 2017; Held et al., 2016; Le Moël et al., 2019; Pegel et al., 2018, 2019).

Furthermore, it is generally accepted that navigation memories are stored in the mushroom bodies. However, Dewar *et al* (2015) suggested that location memories can be formed in the central complex as well.

The connectome of the *Drosophila* central complex was published by Hulse *et al* (2021) and the authors conclude that the discovered networks are suitable for navigational purposes. For Hymenoptera, the projectome of the bumblebee central complex was presented by Sayre *et al* (2021), including a comparison with the fruit fly CX.

## Mushroom bodies

Mushroom bodies (MB) are paired neuropils in arthropod brains, believed to be associative centres. For the purpose of this introduction, in this section I will give a concise and much simplified summary of their organisation and role. There is an extensive body of literature on the mushroom bodies, of which only a few selected articles will be referenced. A historical overview of mushroom body research as well as a review of what is known about their role in the overall neural system of insects can be found in (Strausfeld et al., 1998), together with an extensive selection of references on the subject spanning over a century.

The mushroom bodies, first described over a century and a half ago (Dujardin, 1850), get their name from their shape. There is one in each brain hemisphere. They are believed to be associative memory centres. The mushroom bodies receive olfactory, and in case of Hymenoptera, also direct visual input. Evolutionary studies indicate that originally the mushroom bodies received only olfactory inputs, the addition of some degree of visual data was a development that happened multiple times for certain orders on the taxonomy tree, such as Odonata, Blattoidea, Coleoptera and Hymenoptera (Strausfeld et al., 2009). Recent results seem to indicate that some information of visual origin is also delivered to the mushroom body of *Drosophila* (Diptera) (Vogt et al., 2014, 2016; Li et al., 2020b).

The mushroom bodies have an input and an output region. In Hymenoptera the input region or *calyx* is further divided into a *lip,* a *collar* and a *basal ring.* The lip receives olfactory input, the collar visual, and the basal ring gets both (Gronenberg, 1999, 2001). The output region, the *peduncle,* bifurcates to the *vertical lobe* and the *medial lobe*. A detailed comparative anatomy of the mushroom bodies of arthropods can be found in (Strausfeld, 2012).

The axons delivering information to the MB synapse to the MB intrinsic neurones called *Kenyon cells* after their discoverer (Kenyon, 1896). The number of Kenyon cells in the mushroom body varies widely between species; a fruit fly has about 2500, while a honeybee approximately 180,000 (Groh and Rössler, 2011). For visually navigating ants the figure is closer to that of the bee

(Ehmer and Gronenberg, 2004). No specific data are available for *Myrmecia* but 100-150 thousand neurones is probably a good estimate, considering both the mushroom body volume in microscopic images (see *Fig. 11*) and the fact that as a visually navigating Hymenopteran, its mushroom body needs to process visual as well as olfactory information. In Hymenoptera there are two basic Kenyon cell types (Strausfeld, 2002). *Type-II* or *clawed* ones receive input from at most a few tens of afferents while the *Type-I* or *spiny* ones have a much larger receptive field. The different roles of Type-I and Type-II Kenyon cells were modelled and analysed by Peng and Chittka (2017). They found that Kenyon cells with small receptive fields are better for discriminating between input patterns and the ones with large receptive fields are more suited for generalisation. While the Kenyon cells receive input from a relatively low number of afferents, the mushroom body output neurones are characterised with contacting a large number of Kenyon cell axons.

The number of neurones delivering visual information from the optic lobes to the mushroom bodies is not known for *Myrmecia*. However, for honeybees the figure is roughly 400 per eye (Ehmer and Gronenberg, 2002). As the visual input from the optic lobe is delivered to both mushroom bodies, they both receive about 800 visual afferents. The honeybee visual system is generally considered superior to that of ants (Ehmer and Gronenberg, 2004), so we can take that figure as an estimated upper limit for *Myrmecia*. Therefore, the Kenyon cells involved in visual processing vastly outnumber their sources of input.

Mushroom bodies are associative networks that can learn to recognise and to individually identify certain patterns or pattern combinations of their input (Aso et al., 2014a, 2014b; Falibene et al., 2015; Ichinose et al., 2021; Strausfeld et al., 1998). The structure is based on first changing the densely coded input into a very sparsely coded sample space, which then is turned back into a densely coded recognition space.

A much simplified neural circuitry of a mushroom body is shown in *Fig. 14*.

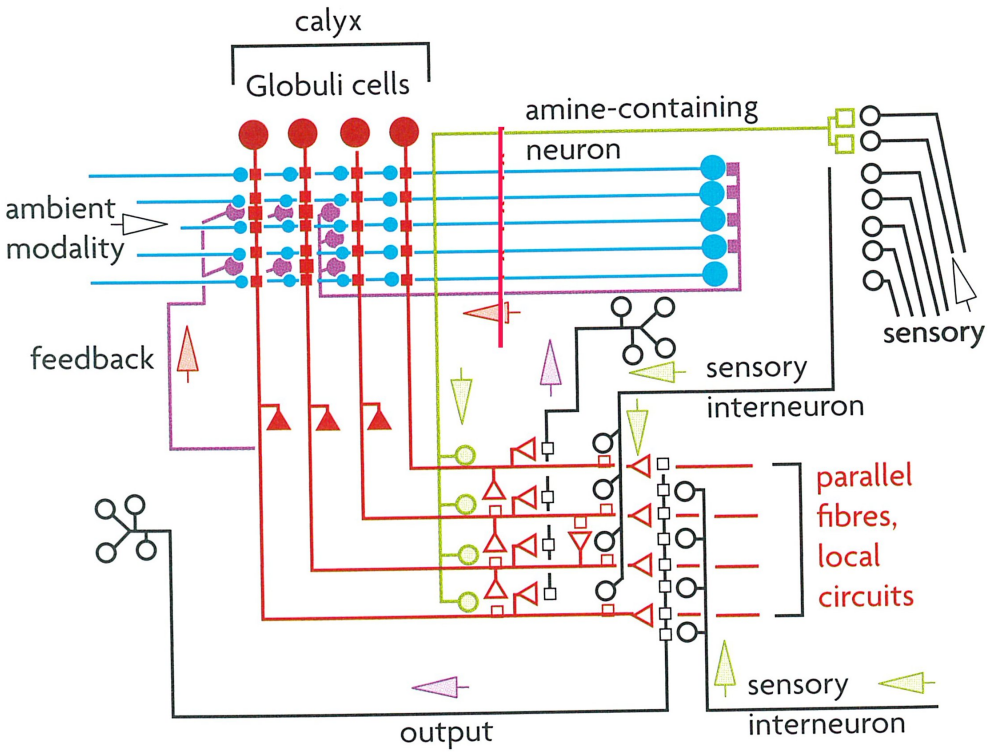


*Fig. 14: Simplified neural circuitry of the mushroom body*

*The input neurones (marked as ambient modality), contact the Kenyon cells (marked as globuli cells) in the calyx. The excitation of the Kenyon cells is affected by the negative feedback neurones, purple in colour. In the peduncle further modulation can happen and there are even connections between Kenyon cells. The output of the mushroom body is generated by neurones that contact a large number of Kenyon cells and can also be affected by external interneurones. Adapted from (Strausfeld, 2012)*

A prominent feature of the mushroom body is a negative feedback loop. Intrinsic neurones synapse on Kenyon cell axons and provide inhibitory input to the Kenyon cell dendritic region. It has been suggested, at least in the context of Drosophila olfactory learning, that the negative feedback maintains the sparsity of the firing pattern of the Kenyon cells, e.g. (Papadopoulou et al., 2011; Lei et al., 2013; Assisi et al., 2020).

It should be noted that this basic blueprint for information processing is not unique to the mushroom bodies. As with the optic lobes and the central

complex, the basic structure can be found in vertebrates as well. Often mentioned is homology with the cerebellum. Albeit using a more sophisticated implementation, it is organised the same way as the mushroom bodies, e.g. (Farris, 2011; Li et al., 2020a). However, similarities with the hippocampus and other vertebrate brain centres were also pointed out (Strausfeld et al., 1998).

While mushroom bodies are considered to be the learning centres, learning is not restricted to the MB peduncles. Evidence indicates that there is synaptic plasticity in the medulla, as well as in the mushroom body input region (Rössler, 2019). Furthermore, the negative feedback neurones in the mushroom body are well known, but there also are feedback loops extrinsic to the MB, where MB output affects external circuitry which, in turn, after combining the MB output with other signals will project back to the MB (Takemura et al., 2017).

Recently artificial neural networks (ANNs) were built based on the basic connection diagram of the mushroom body. Despite their extreme simplicity, these ANNs have shown a remarkable ability of memorising, recalling and associating input patterns. Mushroom body inspired ANNs are described and analysed in, for example, (Mobbs, 1984; Smith et al., 2008; Huerta, 2013; Arena et al., 2013; Peng and Chittka, 2017; Roper et al., 2017). The visual storage capacity of a mushroom body ANN in the particular case of ant navigation is described in (Ardin et al., 2016). In that paper the authors conclude that a 100,000 Kenyon cell MB is capable of storing and recognising about 800 images with only 1% probability of confusion. The model stores the raw image, but Le Moël and Wystrach tested a similar model and found that storing edges instead of raw images significantly increases the storage capacity of the simulated MB (personal communication, 2019). It was also suggested that storing attractive (views towards a goal) as well as repellent (views away from a goal) views increases the navigational performance of a MB inspired ANN (Le Möel and Wystrach, 2020).; see also

The full connectome of the *D. melanogaster* mushroom body was presented in (Li et al., 2020a). It is worth noting that compared to many other insects the fruit fly has a fairly small and simple mushroom body and yet it is still immensely

more complex than even the most sophisticated artificial neural mushroom body model or what was shown in *Fig. 14*.

That the mushroom body plays a role in the visual navigation of ants has long been assumed, if for nothing else, because it receives the optic lobe output. More recently the need for an intact calyx region for visual navigation was experimentally demonstrated for the wood ant *F. rufa* (Buehlmann et al., 2020). The same experiment also established that innate visual behaviour was not affected by the improper operation of the MB, only learned navigation. The need for an intact MB peduncle (or at least its vertical lobe) for view based navigation was demonstrated in *M. midas* (Kamhi et al., 2020) . That is, the animal needs an operational mushroom body to perform visual navigation tasks. In that context what we would like to know is how the mushroom bodies store, compare, individually identify and, would the need arise, forget visual scenes. That is, in general, how visual information is presented to and processed in the mushroom bodies. The role of the mushroom bodies in insect navigation is reviewed in (Webb and Wystrach, 2016).

Unfortunately, most of the mushroom body research centres on olfaction, see for example (Galizia, 2014). This is partly because it is relatively simple to deliver olfactory stimuli to an immobilised animal while optically recording from its brain, and partly because the animal's response is more direct (for example, whether a bee extends its proboscis or not). Visual training, scene presentation and response monitoring of an immobilised animal is much harder. Furthermore, most of the work is done using *Drosophila melanogaster*, due to the extensive genetic toolset available for that animal. Alas, although the mushroom body of the fly receives visual input (Li et al., 2020b), it is from the lobula (where motion and shape detection happens) and not from the medulla (where visual features are extracted). The bee MB, which receives direct visual input from the medulla, is also extensively studied. However, the toolset available for bee physiology studies is much more limited than what is available for the fruit fly. Furthermore, despite the bee being a visual navigator, bee research also focuses on olfaction and not vision.

A hidden danger of the olfaction centric body of research is that it might lead to preconceptions. For example, the mapping of olfactory input to Kenyon cells is essentially random (Caron, 2013; Caron et al., 2013). The artificial models of the mushroom body also assume a random connection scheme for visual data, even though the literature does not provide proof either for or against that assumption. However, the visual processing pathway is retinotopic throughout, thus the data arriving at the mushroom bodies are likely to be spatially organised, containing information which would be discarded by randomisation. While it is possible that the spatial relationships are not important at the mushroom body level, my personal (and I'd like to emphasise that) view is that biological systems are not known to squander information at their disposal. Future research will need to establish whether that is indeed the case.

In any case, as it was discussed previously, artificial neural networks based on very simple MB models are able to memorise and recognise visual scenes. It is generally accepted that mushroom bodies of visually navigating insects can deliver scene familiarity information to the animal, even though the exact neurological details are not known.

## Summary

Ants have compound eyes which, in case of *Myrmecia*, resolve the visual scene to a few thousand pixels. They have a fully panoramic field of view. We know that they have green, blue and UV sensitive photoreceptors although the distribution of different receptors throughout the visual field is not known. A specialised region of the eye, the dorsal rim area (DRA), can detect the UV polarisation pattern of the sky. The animals also have three small camera eyes, the ocelli, that, in honeybee, at last, are also known to contain polarisation sensitive receptors (Ogawa et al., 2017).

The visual information is heavily processed by the optic lobes. Polarisation and motion information is delivered to the central complex, believed to be responsible for behavioural decisions. The pre-processed image itself, most likely represented as a retinotopic but sub-sampled set of features, is delivered to the mushroom bodies, which are associative memory centres. These neuropils can memorise patterns and compare the momentary input against

those memories. The mushroom body outputs are also believed to be sent to the central complex; a direct pathway between the MB and the CX was recently discovered in Drosophila (Li et al., 2020a), although has not yet been shown for other species.

That is the machinery that the animal can use to navigate visually.

Here I would like to stretch the patience of the Reader and finalise this section with a philosophical tangent with some back of the envelope type calculations. As it was said earlier, the brain of *M. pyriformis* is composed of about 500k neurones and both of her compound eyes are built from roughly 3k ommatidia with 9 photoreceptors each. The optic lobes have retinotopic organisation, thus each neuropil contains 3k of repeated units. The exact number of neurones in those units has not been determined, but we can make an educated guess. In the honeybee, another Hymenopteran, each lamina cartridge contains at least 4 morphologically different neurones while a medulla column contains 10 (Ribi and Scheel, 1981). The number is not reported for the lobula in this paper, but it is stated that its volume is roughly the same as that of the lamina. Furthermore, the paper demonstrates that there are also at least 10 different types of tangential neurones in the medulla and at least 6 in the lobula.

Let us conservatively assume that the tangential neurones are a quarter in number compared to the cartridge/columnar neurones. The optic lobe output neurones will simply be ignored. Based on the same volume, we estimate the lobula to contain the same number of neurones as the lamina. Now let us also assume that the optic lobe of the pedestrian *M. pyriformis* has half the complexity (i.e. number of neurones) of that of the flying *A. mellifera*. Thus, the ant optic lobe should contain about 2(la) + 5(me) + 2(lo) = 9 columnar neurones plus a further 2.25 neurones per column as tangential cells, plus the 9 photoreceptors for each ommatidium. That multiplied by 2 eyes and 3k ommatidia per eye gives the figure of 121k neurones. Assuming roughly 130k neurones per mushroom body[6] and also estimating that only half of the mushroom body is devoted to visual information, we then arrive at up to ~250k

6 *No data are available for* Myrmecia*, the figure is for* Camponotus *(Ehmer and Gronenberg, 2004).*

neurones involved in vision. Even if we generously ignore the ocelli and the central complex we can conclude that *M. pyriformis* spend half of their total neural budget on vision.

Such an investment in a single sensation can only be justified with significant selective advantages. Hunting prey while avoiding to become one are the obvious ones. But being able to see also equips the animal with the ability to assess its surroundings close and far simultaneously. As long as there is some light and the air is more or less optically clear, which is normally the case, vision works and it works from the immediate vicinity to practically infinity. It delivers the shapes and locations and colour features of many objects simultaneously. It can measure both self motion and that of external agents. It is not affected by wind and temperature changes. None of the other sensations is so versatile. Those features make vision a unique tool to aide a social forager with one of her most important tasks, namely navigation, to which I turn to next.

## Navigation

Navigation is solving the problem of getting from the current location to a target location. Humans instinctively associate navigation with maps. We know where we are, know where the target is, mark these points on a map and we get the direction and the distance, possibly with waypoints, to satisfy some external constraints. If we want to circumnavigate the globe, or want to find a certain place in an unfamiliar city, we now heavily rely on maps and GPS. An ant's life is much simpler. She needs to go out foraging, sometimes to a familiar place such as a tree, and back to the nest. That task can be solved without a map and even without having a position fix (Graham and Philippides, 2017). As someone summarised it very eloquently[7]:

> *They don't know where they are but they do know where to go.*

Ants have three major tools to aid them in their navigational tasks. It is a common misconception that all ants lay and follow pheromone trails as their sole means of navigation. First, even the species that do use pheromone trails also have other methods in their navigational toolbox. More importantly, many ant families, including *Myrmecia* do not use pheromone trails. As solitary foragers, they need to be able to navigate independently. Myrmecia have antennae with the usual set of olfactory receptors (Ramirez-Esquivel et al., 2014) but there is no evidence that they use them for navigation. For that reason, even though pheromone trail dynamics is a very interesting research topic, olfactory navigation is not examined any further in this thesis.

The other two navigational tools at the ants' disposal are path integration and visual navigation. Before detailing those, first a short summary on what is known about navigation in *Myrmecia* ants is presented.

Different sympatric *Myrmecia* species in Australia occupy distinct temporal foraging niches and their visual system properties reflect their activity schedules

7 *It was paraphrased in (Murray et al., 2020) as "They may know where to go, but they do not know where they are".*

ranging from exclusively diurnal to fully nocturnal (Greiner et al., 2007; Narendra et al., 2011, 2016).

Both the night active *M. pyriformis* and the day active *M. croslandi* are individually foraging along idiosyncratic routes (Reid et al., 2011, 2013; Narendra et al., 2013b; Jayatilaka et al., 2014). *M. pyriformis* begin foraging after sunset throughout the year and their foraging activity is strictly dependent on ambient light intensity (Narendra et al., 2010). *M. pyriformis* are guided by celestial compass information and the landmark panorama during foraging (Reid, 2010; Reid et al., 2011), their foraging efficiency suffers at low light levels (Narendra et al., 2013c) and is affected by changes in landmark panorama (Narendra and Ramirez-Esquivel, 2017).

Freas and colleagues have recently documented very similar reliance on the visual landmark panorama for navigation in another crepuscular/nocturnal *Myrmecia* ant, *M. midas* (Freas et al., 2018; Freas and Cheng, 2019).

The day active *M. croslandi* rely on the landmark panorama within the catchment area of familiar views, ignoring information from path integration (Narendra et al., 2013b; Zeil et al., 2014b). *M. croslandi* foragers are able to home while moving backwards by repeatedly turning around to get a bearing (Schwarz et al., 2017a). They do respond to landmark displacements close to the nest (Jayatilaka, 2014), which they presumably learn to recognize during their learning walks (Jayatilaka et al., 2018). For the general foraging ecology of *M. croslandi*, including activity schedules and individual foraging trips see (Jayatilaka et al., 2014).

## Path integration

Path integration is a method by which an insect can track its position relative to her starting point (e.g. the nest) even in the absence of visual cues. For animals living in visually rich environments it is only used as a last resort mechanism. But even for those animals evolution has retained it, indicating that it is a fundamental tool, thus deserving a brief explanation.

If we know our position at some point in time (called a *fix*), then if we keep track of our movements, we can calculate our momentary position. That is called *dead reckoning*. It is a simple application of Newton's laws of motion:

$$\boldsymbol{s}(t) = \boldsymbol{S_0} + \int \boldsymbol{v}(t)\,dt$$

where the position vector is calculated by integrating the velocity vector over time. Due to that integral, in insect research dead reckoning is called *path integration,* often shortened to PI. The use of dead reckoning by animals was first suggested by Darwin (1873). Path integration for both flying and walking insects is reviewed in (Collett, 2019).

When an animal leaves the nest for a foraging trip, path integration provides a vector that is her momentary position relative to the nest, or, in the opposite direction, the nest's position relative to her. That is why it is called the *home vector*. Furthermore, if the animal discovers an abundant food source, she can memorise her home vector at the source's location. That vector is what will lead her, on a future foraging trip, from the nest straight to the food. In fact, that is the vector that bees describe to other workers in the hive during the waggle dance[8].

8 *The waggle dance is out of the scope of this thesis, however a few literary references are listed here. In short, a bee, after returning from a foraging trip, 'dances' in the hive. The dance conveys the vector to and the quality of the food source that she has found. The dance does it with surprising precision, especially if one takes the motion of the sun (the directional reference) and the presence of wind into account. The precision of the dance language is examined in, for example, (Al Toufailia et al., 2013; Dornhaus and Chittka, 2004; Schürch and Ratnieks, 2015).The waggle dance theory was first published in a paper by Karl von Frisch (1946), and became popularised when he wrote the book "The Dance Language and Orientation of Bees" which was translated to English (von Frisch, 1967). The theory was controversial for a while and there was a long scientific battle between von Frisch (an ethologist) and A. Wenner (a mathematical biologist), starting with the paper (Wenner et al., 1967). The human aspect of that fight is analysed in (Munz, 2005).*

There is a specific terminology associated with path integration. An animal at the nest is called a *zero-vector* animal. As she goes out on a foraging trip, she *builds up* the home vector. When she initiates her return to the nest she is said to have a *full vector*. She then *runs down* the vector and on the way home she has a *partial vector*. When she arrives back to the nest she again becomes a zero-vector animal.

Conceptually, path integration is very simple. However, in real life the measuring equipment is never perfect and there is always noise on the signals, both of which leads to measurement errors. Accumulation of these errors can lead to large deviation from the true goal direction and distance. Mathematical analysis of path integration (Vickerstaff and Cheung, 2010; Cheung et al., 2012; Cheung, 2014) concluded that the path integrator's robustness against noise and sensory errors greatly depends on the choice of coordinate system and the representation of vectors.

Dead reckoning, in its classical form, depends on the knowledge of the momentary velocity vector. Before the advent of GPS, ships and aeroplanes used inertial navigation, a special case of dead reckoning. Inertial navigation systems calculate the velocity vector by integrating the acceleration vector, which, in turn, is derived from measuring inertial forces. Mathematically:

$$\boldsymbol{s}(t) = \boldsymbol{s}_0 + \int [\boldsymbol{v}_0 + \frac{1}{m} \int \boldsymbol{F}(t)\, dt]\, dt$$

Due to the double integration, a single momentary measurement error of the inertial forces triggers a linearly increasing absolute error of the calculated position, while a systemic measurement error results in a quadratic increase of the position error. Extreme efforts of precision engineering went into inertial navigation systems to minimise measurement errors.

Using a platform that could be both linearly and rotationally accelerated, an experiment indicated that some mammals (in this case, gerbils) do use inertial navigation for homing on short excursions from their nest (Mittelstaedt and Mittelstaedt, 1980).

For insects, evolution selected a different, less error sensitive method. Instead of the double integration by time insects learned to measure the velocity vector directly and even to simplify the task to a summation of known displacement vectors.

***Speed measurement***

A flying insect, such as a bee, can determine her momentary speed from the *optic flow*. The optic flow is the term used to describe the change of the visual panorama due to the self-motion of an agent (Lee and Kalmus, 1980). Optic flow analysis can be used to determine position relative to a plane, measuring bank and pitch, calculate self motion, work out distances of objects and several other navigationally relevant factors (Koenderink, 1986). Optic flow is especially important for flying insects. Seminal experiments have shown that honeybees do indeed use it to measure distance flown (Esch and Burns, 1995; Srinivasan et al., 1997). During flight the ventral visual field can be used to determine ground speed, eliminating the problem of wind compensation. On the other hand, optic flow depends on parallax, which in turn depends on altitude; reviewed in (Webb, 2007). The optic flow is also used to stabilise flight (Portelli et al., 2010; Srinivasan et al., 1996). When combined with airspeed sensor data, it is sufficient to control all parameters of the flight (Rutkowski et al., 2011; Taylor et al., 2013).

For walking insects with a stable gait the problem is (seemingly) simpler. The integration of the velocity vector over time can be substituted with a summation of unit length displacement vectors. The natural displacement unit is a single step the animal takes. She can then store distance information as the number of steps to take in a straight line. An elegant experiment was devised to prove that desert ants use step counting instead of absolute distance measurement. Full vector animals were captured and little stilts were attached to their legs, which increased their stride. They overshot their nest by the expected amount. Other animals underwent partial amputation and with the shortened stride they undershot the nest (Wittlinger et al., 2006).

Step counting is a very simple and elegant solution for the vector calculation, but it only works on flat surfaces. If on the outbound journey the animal had to

cross a dune but on the return journey she is on even ground, she would overshoot the nest because during climbing and descending on the dune she needed to spend steps for her vertical displacement. Since ants find their nest, they apparently know how to sum only the horizontal component of each step (Hess et al., 2009).

Although ants could in principle use optic flow for distance estimation, they seem to rely on the step counter instead (Ronacher et al., 2000). On the other hand, if an ant cannot use the step counter, for example while being carried by a nestmate, she updates her path integration using optic flow (Pfeffer and Wittlinger, 2016).

***Direction measurement***

Both speed measurement and step counting works on an absolute scale. Zero means no displacement at all while any positive measure indicates a well-defined displacement, in some arbitrary units.

For the direction component, however, the value of 0 has no absolute meaning. It just denotes the direction that is identical to an arbitrary *reference direction*, relative to which all other directions are measured. Therefore, for an animal to express a direction, she needs an external reference (Cheung et al., 2007) that is stable and independent of the animal's own motion.

*Celestial compass*

To get a directional reference the animal needs something in her environment that is at a constant (or easily calculable) position, independent of the animal's movements, prominent, unambiguous and very reliably present. When an animal obtains such a directional reference from visual cues in the sky, it is said to use a *celestial compass*.

The most prominent object on the daytime sky is the sun, a very reliable reference. It has no observable parallax. While it is constantly moving over the sky, it moves slowly and its motion is completely predictable. The sun is very easy to find because, without exception, it is the brightest object in any daytime natural scene. Ants were shown to use the sun for directional reference in a famous experiment more than a century ago: Animals were returning from the

foraging trip, with the sun at their side. Then a screen was used to hide the real sun from the animal and at the same time a mirror presented the sun at the other side, effectively rotating the sky by 180° around the animal. The ants did an about-face, maintaining their heading relative to the sun, real or projected (Santschi, 1911).

Of course, the sun is not always directly visible. It can be behind a cloud or obscured by foliage. It still can provide a directional reference for the animal.

The upper atmosphere causes sunlight to be linearly polarised to some degree. The more acute the angle with which the sunlight hits the atmosphere, the more polarised it becomes. Moreover, the polarisation pattern is entirely determined by the position of the sun, shown in *Fig. 15*.

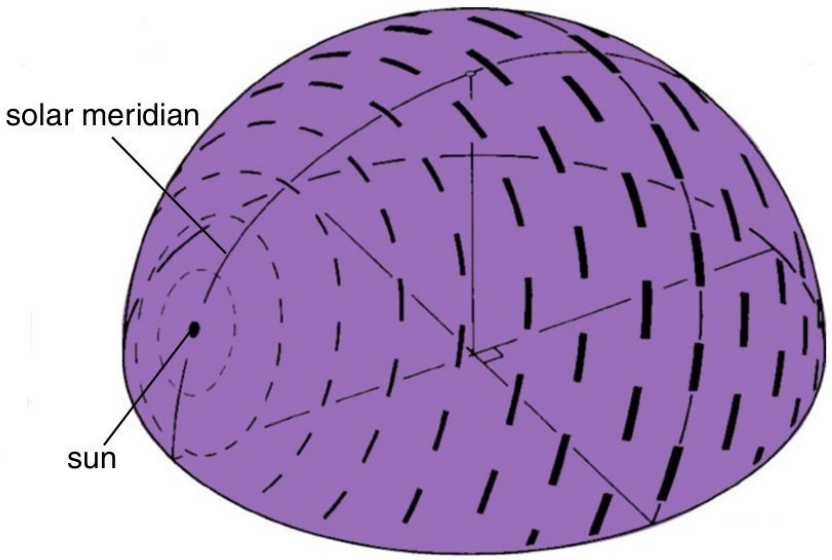


*Fig. 15: Polarisation pattern of the sky*

*The direction of lines indicate the E-vector direction and their thickness indicates the degree of polarisation (polarisation contrast). Adapted from (el Jundi et al., 2014)*

Insects know how to use that pattern to determine the occluded sun's position, more precisely, how to obtain a directional reference from the polarisation pattern. They use the polarisation sensitive dorsal rim area of their eyes (page 35) to establish their own orientation relative to the sky's polarisation pattern.

Although all components of the sunlight become polarised, Hymenoptera chose to use UV. In an open environment under clear skies the UV has no advantage over the visible spectrum, in fact, blue light would offer more polarisation contrast. However, there are clouds and animals often live under vegetation canopies. These scatter light and destroy its linear polarisation. With those effects taken into account, then the remaining polarisation contrast is highest in UV (Labhart, 1999; Pomozi et al., 2001; Barta and Horváth, 2004). Nevertheless, using UV polarisation is not universal. Desert locust, for example, are sensitive to the polarisation of the blue component (Kinoshita et al., 2007) while some beetles use green light polarisation (Labhart and Meyer, 1999).

For many species, the polarisation pattern is also sensed by the ocelli (page 35) (Fent and Wehner, 1985; Narendra and Ribi, 2017; Ogawa et al., 2017; Ribi and Zeil, 2017; Ribi et al., 2011; Zeil et al., 2014a). It is not clear how it is used, but experimental results suggests that path integration uses the information provided by the dorsal rim area while short route segment directions are independently measured by the ocelli (Schwarz et al., 2011c, 2011b).

Labhart (2016) pointed out that there is no need for the eye to extract a polarisation angle as such. Rather, a head direction (relative to the sun's position) should be presented to higher level processing centres.

Indeed, that information can be obtained without even measuring the polarisation direction. As it is shown in *Fig. 15*, how much the light is polarised (the *polarisation contrast*) depends on the angle with which the sunlight hits the atmosphere. At the exact position of the sun, there is no linear polarisation. As one moves on the sky further and further away from the sun, the polarisation contrast monotonically increases until 90° from the sun's position, then it monotonically decreases again. Thus, the polarisation contrast gradient directly points towards or away from the sun, providing the required compass reference. Therefore, even if the animal cannot measure the *direction* of the polarisation, if she can compare the polarisation *contrast* between different spatial directions, she can obtain the direction reference she needs. A biologically plausible method to extract direction information from the polarisation contrast is described in (Gkanias et al., 2019).

The polarisation pattern is a reliable compass even when only part of the sky is visible or in cloudy conditions or in twilight. The cloudy condition polarisation patterns are examined in, for example, (Labhart, 1999; Pomozi et al., 2001) while the twilight patters are analysed in (Cronin et al., 2006; Barta et al., 2014). During the night, however, the Earth completely occludes the Sun. In a moonless, cloudy night the sky simply does not offer any useful directional information. But on a clear night it still presents sufficient navigational cues. The moon, like the sun in daytime, is the brightest object on the sky. It is also the largest, much larger than any other celestial light source. Just like the sun, the moon creates a polarisation pattern and dung beetles (Dacke et al., 2011) as well as *M. pyriformis* (Reid et al., 2011) are known to use that as a compass, even in crescent moon conditions. Unlike the sun, however, the moon is not always present. During new moon it is not visible at all. On those days certain dung beetles rely on the Milky Way as a star compass (Dacke et al., 2013a; Foster et al., 2017). The Australian bogong moth (*Agrotis infusa*) is also known to use the starry sky as a reference (Dreyer et al., 2019). The remarkable navigational prowess of that moth is analysed in the review articles (Warrant and Dacke, 2016) and (Warrant et al., 2016). Further literature references about the use of the night sky as a celestial compass can be found in the review article (Foster et al., 2018).

It is interesting that the celestial compass is so reliable that ants that normally navigate by a mix of pheromone trails and path integration react to polarisation pattern changes even when on the pheromone trail (Freas et al., 2019).

*Magnetic compass*

Some nocturnal insects, such as the bogong moth, rely on the Earth's magnetic field (Dreyer et al., 2018) to assess the reliability of the visual information (landmark panorama).

The magnetic field is also used by naïve desert ants as a directional reference during their initial learning walks (discussed later) outside the nest (Fleischmann et al., 2018a). Experimental proof of insects using a magnetic compass is still scarce. Current knowledge about the topic is reviewed in (Fleischmann et al., 2020).

*Esoterica: light pollution*

Insects are very adaptable. Humans built cities and lit them up. Light pollution levels are so high that there are large areas in Europe where during night time the Milky Way is no longer visible; in fact, only the brightest stars can be seen by the naked eye. One expects that light pollution disturbs nocturnal insect navigation. Indeed, the review paper (Owens and Lewis, 2018) reports the negative effects of artificial lighting on insects in general and fireflies in particular. However, some insects adapted to the new situation. The sky compass is not available so they use in its stead what took it from them: the light pollution itself. Cities are large, immobile and bright. Once established, they have a specific and rarely changing light pattern that can reliably be used for navigation. Dung beetles doing just that is described in the conference paper (Foster and Dacke, 2019) and subsequent journal article (Foster et al., 2021).

***Summing vectors with neurones***

In summary, the celestial compass or some other external source provides a directional reference to the animal. That direction and odometry from optic flow and/or the step counter together supply a stream of displacement vectors to the animal, which she needs to sum to maintain her home vector.

That seems to be a very simple problem. If one decomposes the vectors into two orthogonal components, one parallel to the body axis and one pointing sideways and sums those components independently, then the two sums will be the components of the home vector. The orthogonal decomposition was hypothesised by Mittelstaedt (1962) in the context of flight direction control. But experiments performed with desert ants presented some strange results. On their outward journey the ants were constrained to run in a narrow and straight plastic channel. On the unconstrained return path they showed a systematic error of their path integrator and missed the nest. When they were not constrained on the outward journey and it was not straight, they found their nest without problem (Wehner and Wehner, 1986). That observation led to a long and sometimes heated argument about the neural representation of direction. The main points of that argument can be traced through (Müller and Wehner,

1988; O'Keefe, 1991; Touretzky et al., 1993; Hartmann and Wehner, 1995; Wittmann and Schwegler, 1995a).

An alternative hypothesis is that insect brains represent direction vectors as *phasors* stored and accumulated in *sinusoidal arrays* (Touretzky et al., 1993; Wittmann and Schwegler, 1995b; Goldschmidt et al., 2017). A phasor is the amplitude and phase of a unit frequency sine wave. It is trivial to show that mapping a 2-dimensional Euclidean vector to a phasor is an invertible linear transformation. Consequently, the phasor representing the sum of vectors is the same as the sum of the phasors representing the individual vectors.

Phasors can be easily represented by equidistant samples of one period of the sinusoid. Such a representation, at least in neurobiology, is called a sinusoidal array. Anatomical studies suggested that the neural machinery to store such a representation was readily present in the central complex. However, the phasor model assumes that receptors have a raised sine response to angular stimulus. Intracellular recordings on bees seem to indicate that that is not the case; the excitation curves fall off significantly faster than a sinusoid (Stone et al., 2017). On the other hand, Hulse et al (2021) reports that in *Drosophila* there are neurones (Δ7) that transform the incoming signals to a sinusoidal response before passing the information to the protocerebral bridge and the fan shaped body, supporting the idea of using phasors for heading representation. Thus, at the time of writing, the phasor model is not yet ruled in or out.

An alternative to the phasor model is the *ring attractor* model (of which sinusoidal arrays are a special subset), which uses the same subsampled period representation but not necessarily storing a sinusoid in it. The attractor concept was described on page 45. Anatomical studies and intracellular recordings supporting the model have been published (Kim et al., 2017).

A review article on the neurobiological basis of path integration is (Heinze et al., 2018).

***The reach of path integration***

The odometer, compass and summing circuits are of course not infinitely precise. But they are still capable of amazing navigational performance. The

desert ant *Cataglyphis fortis* lives in a salt pan environment that has very weak terrestrial visual landmark cues. Thus, she relies mostly on her path integrator to find her way home. During a series of experiments testing their response to displacements with and without visual landmarks close to the nest it was reported that the maximum distance an animal reached from the nest was over 300m and the longest foraging trip had a path length of over 1.2km (Huber and Knaden, 2015). In the experiments the animals were allowed to go various distances from the nest. They were then displaced laterally by 60m and were given food to motivate them to return home. Animals from a nest with no landmarks around it ran down their home vector and their angular error was found to be less than ±6° even at the longest distance. Interestingly, animals that ventured far (over 50m) from the nest always kept a straight outward run. It is assumed that they do this because summing vectors that greatly differ in direction is more error sensitive than summing vectors pointing roughly the same way. Animals from the nest with a nearby landmark switched to landmark based navigation and overruled the path integrator, but only when the landmark was clearly visible for the animal.

In a different experiment it was probed how willing ants were to switch from path integration to landmark based navigation. The results suggest that animals living in a feature rich environment are more ready to ignore the path integrator and rely on landmarks than animals living in a visually barren environment, who preferred the path integrator even when landmark cues were present (Schultheiss et al., 2016b), corroborating a similar finding earlier (Narendra, 2007a, 2007b)

Finally, before moving on to visual navigation, an aspect of navigation that I have so far neglected deserves a quick mention. Even the best path integrator or landmark navigation method can only take the animal to the vicinity of the nest. The nest entrance is often deliberately inconspicuous and the animal needs to search for it. Hymenoptera use multiple search strategies, with different statistical properties – and it seems that the method that a species employs depends on their environmental niche and their lifestyle (Schultheiss et

al., 2015). That article, though not a review paper as such, provides a very good summary of all the literature related to search behaviour.

It is believed that chemical signals (e.g. $CO_2$) emanating from the nest also help the animal in finding the entrance, shown, for example for desert ants (Buehlmann et al., 2012a) but no experimental evidence is present for *Myrmecia*.

## Navigation using the visual panorama

Perhaps the most versatile method is navigation by sight of the landscape panorama. It gets the agent to the target, but at the same time it helps with other tasks, like obstacle avoidance. It also offers different strategies depending on the circumstances.

The simplest case is when the target is in line of sight, called *beacon aiming,* e.g. (Graham et al., 2003). All the agent has to do is to walk forward while keeping the target at the centre of its frontal visual field. It does not matter where the agent is and there is no need for any other external directional reference, such as a celestial compass.

But even if the target is not directly visible, the agent can navigate by knowing the target's location relative to objects that are visible. The ant's task is to get back to the nest from anywhere within the foraging range. From her point of view that is literally a life or death problem. The nest is usually not directly visible from afar (in fact, often it is deliberately well hidden) yet she needs to find it. She has the path integrator, but, in itself it does not allow the ant to pinpoint the nest.

If a human would want to describe the location of the nest, something like "about halfway between the tall gumtree and that large bush left of it" would be said. But how can an ant go to "halfway between the tree and the bush"?

The next few subsections examine the various solutions suggested to that problem.

### *Cognitive maps*

One theory is that an agent can build a *cognitive map* of the environment. It knows the location of the target points as well as important *landmarks*, easily identifiable visual objects, on the map. The map can either be a geometric map where the animal knows the actual spatial geometry of the features of the landscape or a topological map, where the precise geometry is not known, but routes connecting landmark points are. In the geometric case, by examining the relative positions of landmarks, the agent is aware of its own location on the map at all times. This idea was originally applied to vertebrates (O'Keefe, 1991)

where it is still an active research topic. In 1986 Gould suggested that bees also build a cognitive map of visual landmarks, because his experiments appeared to show that bees were able to make shortcuts between familiar places (Gould, 1986). Several papers followed; for a review see (Wehner and Menzel, 1990). It was a popular idea for a while, but the results were not all supportive. Based on the available experimental data, a systematic review by (Wehner et al., 1996) came to the conclusion that all the observed behaviour could be explained by simpler means. There was no experimental evidence for the existence of a cognitive map in insects. The authors also made a very important observation, which deserves repeating:

> *"The brain has evolved not to reconstruct a full representation of the three-dimensional world, but to find particular solutions to particular problems within that world."*

The issue did not disappear completely. There are still proponents of the cognitive map (Menzel et al., 2000, 2005; Cheeseman et al., 2014a, 2014b) and those who rebuke them (Cruse and Wehner, 2011; Cheung et al., 2014).

In a more recent review article Webb (2019) argues that if an animal uses visual scene familiarity to set its path integrator into a known state and then uses the PI, then *ipso facto* it is a map-like representation of its surroundings. However, the paper concludes that evidence for such an explicit connection between panoramic views and the PI system is very scarce and further experimentation is necessary to establish whether it exists at all.

***Landmark based navigation***

Landmark based navigation assumes that insects identify prominent features of their visual environment and they use the visual parallax of those selected objects to position themselves and also to determine the location of the target.

The use of landmarks by bees for navigation was systematically examined by Cartwright and Collett (1982, 1983) and led to the development of the *snapshot hypothesis* (Cartwright and Collett, 1987). The idea is that the bee identifies salient objects in the environment and stores a snapshot of their locations on

her retina. When she is at a different place, the location of those objects relative to each other will have changed on her retina and she can use that to derive instructions on how to move to minimise the mismatch between the current view and view that she memorised at the goal.

Desert ants live in an environment that is visually parsimonious[9]. Thus, they rely on their path integrator to find their way back to the nest. However, experiments suggest that if they are offered salient visual cues by means of artificial objects, of which the position relative to the nest is known (learnt) by the animal, the ants not only will use those cues but even overrule their path integrator with them, see for example (Bolek and Wolf, 2015; Narendra, 2007a, 2007b; Wehner et al., 1996). How readily does an animal switch from vector based navigation to landmark based one seems to depend on the probability of any landmarks occurring in her natural environment at all (Buehlmann et al., 2011). Cheung *et al* (2012) present a mathematical analysis of why that is so. The topic is further examined in (Buehlmann et al., 2018) and in (Freas and Cheng, 2018). Computational models for multi-source navigation are suggested in (Wehner et al., 2014) and (Hoinville and Wehner, 2018).

Landmark based navigation of ants who live in visually rich environments was reported, for example, in (Durier et al., 2003; Fukushi, 2001, 2004; Harris et al., 2007) for wood ants, (Narendra et al., 2013b) for *Myrmecia* and in (Aksoy, 2014) for mining ants. Mangan and Webb (2012) have documented the very interesting phenomenon of ants developing multiple simultaneous idiosyncratic routes between the same two points, assumedly using landmarks and snapshot memories e.g. (Collett et al., 2003).

A very detailed review article on landmark based navigation and how it augments or overrules path integration is (Collett et al., 2007). As an exception to the rule, dung beetles ignore landmarks and rely exclusively on celestial compass information when moving along straight lines (Dacke et al., 2013b).

9 *Or, in case of the Australian desert ants, visually cluttered at the ground level, but lacking large features that have measurable, but not too much parallax during a foraging route.*

***Panorama based navigation***

Ants have nearly spherical vision. That gives the animal the possibility of using the entire panorama as a reference, rather than having to segment the scene into identified individual objects.

Assuming a static environment, the panoramic image that an animal sees is determined by two factors. First, it obviously depends on which way the animal is facing. Second, it depends on where the animal is. Zeil *et al* (2003) demonstrated that the difference (which will be formally defined below) between panoramic images in natural environments changes smoothly both by translation and by rotation. They also observed that the change by rotation is more rapid. Thus, if the animal takes a snapshot of the panorama, then after a displacement she can extract compass information from the difference between the memorised image and the current view and she can also use gradient descent to find her way to the spot where the snapshot was taken.

An important characteristics of the method is that there is no need for the agent to identify individual objects in the scene. It is the entire panorama, without any segmentation, that guides the animal.

The central concept is the *image difference function* or IDF for short. It is defined as the pixel-by-pixel (or ommatidium by ommatidium) difference between two images, averaged by some means. The literature usually uses the root-mean-square[10] (RMS) of the pixel differences. Mathematically, for two rectangular images $I_1$ and $I_2$ that are both $W$ pixels wide and $H$ pixels tall, the IDF is defined as:

$$IDF = \sqrt{\frac{\sum_{x=1}^{W}\sum_{y=1}^{H}[I_1(x,y)-I_2(x,y)]^2}{W \cdot H}}$$

By its definition, the IDF is 0 if and only if the two images are identical. It is 1 (assuming that intensity values are between 0 and 1) if and only if both images

10 *The general properties of the IDF do not change if one uses a different method. A simple averaged absolute value works just as well. Also, there is no experimental evidence that insects use pixel by pixel RMS calculation; in fact we do not know how they represent images, let alone how they calculate their differences.*

are composed from pixels that are all black or white (all pixels are 0 or 1) and one image is the negative of the other one. For two images composed of completely random pixel values the expected value of the IDF is the expected value of the RMS of the difference of two independent random variables *X* and *Y* with uniform distribution on the [0,1] interval:

$$E[IDF]=\sqrt{\int_0^1\int_0^1 (x-y)^2\,dx\,dy}=\frac{1}{\sqrt{6}}\approx 0.41$$

or, if the 8-bit pixel value range of 0-255 is used, approximately 100. Any value significantly different from that indicates that the images are in some way correlated, although it cannot be readily established how. For example, two black and white photographs of natural scenes will likely show some level of correlation even if they were taken at independent locations, because natural scenes tend to be lighter at the upper part (sky) and darker on the bottom (ground features). That will result in a lower IDF value than what is expected for two completely random sets of pixel values. A natural image will also have a lower than expected value of the IDF if it is compared against a featureless artificial image, with the bottom half uniform dark grey and the upper half light grey. The actual value of the IDF in general is not very informative, because it is highly context dependent. It is the *change* of the value by moving or rotating at a given scene what is important, that is, how the IDF behaves if we take a snapshot and calculate the IDF between that snapshot and the current view as we move or rotate.

Let us define two functions.

The first one is the rotational IDF, which is calculated by placing the panoramic camera at a reference point, taking a snapshot; then rotating the camera along its vertical axis and for each camera angle calculate the IDF between the reference snapshot and the momentary image seen by the camera. Thus, the rotIDF(φ) at a certain reference location is a function of the rotation angle φ of the camera. The camera's location is not changed.

The other function is the translational IDF. Here we place the camera at the reference location and take a snapshot. We then move the camera to different

locations, but always keep its orientation relative to North to be the same as it was at the reference location. At each location we calculate the IDF between the momentary view and the reference snapshot. Thus, tranIDF(**p**) is a function of the position vector **p** which is relative to the reference location. The camera's orientation remains constant throughout.

We can also combine the two methods. We take a snapshot at the reference location facing a reference direction. Then we move the camera to location **p** and then rotate it at that location and calculate the IDF between the reference snapshot and the momentary view. This combined method gives us the IDF(**p**,φ) function, which gives us a scalar value depending on our position **p** and orientation φ relative to the reference snapshot.

The general behaviour of tranIDF and rotIDF are demonstrated in *Fig. 16*.

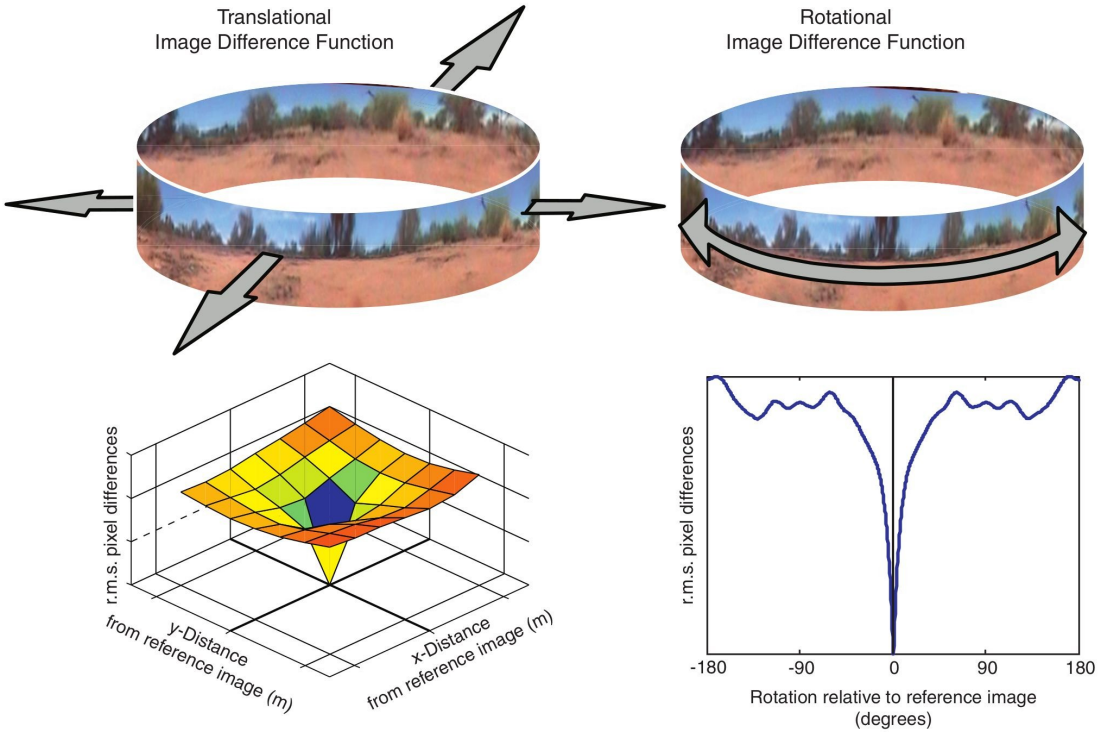


*Fig. 16: Translational and rotational image difference function*

*Adapted from (Zeil, 2012).*

If we examine the combined IDF behaviour, we will find that it is more sensitive to rotation than to translation. That has a simple explanation.

Let us assume that we have a panoramic snapshot and we move forward a small amount. In the anterior view we got a bit closer to the objects in front of us, so they will look a bit bigger. Conversely, in the posterior view objects will be further, therefore somewhat smaller. In the lateral view the objects moved towards the back, without any apparent size change.

Now, if these objects are very close, these effects will be large. However, far objects will change very little as both the looming and parallax motion effects depend not on the displacement itself but on the ratio of the displacement and the distance of the object. Thus, the distant skyline will not change at all, nearby trees will change somewhat and ground features right in front of us will change the most. In addition, the further we move from the reference image's location, the larger the IDF between the reference view and the current view will be, as more and more looming and parallax effects become detectable.

That is not the case for rotation. If we rotate on the spot by $\vartheta$ degrees, all objects will move in the panoramic image by $\vartheta$ degrees (that is, a certain amount horizontally), regardless of whether they are in the immediate vicinity or they are infinitely distant. As we rotate, the IDF will become large quite fast (every object moves in the scene) and stays there until we almost turned a full circle and the objects all get close to their original position in the image again.

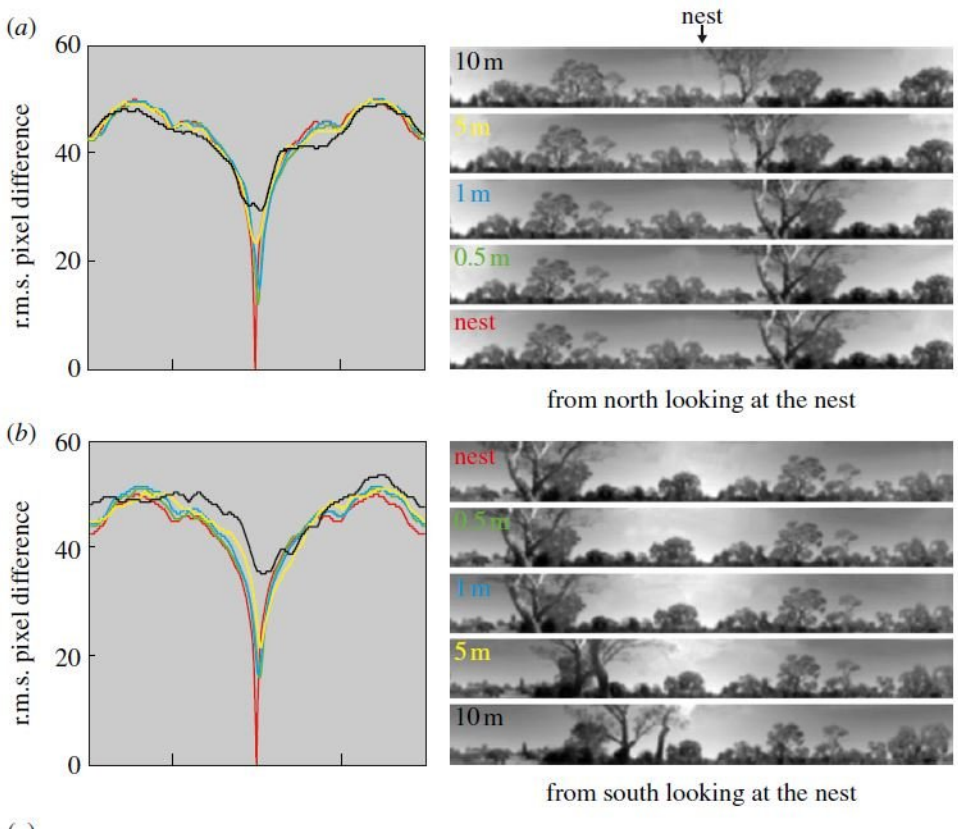


*Fig. 17: rotIDF between a reference image and the momentary view as a function of distance The reference snapshot was taken at the nest (red). Adapted from (Narendra et al., 2013b)*

Since it would be hard to draw an image that shows a function that is dependent on 3 independent variables (the two components of the position vector and the orientation angle), usually it is shown as a collection of rotIDF(φ) functions parametrised with the **p** location. *Fig. 17* shows an example. It demonstrates that at the reference location the rotIDF is zero if the camera is aligned with the reference snapshot's direction. As we start to rotate the camera, the rotIDF quickly rises and remains at a high level until the camera starts to approach the reference direction again. But more importantly the figure also demonstrates that if the camera is moved away, then while the rotIDF will never become zero, its global minimum will be when the camera orientation matches that of the reference image, even at 10m from the reference location. Therefore, if an animal wants to align herself with the orientation she had when she memorised the reference snapshot, all she has to do is to turn on the spot until she finds the global minimum of the IDF.

The figure also demonstrates that once the camera is aligned (i.e. the global minimum is found) the IDF decreases as the camera gets closer to the reference location. Thus, our hypothetical animal can do a gradient descent, moving in the direction where the rotIDF minimum decreases; and when she reaches the value of zero, she knows that she's at the reference point.

There is a rich and growing literature on IDF based navigation; a few selected articles are listed below. The rotational and translational change of the IDF was measured and quantified in natural environments in (Zeil et al., 2003; Stürzl and Zeil, 2007; Murray and Zeil, 2017). The information content of panoramic images in an artificial arena is examined in (Stürzl et al., 2008). The analysis of the actual habitat of *M. croslandi*, for which the homing ability has been previously established experimentally (Narendra et al., 2013b; Zeil et al., 2014b), is presented in (Stürzl et al., 2015). The same paper also analyses another site, from the point of view of ground nesting wasps. Note that this means that it is possible to map the navigational knowledge available for an animal that potentially exploits the IDF for navigation. The change of the IDF as a function of altitude (in the context of an ant climbing a tree) is examined in (Freas et al., 2018). Analysis of the IDF for the foraging paths of *M. midas* is presented in (Freas and Cheng, 2019). Comparison of IDF based and skyline based orientation is discussed in (Philippides et al., 2011).

Stürzl and Zeil (2007) showed that filtering and contrast equalisation makes the IDF more robust and also usable in less feature-rich environments; it is interesting to note that, at least in flies, the lamina performs contrast equalisation (Laughlin, 1981). Computer simulations subsequently demonstrated that the IDF can be used for route following (Baddeley et al., 2012) that its properties are maintained even if the image representation is based on coarse local contrast filters (Baddeley et al., 2011). The IDF can therefore be determined with local image features (which the medulla is believed to extract) rather than with the raw image. It is also important to realise that for view-based navigation a low resolution panoramic visual system has advantages over one that has high resolution but a narrow field of view (Wystrach et al., 2016a).

A very important feature of IDF based algorithms is that the animal does not need to know where she is during the entire process. There is no need for a cognitive map or the understanding of the relative positions of objects. All she needs is a memorised image, the momentary view and neural machinery to calculate a scalar representing the difference between them (and, of course, the machinery to execute the algorithm and turn the result into locomotor response). The algorithms will then guide her back to the location where the memorised snapshot was taken. Furthermore, she does not need to perform geometric calculations about object parallax and loom.

We do know that animals do not store panoramic images as raw pixel values, because on the one hand their mushroom bodies do not have the capacity and, on the other, the visual information they receive has already been heavily processed. Image representation is likely to involve global (Laughlin, 1981) and local (Stürzl and Zeil, 2007) contrast equalisation, filtering and compression. It is likely that the image is represented as retinotopic feature sets rather than subsampled pixel values (see discussion on the Medulla on page 40). We do not know what exactly happens to the image but as long as the transformations preserve the IDF behaviour regarding to translation and rotation, the algorithm is still usable. Since the output of the optic lobes is retinotopic, that is likely the case. Since we don't know how the image is transformed, most ANN models use the raw image.

Baddeley et al (2011) demonstrated that by storing snapshots along a route an animal can re-trace that route based on scene familiarity (i.e. IDF minimisation) without actually knowing where on the route she is. A similar but more sophisticated ANN, based on a spiking neuronal model of the MB, was analysed in (Ardin et al., 2016). The paper investigates the storage capacity of the MB and concludes that ~100 thousand Kenyon cells can store hundreds of images with very little error rate and a very low activation rate. The 100K neurones per MB is realistic for visually navigating Hymenoptera.

Both papers assume that the stored images are treated independently, their sequence along a route is not utilised. However, experimental data (Schwarz et al., 2019) seem to indicate that animals do keep track of the temporal order of

the images along a route and get confused if these are presented to them in an incorrect sequence.

Many studies have addressed snapshot memories and their use in navigation of wood ants (Durier et al., 2003; Graham et al., 2004, 2007; Harris et al., 2007) but most experiments used artificial, high-contrast landmarks and tested the animals' behaviour against a landmark-based (i.e. segmented panorama) snapshot model. While the model worked, the papers did not provide conclusive evidence against full panoramic image comparison; in fact one of them explicitly points out the need for the surrounding panorama for the animal to correctly identify the artificial landmark (Graham et al., 2004).

Both behavioural experiments (Graham and Cheng, 2009a, 2009b) and scene analysis (Philippides et al., 2011) indicate that ants can derive heading direction from the rotIDF even when only the rough skyline is available. Changes in the environment can disrupt ant navigation, and the disruption can be adequately explained by the change of the IDF (Narendra and Ramirez-Esquivel, 2017) and the availability of skyline (Reid et al., 2011). The IDF can supply directional information even when the animal is high above the ground on the foraging tree (Freas et al., 2018).

***Learning walks***

Visually navigating ants perform so-called *learning walks*. Before going out to forage the animal walks around the nest several times in increasing loops in several directions. At regular intervals she alternates between looking towards the nest or away from it.

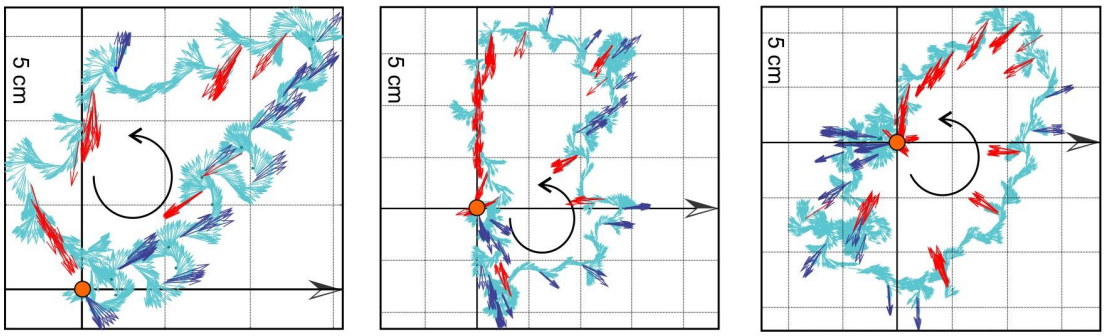


*Fig. 18: Learning walks of M. croslandi*

*The learning walk paths of three different animals. North is to the right. The orange circle is the nest. The black circular arrow indicates the walking direction. The light blue arrows indicate the gaze direction, taken at 40ms intervals. The colour of the arrow is changed to red when the animal looks towards the nest and to dark blue while she is looking away from it. Adapted from (Jayatilaka et al., 2018)*

The consequence of that choreography is that the ants experience nest-directed (attractive?) views from different compass directions and different distances from the nest, alternating systematically with (repellent?) views away from the nest. This choreography is very similar in different species of ants, as reviewed in (Zeil and Fleischmann, 2019).

It is believed that ants memorize views tagged with the nest direction, an information that must be supplied by path integration, during these learning walks, which they subsequently use to pinpoint the nest when returning from foraging excursions, see e.g. (Graham et al., 2010; Fleischmann et al., 2016; Jayatilaka et al., 2018; Zeil and Fleischmann, 2019). For instance, nest-directed views from a few metres away from the nest can provide nest-directed heading direction information over a large area around the nest, depending on the structure of the environment (Narendra et al., 2013b; Stürzl et al., 2015), as demonstrated in *Fig. 19* .

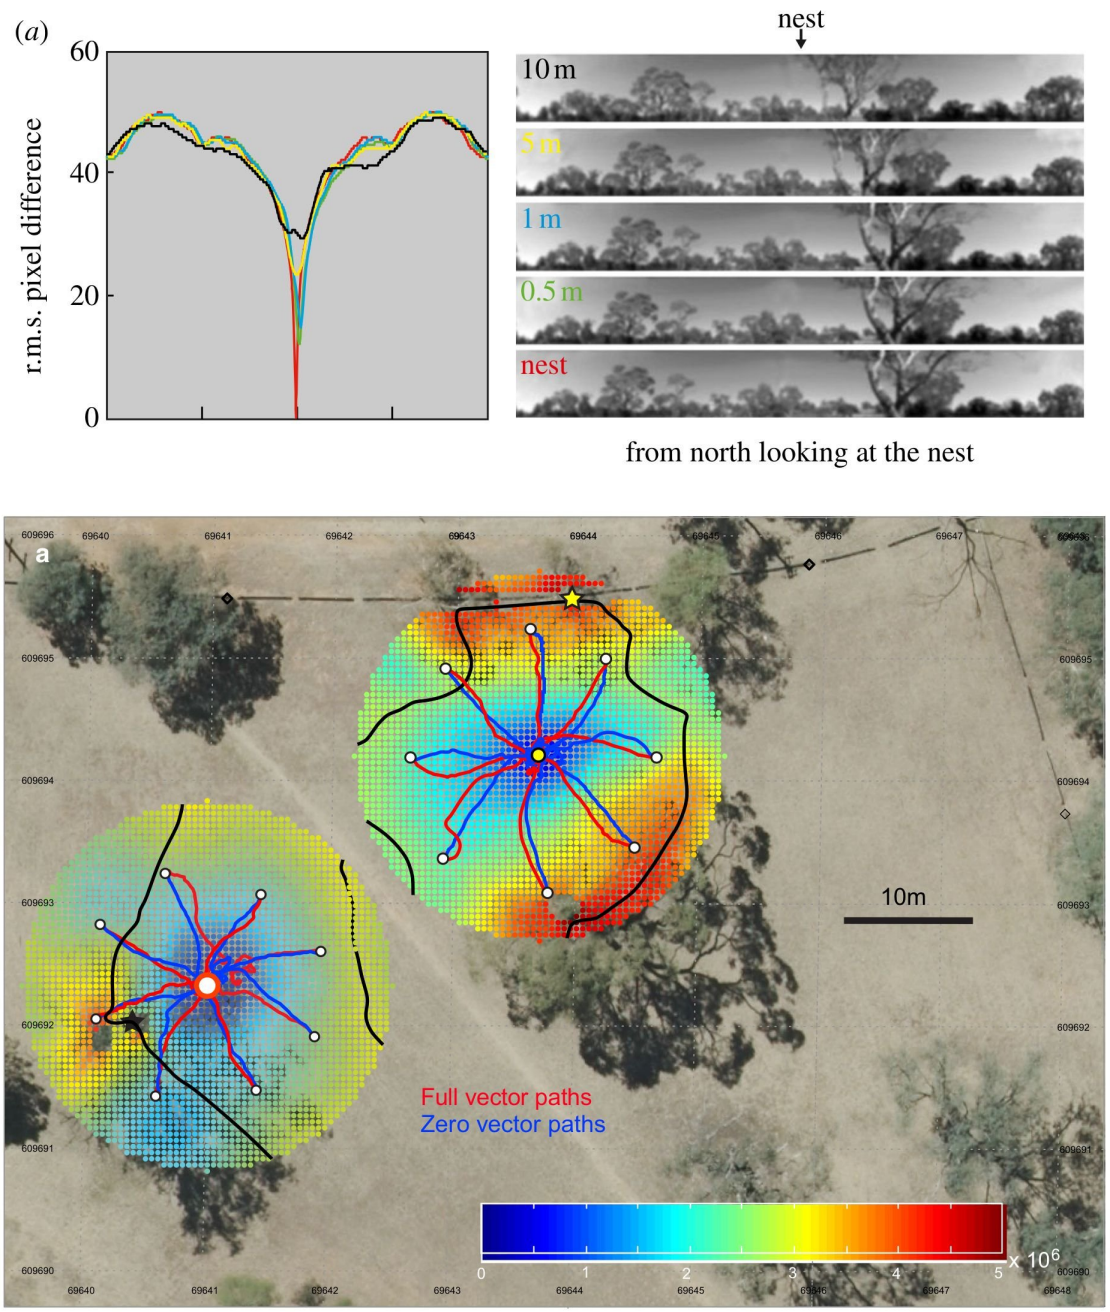


*Fig. 19: rotIDF and transIDF at a nest site*

*Top panel: the rotIDF between a reference image taken at the nest and an image taken at a location 0, 0.5, 1, 5 and 10 metres North of the nest. Note that the global minimum is always at 0 rotation, that is, when the animal looks towards the nest.*

*Bottom panel: transIDF at the same area (two nests are shown). Coloured dots are the transIDF, blue is very low and red is very high, the actual numerical value is irrelevant. Red and blue lines are actual homing paths. The black lines are the limit of rotIDF providing useful direction information.*

*Adapted from (Narendra et al., 2013b; Stürzl et al., 2015).*

Modelling inspired by the choreography of learning walks, by the behaviour of tethered ants at the nest location, and by some known neural processing mechanisms indicates that the use of attractive and repellent scene memories may make both route navigation and homing more robust (Jayatilaka et al., 2018; Le Möel and Wystrach, 2020; Murray et al., 2020).

However, it remains unclear how exactly ants are able to pinpoint the often visually very inconspicuous nest during their final approach, and whether in addition to guidance provided by memorized views ants utilise other indicators of the nest location, such as olfactory cues, e.g. (Buehlmann et al., 2012a).

## Where does that leave us?

In this chapter I described the visual system of ants, the parts of their brain that are currently known to be involved in their navigation and listed the methods the animals are believed to use to find their way in their environment.

As a quick summary, we do know that they use the Sun, or other celestial bodies, as a compass reference. We know that they count steps and that they perform vector integration with surprising precision. In visually parsimonious habitat that is the only tool they can rely on. We also know that if there are landmarks (salient features) in their environment, they use those for navigation. Furthermore, we know that in feature rich habitats they utilise the entire panorama to determine their heading. It seems that animals can use all three forms of vision-based navigation, and depending on their circumstances they can select the most reliable method.

While we have fairly good anatomically mapped neural models for path integration, we have no knowledge of how landmarks or panoramas are stored, recalled and analysed by the animal. We know that the optic lobe has very much fewer output neurones than the number of pixels in the compound eyes. We assume that those output axons carry extracted image features rather than a compressed representation of pixels. We do not know how many and what those extracted features are and how they are used in navigation. We do know that the mushroom bodies are associative memory centres, but we don’t know what they associate with what and how they do it. Our current mushroom body

based navigation models are fairly simple; they use pixels and edge detection, not sets of unknown features extracted from the image.

However, as behavioural experiments, neuroanatomical and neurophysiological studies advance and new results pour in, the models get more and more sophisticated. They should predict what the animals actually do and at the same time they have to fit within the neurological constraints.

For example, the relative simple computer models of panoramic snapshot based navigation by Baddeley *et al* (2011, 2012), did not determine whether their models were neurologically plausible nor have they used natural images. Webb and Wystrach in their review paper (2016) examined the role of the mushroom bodies and the central complex (using simplified neural models) in multimodal navigation. Stone *et al* (2017) suggested a model for path integration that they mapped to neurones in the bee brain. A refined model based on polarisation contrast instead of E-vector direction, complete with compensation for the sun's movement and the tilt of the animal was described in (Gkanias et al., 2019), readily mapped to known anatomical structures. A sophisticated model was introduced in (Hoinville and Wehner, 2018) which deals with the dynamic prioritisation of navigation information sources depending on the momentary assessment of their reliability. Note that at that level the entire path integrator or the neural circuitry calculating panoramic image IDF are just processing units among others, under the supervision of higher level neural circuitry. The concept is extended in (Sun et al., 2020) where the animal's motivation (homing, route following, searching etc.) is also taken into account.

Thus we can conclude that the models we currently have are getting larger, encompassing several behavioural, environmental and neural properties simultaneously; and more complex, involving more and more neural circuitry. The question then becomes, how to test these models. Are they robust? Are they neurally realistic? Can we experimentally show that living animals do use the model's core algorithm(s)?

To be able to do that, we need to pose the right questions to the animals. Showing them black and white stripes is no longer enough, because the neural

architecture the details of which we are trying to probe has been evolutionarily optimised to recognise natural scenes and not artificial high contrast stripes[11]. To further complicate matters, we try to understand visual navigation while we don't really understand visual processing in animals. As Gronenberg (2008) pointed out: "*Almost nothing is known about central visual information processing in ants – no neurophysiological data exist on visual interneurons in ants*" and that situation has not significantly improved since the paper was published.

To get useful answers we thus need to be able to present views of the natural environments to animals and be able to manipulate these visual environments in various ways, then monitor the behavioural or electrophysiological response to that change.

That is problematic. Doing electrophysiology on small insects in their natural environment is simply not practical. Furthermore, changing the environment in arbitrary ways is not that simple. We cannot move trees around, make hills disappear from the skyline, rotate the celestial light distribution relative to the landmark panorama, or, for a challenge, change the fractal dimension of a scene (Zahedi and Zeil, 2018).

It is really hard to manipulate the visual appearance of distant objects, or to present an arbitrary UV polarisation pattern, or to distort the world in some complex, non-linear way. Yet it seems that that is exactly what we need to do to test our models.

The solution is reconstructed visual reality, at least a partial one. A device that can recreate the visual component of the natural habitat of the animal in the laboratory and which allows arbitrary changes to the reconstructed world.

11 *Point in case: in an experiment that measured head stabilisation, wasps were oscillated around their roll axis in their natural environment, in a black and white striped cylinder and in darkness. One would think that the cylinder supplies the best clues, with very high contrast and solid, straight edges. Yet, the animals gave the best result outdoors, in their natural environment (Viollet and Zeil, 2013).*

In addition, the devide allows us to move the animal within the world without actually interfering with her physically. That ability removes all the disturbance introduced by physical manipulation of the animal.

During my candidature I designed and built a virtual reality arena optimised specifically for ant vision. The focus of this thesis is that device, which I named the Antarium.

The rest of this thesis is organised in the following topics:

- Chapter 2 presents the design of the Antarium.
- Chapter 3 describes outdoor experiments that provided the reference data for testing the device.
- Chapter 4 details the experiments that were designed to assess the feasibility of the Antarium and identify possible shortcomings in its implementation.
- Chapter 5 presents the results of Antarium experiments where the visual reality was manipulated in ways that would be very hard to achieve without a reconstructed reality projector.
- Chapter 6 demonstrates how the next generation Antarium will overcome all the shortcomings of the current Antarium device that were discovered during the experiments.
- Chapter 7 summarises the Antarium work described in the previous chapters and sketches some future research directions opened up by the unique capabilities of the Antarium device.

# The design of the Antarium

As the Introduction chapter concluded, for more sophisticated experiments we need an arena that can present the visual component of the natural environment of the ant. The animal can be put on a trackball and placed inside the arena. In a closed-loop configuration she can freely explore her natural visual environment all the while the scenery can be manipulated in arbitrarily ways for experimental purposes. Such an arena has three basic elements: the projector device, the three-dimensional (3D) rendering software and the 3D model of the environment. Our group decided to build such an arena, specifically tuned to the visual apparatus of ants. I have a background in electronics engineering, allowing me to design the projector while Trevor Murray's experience in computer game programming made him the ideal candidate to code the 3D rendering engine.

The device, which I named Antarium, has been built. Its technical details as well as the initial findings regarding to its use have been published:

*Kócsi, Z., Murray, T., Dahmen, H., Narenra, A., and Zeil, J. (2020)* The Antarium: A Reconstructed Visual Reality Device for Ant Navigation Research. *Frontiers in Behavioural Neuroscience, 14:599374.*

The rest of this chapter contains the final manuscript verbatim. The following table details the authors' contributions to the design of the system:

| | TM | ZK | HD | AN | JZ |
|---|---|---|---|---|---|
| **Conceptualisation** | | • | | | • |
| **Trackball system** | | | • | | |
| **Trackball testing outdoors** | • | • | • | • | • |
| **Projector design** | | • | | | |
| **Electronics design** | | • | | | |
| **Manufacturing supervision** | | • | | | |
| **3D rendering software** | • | | | | |
| **3D model** | • | | | | |
| **System testing & debugging** | • | • | | | |
| **Behavioural experiments** | • | • | | | |
| **Funding** | | | | • | • |
| **Project supervision** | | | | | • |
| **Project administration** | | | | | • |
| **Funding acquisition** | | | | | • |

*Table 1: Article authors' contribution to the Antarium*
*The name abbreviations are: TM, Trevor Murray; ZK, Zoltán Kócsi; HD, Hansjürgen Dahmen; AN, Ajay Narendra; JZ, Jochen Zeil.*

# The Antarium:
# A Reconstructed Visual Reality Device for Ant Navigation Research

Zoltán Kócsi[1], Trevor Murray[1], Hansjürgen Dahmen[2], Ajay Narendra[3]

Jochen Zeil[1]

[1]Research School of Biology, Australian National University, Canberra ACT2601, Australia; [2]Cognitive Neuroscience, University of Tübingen, Tübingen, Germany; [3]Department of Biological Sciences, Macquarie University, Sydney NSW 2109, Australia.

**Correspondence**: Zoltán Kócsi zoltan@bendor.com.au

**Keywords**: Visual navigation, virtual reality, reconstructed visual reality, ants, LED arena

## Abstract

We constructed a large projection device (the Antarium) with 20 000 UV-Blue-Green LEDs that allows us to present tethered ants with views of their natural foraging environment. The ants walk on an air-cushioned trackball, their movements are registered and can be fed back to the visual panorama. Views are generated in a 3D model of the ants' environment, so that they experience the changing visual world in the same way as they do when foraging naturally.

The Antarium is a biscribed pentakis dodecahedron with 55 facets of identical isosceles triangles. The length of the base of the triangles is 368 mm resulting in a device that is roughly 1m in diameter. Each triangle contains 361 blue/green LEDs and 9 UV LEDs. The 55 triangles of the Antarium have 19,855 Green & Blue pixels and 495 UV pixels, covering 360° azimuth and elevation from -50° below horizon to +90° above the horizon. The angular resolution is 1.5° for Green and Blue LEDs and 6.7° for UV LEDs, offering 65,536 intensity levels at a flicker frequency of more than 9,000 Hz and a framerate of 190 fps.

In addition, the direction and degree of polarisation of the UV LEDs can be adjusted through polarisers mounted on the axles of rotary actuators.

We build 3D models of the natural foraging environment of ants using purely camera-based methods. We reconstruct panoramic scenes at any point within these models, by projecting panoramic images onto 6 virtual cameras which capture a cube-map of images to be projected by the LEDs of the Antarium.

The Antarium is a unique instrument to investigate visual navigation in ants. In open loop, it allows us to provide ants with familiar and unfamiliar views, with completely featureless visual scenes, or with scenes that are altered in spatial or spectral composition. In closed loop, we can study the behaviour of ants that are virtually displaced within their natural foraging environment. In future, the Antarium can also be used to investigate the dynamics of navigational guidance and the neurophysiological basis of ant navigation in natural visual environments.

## Abbrevations

FPGA: field programmable gate array

PWM: pulse-width modulation

PCB: printed circuit board

SPI: serial peripheral interconnect

UDP: user datagram protocol

## Introduction

Ample experimental evidence now makes us confident that central-place foraging insects, such as ants, bees and wasps navigate predominantly visually, relying on both scene memories and celestial compass information e.g. (Zeil, 2012; Collett et al., 2013; Wystrach et al., 2014; Graham and Philippides, 2017; Wehner, 2020). Visual navigation is supported by path integration (Heinze et al., 2018) which runs in the background, providing a failsafe, and in some cases and situations, also by olfactory, tactile and magnetic cues (Buehlmann et al., 2012b; Knaden and Graham, 2016; Fleischmann et al., 2018a). Evidence from

behavioural studies and increasingly detailed knowledge of neural circuits relevant for navigation e.g. (Stone et al., 2017; Buehlmann et al., 2020; Kamhi et al., 2020) are beginning to feed into neurally constrained and experimentally informed models of navigation e.g. (Ardin et al., 2016; Webb and Wystrach, 2016; Hoinville and Wehner, 2018; Stone et al., 2017; Differt and Stürzl, 2020; Sun et al., 2020) and into robotic implementations e.g. (Lambrinos et al., 2000; Möller, 2000; Stone et al., 2016, 2017; Dupeyroux et al., 2018; Sabo et al., 2017).

It is likely that the predictions of these models will become increasingly hard to test in behavioural experiments. The main reason being that controlled manipulations of complex visual cues, such as the full landmark panorama or conflict experiments between different compass systems are difficult to perform in natural navigation environments. Equally, investigations of the real-life computational properties of navigation-relevant neural circuits are currently hampered by limitations in the way visual information can be presented in electrophysiology rigs (see e.g. Table 1). There are currently no projection devices that can convey the full information content of the spatial, spectral and polarization signal patterns that characterize natural navigation environments. And lastly, the navigational competence of insects is based on active learning processes e.g. (Collett and Zeil, 2018; Jayatilaka et al., 2018; Zeil and Fleischmann, 2019) and relies on the active comparison between remembered and currently experienced input patterns e.g. (Zeil, 2012; Murray et al., 2020; Le Möel and Wystrach, 2020). It is thus likely that the neural machinery underlying navigation is heavily state-, context- and activity-dependent, requiring closed loop control of the visual scene by the insect and control by the experimenter over the experience (what has been learnt?), the motivation (what is the navigational goal?) and the state of the animal (whether it holds information from path integration or not).

With this in mind, we designed the Antarium, a panoramic projection device that would allow us to present ants walking on a trackball with views of their known foraging environment and to give the insects full control over the view transformations by feeding their intended movements back onto the panorama.

Beside the engineering challenges of the device itself, there are two pre-conditions for this to work: a need to know the movements of the ants in their natural foraging environment and a way of reconstructing the views they will have encountered under natural conditions. To satisfy the first condition, we rely on several years of tracking ant movements with differential GPS, both during their normal foraging activity and after systematic displacement experiments e.g. (Narendra et al., 2013b; Zeil et al., 2014b). We secondly used LIDAR and camera-based methods to build 3D models of the ants' foraging environment e.g. (Stürzl et al., 2015, 2016; Murray and Zeil, 2017), which we now can use to render panoramic views at any location within the foraging range of the ants and project them in the Antarium.

The Antarium is not the first 'Virtual Reality' device in insect research but it is the first one which has been designed with the specific aim of enabling the presentation of natural, in contrast to synthetic, visual navigation environments e.g. (Van De Poll et al., 2015). We summarize the features of some devices described in the literature in *Table 2* and briefly describe their properties below, see also (Dombeck and Reiser, 2012; Fry et al., 2004, 2008; Stowers et al., 2017; Schultheiss et al., 2017).

| | Strauss et al. 1997 | Gray et al. 2002 | Lindemann et al. 2003 | Reiser & Dickinson 2008 | Takalo et al. 2012 | Paulk et al. 2014 | Koenig et al. 2016 | Kaushik et al. 2020 | Antarium |
|---|---|---|---|---|---|---|---|---|---|
| **Colour** | Green | RGB | Green | Green | White | Green | RGB | LGB | G, B, UV |
| **Polarised** | No | No | No | No | No | No | no | No | Yes (UV) |
| **Technology** | LED | Projector | LED | LED | Projector | LED | Projector+ l. guides | LCD | LED |
| **Azimuth [°]** | ±180 | ±125 | ±125 | *Depends* | ±135 | ±180 | ±180 | ±180 | ±180 |
| **Elevation [°]** | -0,+45 | ±125 | -90,+70 | *Depends* | -64,+57 | -35,+45 | -45,+45 | -58,+72 | -50,+90 |
| **Number of pixels** | 5,760 | 307,200 | 7,168 | 64N | 480,000 | 4,096 | 5,760 | 11 million | 19,855 GB 495 UV |
| **Intensity levels** | 2 | 256 | 8 | 8 | 256 | 2 | 256 | 256 | 65,536 |
| **Flicker [Hz]** | 1,000 | 60 | ∞ | 372 | 360 | >300 | ? | ? | >9,000 |
| **Frame rate [Hz]** | 1,000 | 60 | 370 | 372 | 360 | ? | 300 | 165 | 190 |
| **Angular resolution [°]** | 20 | <10 | 2.3 | *Depends* | a.so | 3.5 | 2-3 | 0.14 | 1.5 GB 6.7 UV |
| **Light level** | 60 $cd \cdot m^{-2}$ | 14 lux | 420 $cd \cdot m^{-2}$ | *Depends* | 4 $W \cdot m^{-2}$ | 168 lux | ? | ? | N/A |
| **Closed loop** | Yes | Yes | No | *Depends* | Yes | Yes | *Depends* | Yes | Yes |

*Table 2: Parametric comparison of existing insect research VR systems and the Antarium*

Disckinson and Lighton (1995) built a cylindrical arena with green LEDs which was limited to display a dark vertical bar that could be rotated around the animal. The device could not display an arbitrary scene. Similarly, Strauss *et al* (1997) designed a projector for walking *Drosophila* experiments. It is a cylindrical device, with monochrome (green) LEDs. A full-colour computer projector with a hemispheric back-projected screen was built by Gray *et al* (2002) and combined with a wind tunnel for moth research. The FliMax device (Lindemann et al., 2003) is an LED projector designed for fly research. It delivers a monochromatic (green) image for the tethered insect in its frontal visual field and was used to present reconstructed, outdoor view-sequences in electrophysiological experiments (Boeddeker et al., 2005). Reiser and Dickinson (2008) designed a modular projection device consisting of small identical square panels of monochromatic (green) LEDs. These modules can be used to tile a surface that has curvature around at most one axis, for example a cylinder[12]. The projection system designed by Takalo *et al* (2012) is based on a modified video projector with elaborate optics. Paulk *et al* (2014) used four LED panels to build a square well around the animal on the trackball. The panels are approximately 20cm squares, with a 32 by 32 matrix of RGB LEDs on each. Only the green channel was utilised and only vertical bars were shown to the animal. Commercial projectors beamed onto a hemisphere were used by Peckmezian & Taylor (2015) who presented artificial 3D environments to trackball mounted jumping spiders. Koenig *et al* (2016) projected simple shapes onto a rectangular array of light-guides, the other ends of which lined the walls of a cylindrical arena. More recently Kaushik *et al* (2020) built an arena where the tethered insect is placed in the geometric centre of a triangular prism formed by 3 high-speed commercial computer monitors turned on their side, delivering full-colour video of a 3D modelled landscape.

The aim of the Antarium project was to design a projection system for experiments on ant navigation which must be capable of presenting panoramic

12 *A bi-colour (green and blue) version of that device is now commercially available and a trichromatic (green, blue and UV) device is being designed (M. Reiser, 2019, personal communication).*

views of the natural foraging habitat of ants in a way that addresses their spectral and polarization sensitivities while also allowing the ants to interact with the scene and the experimenter to modify it in arbitrary ways.

None of the existing projection systems could deliver on all these points. The following constraints were considered at the outset:

- Since ants have panoramic vision e.g. (Zollikofer et al., 1995; Schwarz et al., 2011a), the arena must cover 360° azimuth and the whole celestial hemisphere. Similarly, the arena must be able to project ground features down to -45° elevation.
- At the time the Antarium was designed, the spectral sensitivities of *Myrmecia* ants were not known, but scattered reports made it likely that ants in general possess UV, blue and green receptors, see references in (Ogawa et al., 2015).
- The Antarium must be able to deliver light of sufficient intensities at these wavelengths. On a sunny day, the brightness in a natural scene can vary by 5 log units. The Antarium should be able to deliver a similar intensity range.
- Like most insects, ants possess a dorsal eye region with UV and polarization sensitive receptors that feed into the skylight polarization compass system. The Antarium, therefore, would need to provide adjustable polarisation covering the celestial hemisphere.
- We work with Australian bull ants. One of the largest bull ants (*Myrmecia pyriformis*) has around 3,500 ommatidia per eye (Narendra et al., 2011). Therefore, in order to avoid aliasing, the number of pixels must be at least 20 000.
- The critical flicker fusion frequency (CFFF) has been determined for two *Myrmecia* species, for the nocturnal *M. midas* at 84.6 ± 3.2Hz and the diurnal-crepuscular *M. tarsata* at 154.0 ± 8.5Hz, cited in (Ogawa et al., 2019). For the Antarium, we opted for a minimum flicker rate of 300Hz. The minimum frame rate for ants to observe continuous motion is not

known, but it cannot be higher than the critical fusion frequency. Therefore, a frame rate close to 200fps should be sufficient.

- We decided to use the trackball system designed by Dahmen (Dahmen et al., 2017) that records the rotations of a hollowed-out, air-supported Styrofoam sphere using optical mouse sensors. Beside a very high sampling rate, the advantages of this system is that it can be used in two ways: with the tethered animal free to rotate around the yaw axis and the trackball recording the animal's translational movements only and with the tethered animal fixed, so that the trackball movements reflect both the yaw rotations and the translational movements of the animal.
- Finally, we had to operate within tight budgetary constraints.

The Antarium offers unique and crucial opportunities to investigate visual navigation in ants and to test models of visual navigation. It allows us to confront ants in both open and closed loop with familiar and unfamiliar views of their natural environment, but also with completely featureless visual scenes, or with scenes in which dominant objects have been removed or displaced or that are altered in spatial or spectral composition. Most importantly, the Antarium can also be used in future to investigate the neurophysiological basis of ant navigation in natural visual environments.

## The Antarium design

***Geometry***

Although an ideal projector would be spherical, there are a number of practical constraints that make this untenable. For example, if LEDs were drilled and glued to the inside surface of a sphere, the optics would be ideal, see e.g. (Koenig et al., 2016). However, hand-soldering thousands of LEDs to their drivers is error prone and extremely labour-intensive, and thus prohibitively expensive. A faster and cheaper alternative is to have machine assembled printed circuit boards (PCB). PCBs can be any shape, but must be flat, which constrains the projector to be a polyhedral approximation of a sphere.

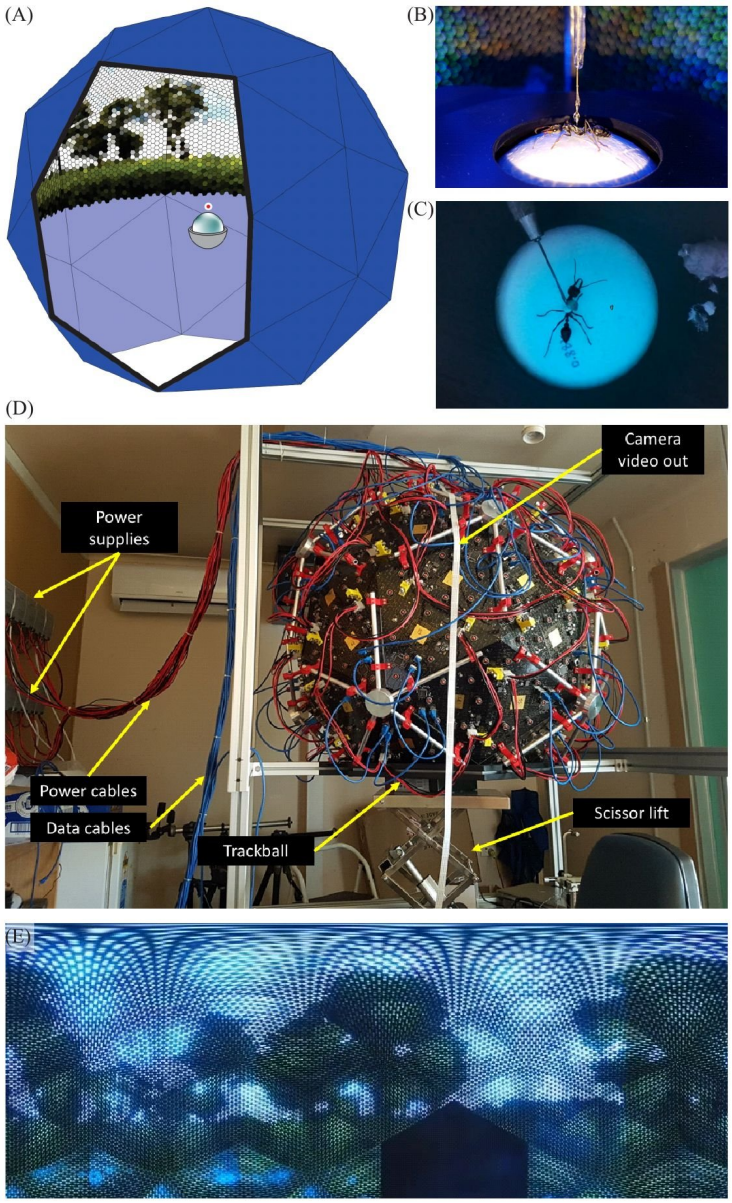


*Fig. 20: The Antarium*

***(A)*** *Concept schematics of the biscribed pentakis dodecahedron with 55 facets of identical isosceles triangles carrying LEDs and control electronics and the trackball device.* ***(B)*** *Thethered ant on an air-cushioned trackball. The ants are free to rotate around the yaw axis, but its translational movements are registered by monitoring the rotations of the styrofoam ball.* ***(C)*** *The tethered ants as seen by the Antarium camera.* ***(D)*** *The fully assembled Antarium.* ***(E)*** *The landscape panorama projected by the Antarium LEDs seen at 1.5° resolution, about twice the average resolution of ants.*

Since PCB manufacturing has a large NRE (non-recurrent engineering) cost, it is significantly cheaper if the polyhedron can be built from identical facets. Facet number is then a trade-off between optical properties and cost, with larger

numbers leading to a better approximation of the sphere, but higher manufacturing and labour costs. To guarantee that each facet has identical properties, i.e. that the LED arrangement can be identical on them, all of the polyhedron's vertices should lie on a sphere.

We chose the biscribed pentakis dodecahedron (*Fig. 20A*) as our spherical approximation for the Antarium. It has 60 facets of identical isosceles triangles. Five triangles form a pentagonal pyramid and 12 of such pyramids comprise the solid. For the Antarium one such pyramid is removed at the bottom, providing an opening where a trackball with the tethered animal can be inserted.

The physical size of the Antarium is constrained by electronic circuit board density, mechanical limitations, and by the need for the opening at the bottom to be sufficiently large for the insertion of the trackball apparatus. With all those factors considered, the length of the base of the triangle was chosen to be 368 mm. All other dimensions are determined by the geometry of the pentakis dodecahedron, resulting in a roughly 1m diameter device (*Fig. 20D*).

***Pixel arrangement***

Ideally the LEDs should be as evenly distributed on the surface of the polyhedron as possible, which is challenging, because the pattern continuity between adjacent panels needs to be addressed. A pattern was found where the LEDs are on the vertex points of a hexagonal lattice. A computer program was written that calculated the pixel positions and minimised the inter-pixel angle variation while taking the technological constraints of manufacturing into account.

Two such hexagonal grids were calculated, one for the GB (green/blue) pixels and another for the UV pixels. The angular acceptance function and the spacing of ommatidia in the dorsal rim area are much higher than in the rest of the eye. It was decided that the UV LED pattern therefore should be made significantly sparser than the BG pattern, especially because of the high cost of UV LEDs and the need for their adjustable polarisation.

Each triangle contains 361 blue/green pixels and 9 UV pixels (*Fig. 21A,B*). Therefore, the 55 triangles that form the Antarium all together have 19,855 GB pixels and 495 UV pixels. Because no spectral sensitivity information was available at the time, the LEDs were chosen based on their price, availability, physical size, brightness, and beam angle. The selected LEDs were LTST-C930KGKT (Lite-On, Inc.), LTST-C930TBKT (Lite-On, Inc.) and VLMU3100 (Vishay) for the green, blue and UV, respectively. As can be seen in *Fig. 20C*, the current LEDs' spectral emissions are ill matched to the photoreceptor spectral sensitivities that have since been determined in *Myrmecia* ants (Ogawa et al., 2015). This problem will be fixed in Antarium Mark-II, which is currently under construction (see Outlook section below).

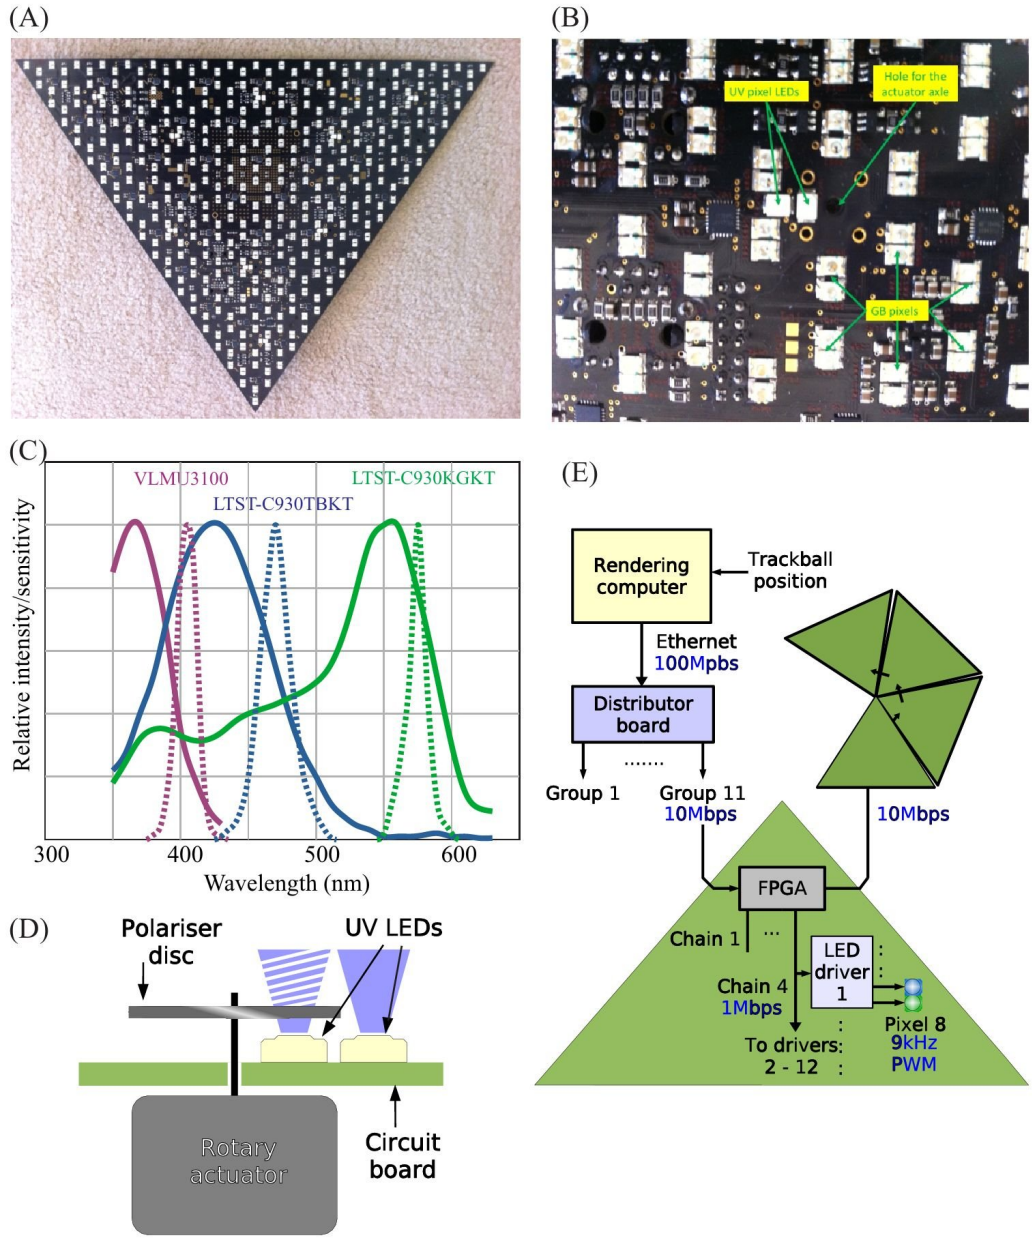


*Fig. 21: The design of individual Antarium panels*

*(A) Photograph of one of the panels with LEDs seen as white rectangles. (B) Detail of the panel LED locations on the printed circuit board and the actuator axle location for polarizer disks. (C) Spectral sensitivities of Myrmecia ants compared with current LED emission spectra. Continuous lines: Normalized spectral sensitivities of the nocturnal Myrmecia vindex recorded intracellularly (redrawn from Ogawa et al. 2015). Dotted lines: Emission spectra of the LEDs used in the current version of the Antarium as per manufacturer specifications. (D) Schematic of the way in which light polarization is achieved. (E) The data path of the Antarium.*

Preliminary experiments revealed substantial internal reflections within the Antarium, which were subsequently minimized by fitting a low reflection black cardboard cover to its internal surface. We measured the reflectance of the black cardboard with a USB-4000 Ocean Optics spectrometer against a certified reflectance standard reference from LabSphere illuminated by natural light. For all wavelength points the cardboard intensity was divided by the reflectance standard's intensity. Between 400 nm and 700 nm the cardboard reflects between 5% and 7% of the light, without dips and peaks.

***Polarisation***

The adjustable polarisation of the UV LEDs is based on each UV pixel being composed of two UV LEDs (*Fig. 21D*). One of them is not polarised at all. The other one is placed behind a linear polariser. The polariser is a small disc mounted on an axle of a rotary actuator. The actuator can rotate the disc and therefore its plane of polarisation can be at any angle. By varying the relative intensities of the polarised and unpolarised LEDs, the polarisation depth can also be controlled.

The actuator needs to be fast as it must to be able to follow scene changes. Stepper motors and servos are too slow. The chosen actuator is an aircore, comprising of a small permanent magnet rotor and a stator with two coils arranged orthogonally. The combined magnetic fields of the two coils can have constant strength but set to any direction by driving one coil with a current that is proportional to the sine of the desired angular position while the other with its cosine. The permanent magnet rotor will always align with the magnetic field direction. Because the rotor is low mass, an aircore can be driven into a new position quite fast. It has a tendency of oscillations while it settles, but manufacturers also offer devices with a small droplet of silicone oil in the rotor bearing. The oil acts as a damper and the time constant of the damping depends on the viscosity of the oil used. With the correct viscosity, the settling can approach the theoretical optimum. The chosen aircore, MicroAirCore 2022-715 from Simco, Ltd. was tested in the laboratory and it was fast settling, with very little oscillation. A 180$^{o}$ rotation can be achieved in less than 200ms.

***LED driving***

To guarantee constant brightness the LEDs must be driven by a constant current source. The brightness of an LED is a function of the current flowing over it. LEDs are semiconductor diodes with nonlinear I-V characteristics. In addition, like with all semiconductor devices, the characteristics are dependent on the temperature of the chip. Although a laboratory is usually an air-conditioned room, LEDs generate waste heat which warms them up. An LED that was bright for a while will be significantly warmer than one that ran at low intensity.

To mimic natural conditions, the intensity range of the arena should span close to 5 log units. A 16-bit linearly spaced intensity regime (65536 levels) corresponds to 4.8 log units. We used a commercially available LED driver chip, the MBI5040 from Macroblock which satisfies all these criteria. It can drive 16 LEDs with constant current. It uses a 16-bit pulse-width modulation (PWM) scheme to set the intensity of each LED individually. It can also apply a correction scheme to compensate for LED brightness variation. The correction scheme can vary the drive current from 0 to the nominal maximum in 1% steps for each LED separately. In addition, it can detect and report short circuit and open circuit LED failures. Furthermore, the chip can operate with only a 0.5 V drop across its driving circuitry, an important feature from a power consumption point of view. The maximum drive current is 30 mA per LED; the LEDs used in the Antarium use only 20 mA drive current, far below the chip's limits.

There are 361 BG and 9 UV pixels on a triangle and the MBI5040 can drive 16 LEDs (i.e. 8 pixels), therefore each panel contains 47 chips.

***Flicker considerations***

Using PWM to set the LED brightness introduces flicker. PWM works by turning the LED full brightness for a short time then completely dark for some other time; the average intensity is the ratio of the ON time and the PWM period (the sum of the ON and OFF times). Thus, the LED flickers with the PWM period. Using discrete time increments, the number of levels that can be displayed is the number of increments per PWM period. To ensure ants do not see the flicker, the Antarium needs a flicker frequency of 300 Hz or more. Thus, the

PWM period needs to be no more than 3.33ms which with 65,536 levels gives an elementary time increment of 50.86 ns, and a clock frequency of 19.7 MHz. We chose to run the PWM on a 20MHz clock, even though the MBI5040 chip could run on up to 30MHz.

However, another method allows us to reach a much higher flicker frequency far beyond what would be detectable by any biological system. The MBI5040 implements what is called scrambled PWM, a scheme designed to increase the flicker frequency above the PWM period. Instead of turning the LED on for the ON time then extinguishing it for the OFF time, the scheme spreads those times around within the PWM period. For example, if the period is 10 time units and the LED has a brightness of 30%, a simple PWM will turn it on for 3 units then off for 7 units. However, a scrambled PWM system might turn the LED on for 1 unit, then off for 2 units, on for 1, off for 2, on for 1, off for 3. Since the LED was on for 3 units and off for 7 the average brightness is still 30%, but now the LED blinked 3 times during the period instead of once. There are various ways to perform the spreading. The MBI5040's method becomes active when the brightness level increases above 32 units out of the 65,536. The Antarium uses a 20 MHz clock, thus if the LED brightness is higher than 0.05% of full scale, the flicker frequency will be more than 9 kHz, while below this threshold, for very dark LEDs, the flicker will be 305 Hz. Photodiode tests using an oscilloscope confirmed flicker at 9 kHz.

***Video delivery and frame rate***

Since the Antarium's LED array is simply a display device, the method of data delivery from the rendering computer must be defined. All together the Antarium has 20,350 pixels, each of which needs 2x16 bits of data to set the brightness, giving a total of 651,200 bits per video frame. The most common communication links on a computer are USB and Ethernet. When the Antarium was designed, the fastest USB was 450 Mbps (USB-2.0 full speed), the next step down was 12 Mbps (USB-2.0 high speed). The most common Ethernet interface was the so-called 100BASE-TX, delivering 100 Mbps over the ubiquitous "blue cable" (officially named Category-5 twisted pair cable). Full-speed USB interface chips were not readily available at the time and the high-

speed USB was simply not fast enough. We therefore chose the 100 Mbps Ethernet link as the delivery medium for the video stream.

If a full frame is 0.6512 Mbits, then the 100 Mbps link has a theoretical limit of 153 frames per second. In reality it is less, as there are protocol overheads. That does not meet our goal of 200 fps and so we needed to find ways to compress the video stream.

The compression scheme must be relatively simple so that the panels of the Antarium can decode it and so that any computer can encode it without special hardware. The solution we chose is to subsample the colour information. Instead of delivering 16-bit resolution green and blue values for a pixel independently, a 16-bit luminance value and an 8-bit chromaticity value can be delivered. That saves 25% of the video bandwidth (24 bits per pixel instead of 32). It does not compromise the 4.8 log unit brightness range, however it does limit each pixel to 256 available hues.

The simplest way of sending data from a computer over an Ethernet link is using a standard protocol that is supported by any operating system. One of those is UDP (user datagram protocol), where blocks of data (packets) are sent from one machine to another. UDP is advantageous in that it has smaller overhead than other protocols. On the other hand, it does not guarantee delivery and gives no feedback whether the packet ever arrived. UDP is often used in situations where occasional loss of a packet is acceptable, but the unpredictable delays arising from confirming the reception of every packet and re-sending lost ones is not. These strengths and limitations are well suited for video streaming, since if a single video frame gets lost, most of the time the observer will not even notice. Whereas if the streaming stopped while the sender and receiver negotiate the retransmission of a single packet, the video quickly becomes unwatchable. The Antarium therefore uses UDP for video delivery, with a dedicated Ethernet link to ensure that packet loss is rare.

An Ethernet frame contains up to 1500 bytes of actual data (usually called the payload) and a further 38 bytes of addressing, synchronisation and other ancillary information. Furthermore, UDP adds 24 bytes of protocol information to the data portion of the packet. The protocol overhead is thus 62 bytes for each

Ethernet frame with a UDP packet in it. In a full video frame a single Antarium triangle is represented by 1,110 bytes. Two extra bytes are added to the raw data, for reasons explained later. Therefore, the payload is 1,112 bytes. If each packet contains one triangle's worth of video information, then 1174 bytes need to be transferred per triangle. A video frame contains 55 such Ethernet frames, resulting in a maximum theoretical video rate of 194 fps over a dedicated Ethernet link. Indeed, in practice the Antarium sustains around 190 frames per second.

***Architecture***

Driving the 9 polarisation actuators exceeds the capacity of available microcontrollers, so the Antarium's panels are equipped with a field programmable gate array (FPGA) instead. The processing unit of each triangle must receive video frames and send the brightness data to the 47 LED driver chips. In addition, it must control the drive current of the 9 actuators for the polarisers which each have 2 coils (18 total drive lines). Using pulse width modulation (PWM) to set the current necessitates a device with 18 PWM units which no commercially available microcontroller can support. Instead, we chose to use an FPGA. An FPGA is just a large collection of simple digital logic building blocks, which then can be connected to each other inside the chip to form a digital circuit that performs a specific function. Microcontrollers are well suited for tasks that work on fewer hardware signals at a time and where the decision making logic or calculations are complex. For tasks where there are many hardware signals and the calculations and decision making are relatively simple, but must be performed at high speed and with precise timing, FPGAs are often a better choice. The large number of PWM signals makes the FPGA a better solution for the Antarium. As such, each triangle panel contains an XC3S50AN chip from Xilinx, Inc. The chip has 50,000 logic gate's worth of resources and can handle more than 80 input/output digital signals at high speed.

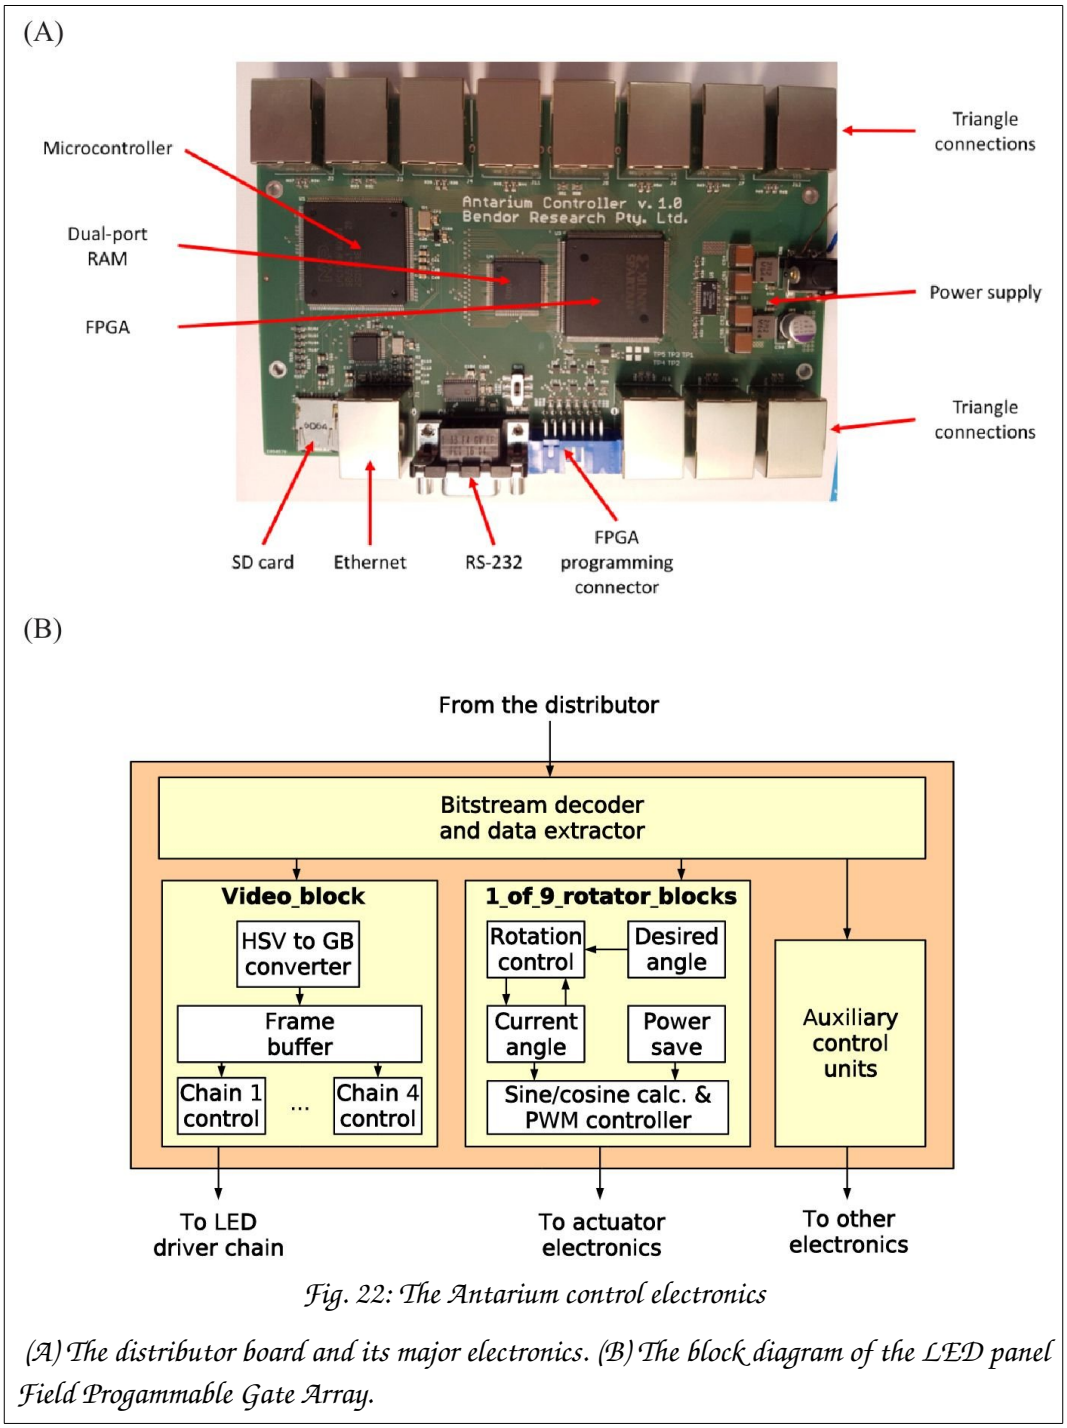


*Fig. 22: The Antarium control electronics*

*(A) The distributor board and its major electronics. (B) The block diagram of the LED panel Field Progammable Gate Array.*

For our triangular panels the FPGA needs to buffer a video frame, decode the compressed chromaticity, send the decoded data to the LED driver chips and run 18 PWM controllers for the actuators, which consumes about 60% of its gates. The remaining 40% are not sufficient to also run Ethernet and UDP protocols as a logic circuit. While we could have used a more powerful chip, the added cost for each 55 panels would have been a significant expense. We

instead chose to design a single interface board, with an associated one-off cost, that receives the video feed from the computer and distributes it to the triangles in a simpler way.

When the FPGA on each triangle panel receives a frame, it decodes the chromaticity encoding and collects the 16-bit intensity values for each LED in a buffer. At the end of the video frame, the buffer is sent to the LED driver chips. The drivers have an SPI (serial peripheral interconnect) interface, a standardised serial bus. The LED driver chips are designed to be daisy-chained. Since very long SPI chains are technically problematic, we divided the LED drivers into four chains. The FPGA delivers the video data to the chips on the four SPI chains simultaneously, which allows us to use a lower speed on the buses.

We use an H-bridge design for the PWM controller of the polariser's actuators, which provides a large reduction in energy usage when the actuators are idle. To drive a single H-bridge the FPGA needs to produce two signals, so for the 2 coils of 9 actuators each, 36 output signals are generated. This design allows energy to be saved since the FPGA reduces the current on both coils by the same factor (thus keeping their ratio, and therefore the angle of the actuator intact) when the actuator is stationary. This holding current is one quarter of the current used for moving the actuator. If the actuator needs to be re-positioned, the FPGA switches the drive current back to nominal and when the position has not changed for a while, it slowly reduces the current to the one quarter holding value.

Finally, we placed thermal sensors on each triangular panel which are also controlled by the FPGA. The data from these sensors can be sent back across the network, which is important given the large amount of heat which can be produced when the full device is running at maximum brightness.

***Power distribution***

Since the Antarium consumes a significant amount of power, ensuring adequate power supply was integral. Each LED needs 20mA for full brightness. A typical blue or UV LED has a voltage drop of around 3.4 V. The driver chip needs an

extra 0.5 V, resulting in a minimum power supply voltage of 3.9 V. To cater for variations and to provide a safety margin, the LED driver circuitry operates from a 4.2 V supply. Due to the use of the intensity/chromaticity encoding, a pixel never needs more than 20 mA. Therefore, a triangle panel's 370 pixels draw 7.4 A. In addition, the driver chips themselves also consume approximately 30 mA from the same supply. With 47 driver chips per panel that adds 1.4 A to the load. The FPGA and its support circuitry need to be supplied as well, although that supply current is negligible compared to that of the LEDs and the drivers. The actuators run from 12 V and the nominal coil current is 54 mA. Due to the sin/cos driving scheme, however, the two coils of an actuator together have a maximum current consumption of 77 mA. The maximum current therefore is 0.7 A.

All together the board needs about 9 A from 4.2 V and 0.7 A from 12 V. The boards have two high-efficiency switch-mode power supplies that generate the 12 V and 4.2 V from a 24 V supply. The efficiency of these supplies is close to 90%, thus the board draws a maximum of 2.13 A from 24 V. Since under no circumstances will all LEDs of all triangles be on full power while all actuators being also set to their most power-hungry position, it was decided that a commercially available 24 V, 10 A power supply unit from MeanWell can safely power 5 triangles forming a pentagon. 11 such units power the Antarium. Power losses on the cabling are minimised by using sufficiently thick wires.

***Thermal considerations***

The Antarium’s maximum power consumption is 2.5 kW, making its heat generation roughly equivalent to a portable oil radiator, enough to warm a small room with a volume of 16 $m^3$. If that thermal energy were concentrated inside the Antarium's less than 1 $m^3$ volume, the temperature would rise to uncomfortably high levels for any subject very quickly. There are three ways to mitigate that risk: reducing the dissipated power, ensuring that heat radiates outwards rather than inwards, and ensuring convection between interior and exterior spaces.

Consumption is minimized due to our use of natural scenes, which are highly varied and contain many dark objects, such as trees trunks, buildings, and

shadows on the ground. Furthermore, to compensate for the intensity variation due to parallax arising out of the Antarium's geometry, the central area LEDs of each panel are artificially darkened. Together these two factors more than halve the overall power consumption.

Unfortunately, most of the heat is generated by the LEDs, which are on the inside of each panel. To minimise the amount of heat inside the Antarium we made use of the fact that each LED is connected to a solid copper plane near the outer surface of the PCB. While normally the thickness of copper in PCBs is 35 µm we used 70 µm copper for the Antarium to improve heat conductance. To further augment each panel's heat conduction, we added a large exposed copper square to the exterior of each panel, which is thermally connected to the inner plane. This allows us to attach a Peltier cooling element with a heatsink and a fan, which can even more effectively suck the heat out and dissipate it. However, after testing the Antarium in its final form it turned out that there was no need for such additional cooling of the panels.

The lack of the need for a cooling element was perhaps facilitated by ensuring good airflow between the interior and exterior of the Antarium. This convection is assisted by a small table fan placed under the Antarium when it is operational, which supplies fresh air into the internal volume and forces the warm air out. In addition, an air-conditioned room helps to keep the internal temperatures at comfortable levels, and also ensures comfortable working temperatures for operators when set to 19°C.

We measured the temperature inside the Antarium at the position where the ant would be on the trackball using a Kestrel 5500 Weather Meter (Kestrel Australia, East Melbourne, Victoria), the room air conditioning set at 19°C and after allowing temperatures to stabilize for 1 hour. The temperature was recorded when it stopped changing over a 3 min period. We measured: Ambient room temperature: 20.5°C on a 26°C day; all LEDs on maximum output, no fan: 61.3°C; natural image, no fan: 28.3°C; natural image, with fan: 25.1°C; ambient room temperature re-tested after the Antarium measurements: 20.5°C. This is well within natural foraging temperatures for both day- and night-active Myrmecia ants (Jayatilaka et al., 2011).

*Fig. 1D* shows the fully assembled Antarium.

***Distributor board***

The distributor board, as its name implies, distributes the video signal to the triangles (*Fig. 21E, Fig. 22*). It contains an LPC1788 microcontroller from NXP, Inc. The microcontroller has an ARM Cortex-M3 core running at 120 MHz, 512 KB internal FLASH and 96 KB internal RAM. It also has built-in peripherals, including an Ethernet protocol engine, an SD card protocol engine, several other serial communication blocks, timers and user-programmable digital I/O ports. Its Ethernet engine, augmented with an external media access controller (TLK110, Texas Instruments) provides the 100Mbps Ethernet interface.

The microcontroller shares its work with an XC3S500E (Xilinx, Inc.) FPGA containing half a million gates worth of logic. Between the microcontroller and the FPGA there is a 128 KB dual-port static RAM chip (IDT70V28L, Integrated Devices Technology). All received Ethernet frames are written into the dual-port RAM. Then the microcontroller decodes the protocol and analyses the packets. Packets related to connection maintenance are processed and responded by the microcontroller. If the packet contains video data, then the microcontroller sends a message to the FPGA that the data should be delivered to a triangle. The FPGA examines the packet data, decides which pentagon it belongs to and queues it for transmission on one of its 11 output links to the pentagons. After delivering the packet to the triangle the FPGA sends a message to the microcontroller informing it that the data is out and the given dual-port RAM region can be released.

If a triangle sends some data, then the FPGA holds the message in a temporary internal storage and signals the microcontroller that a message is available. When the microcontroller indicates that it is ready, the message is passed to it through the dual-port RAM.

The communication between the FPGA on the distributor board and the FPGAs on the triangles uses differential signalling. The data rate is 10 Mbps and the signal is subjected to the so-called Manchester encoding. That data speed and encoding is used by the 10BASE-T Ethernet standard, which facilitates the use

of low cost Ethernet connectors, magnetics and cables. While the data speed and encoding method is the same, the protocol which the Antarium uses is much simpler than Ethernet. Each data frame starts with a preamble, followed by a synchronisation byte, followed by a byte that indicates the type of the packet and its destination (or source) triangle within the pentagon. The next byte contains additional information about the packet content. The data follows and the packet is finished with a 2-byte long data integrity check. That protocol is simple enough so that even the resource-limited FPGAs on the panels can handle it.

The configuration bitstream of the distributor board's FPGA is stored on a micro-SD card. The board has an SD card socket and the microcontroller drives it. The controller implements the SD card protocol as well as the Microsoft FAT file-system, thus the FPGA bitstream can be written to the card using any computer. When the board is powered up, it first reads the SD card and loads the bitstream into the FPGA.

From the TCP/IP network stack, the firmware of the microcontroller also implements the UDP (user datagram protocol), IP (internet protocol) and ARP (address resolution protocol). Those are the necessary and sufficient components to be able to communicate with a machine with a standard network stack, regardless of the operating system it runs.

The distributor board also has a secondary function: to program the FPGAs on the triangles. The FPGA on the distributor board forgets its configuration when it is powered down. When the board is turned on, the microcontroller needs to load the configuration from the SD card. The FPGA on the triangle has a built-in non-volatile storage to hold its configuration, thus it wakes up fully configured. However, the configuration first needs to be programmed into the non-volatile storage. Xilinx offers a free tool to do that, but the tool was slow and unreliable. Fortunately, the programming algorithm could be reconstructed from various application notes (engineering advisory articles). We then created our own implementation of the algorithm on the distributor board and it can program the triangle’s FPGAs in a few seconds, with 100% reliability.

The distributor board is powered from a commercially available 12 V power module (plug-pack). The actual supply voltages for the electronics are generated from that 12 V using an LT3824 (Linear Technology) dual switch-mode regulator. To aid software development and the initial programming of the board also contains an RS-232 serial port.

***Design tools***

All design work was performed on a computer running the open-source GNU/Linux operating system. To aid engineering, several programs were written in-house to calculate or optimise certain parameters, to assist debugging, or to automate tasks. These programs were all written either in the C or in the Tcl language. Tcl/Tk is an open-source, interpreted scripting language with graphical capabilities. C programs were compiled using the open-source gcc tool chain. Building the final binary image or bitstream was controlled by the open-source gmake tool. The open-source Fossil distributed version control system was used to keep track of changes during development.

The schematic entry and the PCB design for the triangles and the distributor board were done using the commercial Eagle EDA package from CadSoft GmbH (recently taken over by Autodesk), version 6.4, professional edition, for Linux. The PCB manufacturing files were visually checked using the gerbv open-source Gerber viewer tool.

The code for the FPGAs was written in the Verilog hardware description language. The logic simulations utilised the Icarus Verilog open-source simulator and the GtkWave open-source waveform viewer programs. Logic synthesis, technology mapping, place-and-route, and bitstream generation were performed by the ISE 14.7 toolchain from Xilinx, Inc. The tool is closed source but Xilinx provides it free of charge.

The firmware for the microcontroller on the distributor board was written in the C language. The code was compiled using gcc in a cross-compiler configuration. The open-source Armlib library from Bendor Research Pty. Ltd. was used for most low-level functions and for the task scheduler. The Ethernet driver, SD

card driver and the FAT filesystem utilised routines donated by Arthur Digital Solutions Kft (Hungary).

The component sourcing, purchasing, PCB manufacturing and assembly were ordered from Albacom Kft. (Hungary). Quality control and thorough testing of the boards before shipment to Australia was performed, gratis, by Arthur Digital Solutions.

The mechanical design and the manufacturing of the scaffolding was done by the ANU workshop. The power cables were manufactured by hand; the Ethernet cables, wires and sundry electronics items were purchased from Jaycar, a local electronics store.

***3D rendering and driver software***

The software that generates the video stream for the projector makes use of the commercially available three-dimensional (3D) rendering engine Unity (Unity Technologies) running in Microsoft Windows. The primary market for the engine is computer games and as such it is best suited for planar projections. The Antarium has a low pixel count compared to most commercial video games and it is, therefore, possible to render 6 or more game views simultaneously at a high frame rate, on modern graphics cards. The 6 views have the same camera position in the 3D virtual world, but the cameras look in 6 orthogonal directions (up, down, left, right, front, back), essentially creating a projection onto a cube. A custom shader uses a spherical transformation known as cube-mapping to map the pixels of our rendered cube onto any arbitrary 3D model. By applying this shader to a 3D model that represents each LED in the Antarium as an individual face, with the same azimuth and elevation as the LEDs real world coordinates, we are able to render the scene as it would appear if projected onto the Antarium. We then use a compute shader to sample each face of our virtual Antarium using its normal as a lookup into the now spherical cubemap (using DirectX SampleLevel function). Finally we encode and package these as pixel data to send over UDP to the distributor board.

The aim of the Antarium is to display views of the natural habitat of the animals. We, therefore, constructed a 3D model of that habitat using camera-based

reconstruction methods, see (Stürzl et al., 2015; Murray and Zeil, 2017). Thousands of photographs were taken with a Panasonic Lumix DMC-FZ200 camera at 4000x3000 pixel resolution while walking around in the area surrounding the nests of the experimental ants. Multiple voxel clouds were created from these photographs with the software Pix4D (Pix4D SA) and exported as 3D models before being combined into a single unified and aligned 3D reconstruction of the ants' foraging environment. Since the very distant panorama does not have enough parallax to be processed by the 3D reconstruction software, we added the distant panorama later as a static background image at 1km (approximately infinite) distance. We captured this panorama with a Ricoh Theta S panorama camera (Ricoh Company Ltd, Tokyo, Japan).

This procedure allows us to capture views from within our 3D model, or from within projections of panoramic photographs, to edit the 3D model (using Blender) or photographs to fix errors (using Paint.net), and finally to generate experimental treatments (using Unity3D). For example, *Myrmecia* ants regularly visit trees for foraging, e.g. (Narendra et al., 2013b) and we are now able to extract such foraging trees from the photograph and the 3D model, allowing us to move the foraging tree to any arbitrary location or bearing in the model/photograph as an ant is viewing the scene inside the Antarium. We can then ask, whether the ants treat trees as individual landmark beacons, or get their bearing from the whole landmark panorama.

***The trackball system***

The ants are placed on an air-cushioned, light-weight, 10 cm diameter track ball (*Fig. 20B*) on which they are free to rotate around the yaw axis but that allowed us to record their intended translational movements as described in detail by Dahmen *et al* (2017), see also (Murray et al., 2020). The trackball sends the position data to the rendering computer using USB. In a departure from the original, we now maintain and compile the trackball code using Microsoft Visual Studio in the C language (Microsoft, Inc. 20XX). The USB connection relies on the open-source usblib library. The system response is linear up to speeds of 1.2 m/s. For detailed system properties see (Dahmen et al., 2017).

Since the trackball is connected to the computer running the 3D engine, we can use the movement data it generates to update the position of our virtual cameras in the 3D world, thus providing our ant subjects with closed loop control of the visual scene. When running in open loop, 3D scenes or panoramas can be presented either statically or in sequence. For closed loop we use Kernel32 to share a file in shared memory between the trackball program and the game engine. In this file we write the current offset of the trackball from its starting location, and accept commands to reset the starting location, such as when a new treatment begins. In both modes the human operator, or their code, can arbitrarily change the ant's virtual position and heading at any time. However, in closed loop mode this trackball offset can be used to update the position of to the six cameras inside the 3D model, thus updating the view that is presented to the ant subject, based on its own movement on the trackball. It should be noted that due to the complexity of this setup significant care must be taken to ensure all real-world and virtual objects are rotationally aligned, so that the visual consequences of the ant's movements are accurately represented.

***Antarium camera***

To record in addition to the ants' intended paths also the scanning movements of their head, we mounted a Raspberry-Pi V1 camera at the apex of the Antarium. The camera is connected to a Raspberry-Pi single-board computer (Raspberry Pi Foundation, UK). It records 1280x960 pixel video at 30 fps to an external USB disk. The recoding format cannot be played back with commercially available software on Windows, thus the recorded footage is transcoded to MP4 format using the open-source ffmpeg package on a Linux computer.

## Proof of concept

To date we have conducted a number of experiments demonstrating that ants recognize familiar scenes in the Antarium and derive navigational instructions from them. We will present these behavioral results in a separate publication.

In brief, we confronted ants tethered on the trackball with four different views (*Fig. 23A*): a familiar view half-way toward a tree along their normal foraging corridor (*Familiar*), the view from the nest (*Nest*), an unfamiliar view from a location about 5 m offset from the foraging corridor (*Unfamiliar*) and a scene that consisted of a horizon line only (*Unstructured*). As the ants walked on the trackball in these four situations, we instantaneously rotated the scenes several times through 90 degrees randomly clock- or counter-clockwise to test whether the insects took note of panorama information. They indeed changed path direction in response to such rotations when confronted with any of the structured, but not the unstructured scenes as shown for two examples of the Familiar scene in *Fig. 23B* (*Familiar*) and *23C* (*Unstructured*), with 15 s long segments before rotations labelled red and 15 s segments after rotations labelled blue. Instances of rotations are marked by a blue dot. Note that the ants' speed is not constant, but indicates that the ants move in spurts (*Fig. 23B, C*) and that their path direction oscillates with smaller amplitudes when confronted with a familiar scene and larger amplitudes when confronted with an unstructured scene.

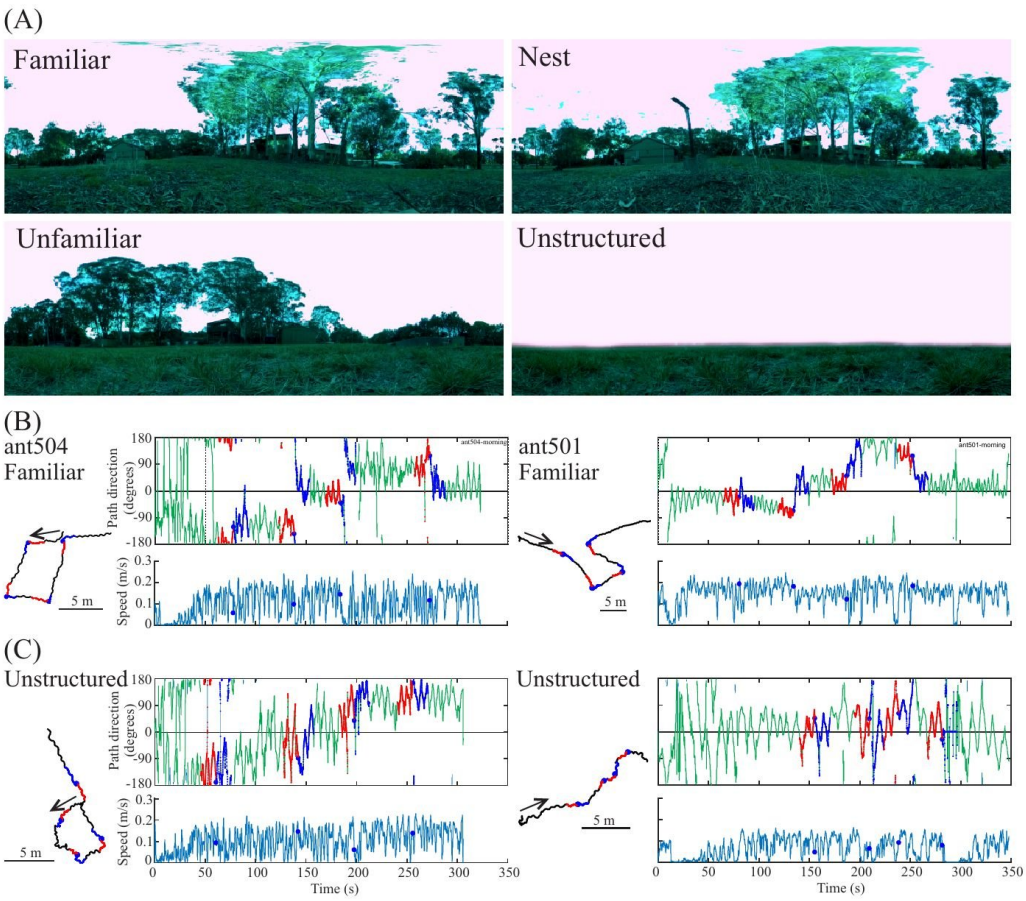


*Fig. 23: Proof of concept experiments*

*(A) Four panoramic views from the ants' foraging habitat. Familiar is located on the ants' foraging corridor, half-way toward their foraging tree; Nest is the view from the ants' nest entrance; Unfamiliar if the view from a location about 5 m to the side of the foraging corridor and Unstructured is a synthetic view without landmark panorama. (B) Two examples (left and right) of ants responding to familiar scene rotations. Instances of rotations are marked by blue dots in the time course of path direction (top) panels and of speed (bottom panels). 15 s segments before (red) and after rotations (blue) are in addition marked on the intended paths of the ants (shown on the left) and on the time course of path direction (top panels). Paths are shown in the track ball coordinate system. (C) Same as (B), but in the presence of the unstructured scene. Note the difference in path direction oscillations in (B) and (C).*

For another example of responses to the familiar scene rotations (*Fig. 24*), we extracted the head- and longitudinal body axis orientation of the ant from the Antarium camera footage 15 s before to 15 s after the rotation (*Fig. 24B*). Following a rotation, the ant's head- and body scanning movements tend to increase (*Fig. 24B*) as she changes her heading direction in the three instances in which she responded to the rotation.

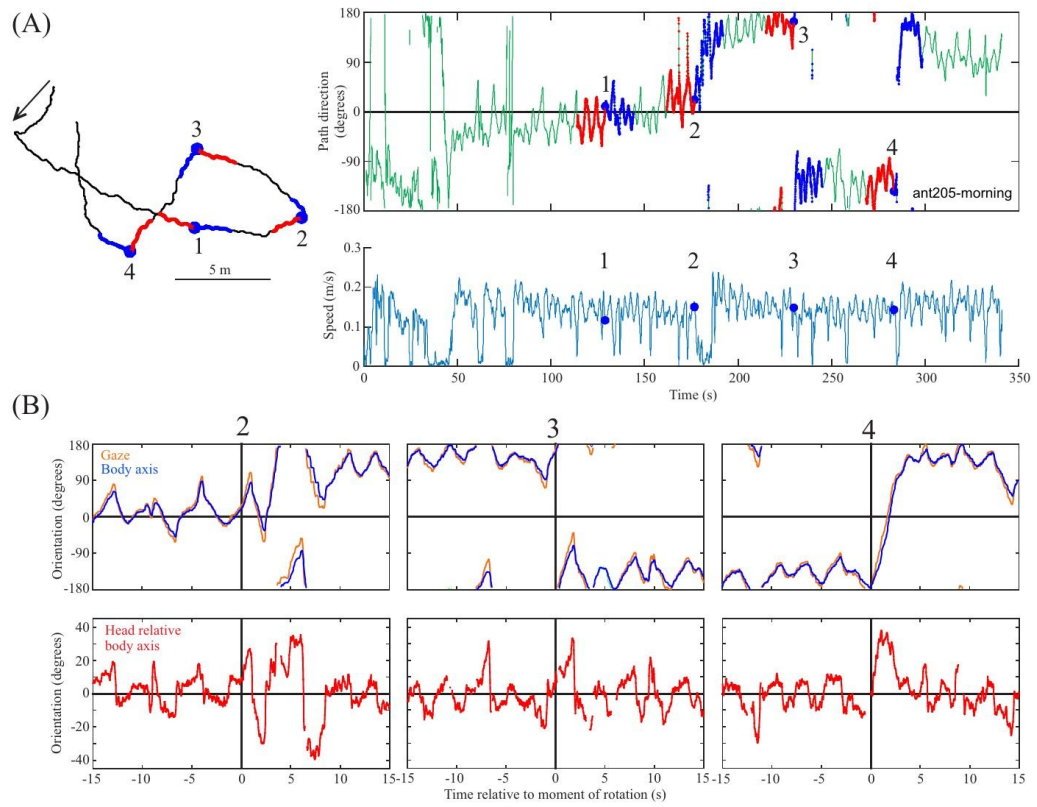


*Fig. 24: Proof of concept experiments.*

*(A) The path (left), time course of path direction (right top) and time course of speed (right bottom) for an ant in the presence of the familiar view. Successive instances of scene rotation are marked by blue dots and numbered. Otherwise conventions as in Fig. 4. (B) Top row: Gaze (head, orange) and longitudinal body orientation (blue) over time from 15 s before and 15 s after rotation 2-4. Bottom row: Head orientation relative to longitudinal body axis for the same segments. Vertical black line marks the moment of rotation.*

## Outlook

The Antarium is a unique reconstructed visual reality arena for ants. No projection system before it has offered a completely panoramic projection tuned to an insect's vision, including arbitrary polarisation patterns. Furthermore, the Antarium can deliver accurate recreations of the visual reality of animals, by projecting imagery captured from their natural habitat rather than artificially generated scenes, e.g. (Stowers et al., 2014; Kaushik et al., 2020). We see the ability of presenting natural views that are familiar to an insect as an important condition for answering many questions about the neural mechanisms underlying visual navigation.

The Antarium not only allows us to compare responses to familiar and unfamiliar natural scenes, but we can also add, remove or dislocate landmarks, set up conflicts between different visual information (i.e. celestial vs terrestrial), and manipulate the intensity, the colour, or the spatial frequency composition of scenes. In closed loop, we can investigate the dynamics of visual navigation, such as the relationship between navigational decisions and scanning movements, or the frequency with which ants check and update their heading direction.

Since the initial conception of the Antarium, many advancements have been made, both in the development of LEDs and in our knowledge of the neural and visual systems of ants. These advancements combined with lessons from our experiments with the Antarium, have led us to design a second version, the Antarium Mark II to improve upon the original. For instance, we now know that the spectral sensitivities of *Myrmecia* photoreceptors in both day- and night-active species have peak sensitivities around 375, 430 and 550 nm (*Fig. 21C*); (Ogawa et al., 2015). As LEDs with expanded emission in the UV range have become available and have dramatically decreased in cost, we can now much more precisely match LEDs to ant spectral sensitivities and increase the density of UV LEDs. Antarium Mark II will thus provide much improved UV contrast of the landmark panorama, which has been shown theoretically and in behavioural experiments to be important for providing information on heading direction, e.g. (Möller, 2002; Graham and Cheng, 2009a, 2009b; Kollmeier et al., 2007; Differt and Möller, 2015; Stone et al., 2014, 2016; Schultheiss et al., 2016a).

***Conflict of Interest***

ZK is a director of Bendor Research Pty. Ltd., an embedded systems consultancy company. To keep cost low, the company's existing contacts were utilised in the component sourcing and manufacturing of the Antarium and its test equipment was used during development. Bendor Research has not charged any fees for its services and had no gain, financial or otherwise, from providing them. All other authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

### *Author Contributions*

ZK: Developed design, architecture, electronics & software; supervised manufacturing; conducted tests & behavioural experiments; wrote the first draft of the manuscript. TM: Conducted 3D modelling & developed rendering software and interface pipeline to the Antarium; conducted tests & behavioural experiments. H-JD: Provided trackball system & training. AN: Conducted initial experiments with trackball system & provided funding. JZ: Project supervision & provision of funding. All authors contributed to several revisions of the manuscript.


### *Funding*

We acknowledge financial support from Australian Research Council (ARC) Discovery Project Grants (DP150101172 & DP150102699), an ARC Future Fellowship (FT140100221), the Hermon Slade Foundation (HSF 10/7), the Australian National University Endowment Fund and private funding. We are grateful to Ken Cheng for providing initial funding during the start phase of this project. TM was supported for part of the work by the Australian Government, via grant AUSMURIB000001 associated with ONR MURI grant N00014-19-1-2571.

### *Acknowledgments*

We are grateful to the mechanical workshop of the Research School of Biology for constructing the housing for the trackball contraption and the scaffold for the Antarium panels, and to Jesse Wallace for providing the raspberry camera.

# Outdoor trackball experiments

## Introduction

While the Antarium was designed and being built, a series of experiments were performed outdoors. The aim was two-fold: on the one hand, they were to verify that the trackball apparatus worked as expected and on the other, to collect reference data against which the behaviour of the animals in the Antarium can be compared.

The trackball system was supplied by Hansjürgen Dahmen. Its design is described in detail in (Dahmen et al., 2017). The device has two versions, with the only difference between them is the diameter of the ball, 5cm and 10cm, respectively.

During an initial field trial of the trackball it came to light that the very low weight ball floating on an air cushion and having a relatively large exposed surface area is too sensitive to wind pressure. Even a light breeze would wobble the ball and a small gust of wind could completely dislodge it from its air cup. The ANU workshop was commissioned to build enclosures for the trackball units so that only the top of the ball is exposed. Subsequent trials have proven that with the enclosure the trackball was insensitive to wind.

The outdoor experiments were performed at the ANU Campus Field Station in 2017 to 2019. *Myrmecia croslandi* was used; a diurnal ant with a body length of approximately 10-15 mm. Due to the body size, the smaller (5cm) ball trackball unit was selected.

This chapter contains the outdoor experiments described above. Furthermore, to clarify an earlier finding, another experiment, without the trackball, was devised and performed, also described in this chapter.

## Published work with the trackball

The experiments verified the trackball apparatus and supplied the expected behavioural baseline for the Antarium work; that baseline will be detailed in the next chapter. In addition, the experiments supplied enough data for a publication not directly related to the Antarium:

*Murray, T., Kócsi, Z., Dahmen, H., Narendra, A., Le Möel, F., Wystrach, A. and Zeil, J. (2020)* The Role of Attractive and Repellent scene memories in ant homing (*Myrmecia croslandi*). *The Journal of Experimental Biology 223:3:jeb210021*.

The following table shows the contribution of the authors:

| | TM | ZK | HD | AN | FL | AW | JZ |
|---|---|---|---|---|---|---|---|
| **Conceptualisation** | • | • | | • | | • | • |
| **Methodology** | • | • | • | • | • | • | • |
| **Software** | • | • | • | | • | • | • |
| **Validation** | • | • | | | | | • |
| **Formal analysis** | • | • | | | • | • | • |
| **Investigation** | • | • | | • | • | • | • |
| **Resources** | | | | | | | • |
| **Original text** | • | | | | | • | • |
| **Review text** | • | • | • | • | | • | • |
| **Visualisation** | • | | | | | • | • |
| **Supervision** | | | | | | | • |
| **Project administration** | | | | | | | • |
| **Funding acquisition** | | | | | | | • |

*Table 3: Author contributions to the published paper*

*The authors' name abbreviations are: TM, Trevor Murray; ZK, Zoltán Kócsi; HD, Hansjürgen Dahmen; AN, Ajay Narendra; FL, Florent Le Möel; AW, Antoine Wystrach; JZ, Jochen Zeil.*

The final manuscript of the article, with only minor formatting differences, is provided below.

## Abstract

Solitary foraging ants rely on vision when travelling along routes and when pinpointing their nest. We tethered foragers of *Myrmecia croslandi* on a trackball and recorded their intended movements when the trackball was located on their normal foraging corridor (on-route), above their nest and at a location several meters away where they have never been before (off-route). We find that at on- and off-route locations, most ants walk in the nest or foraging direction and continue to do so for tens of metres in a straight line. In contrast, above the nest, ants walk in random directions and change walking direction frequently. In addition, the walking direction of ants above the nest oscillates at a fine scale, reflecting search movements that are absent from the paths of ants at the other locations. An agent-based simulation shows that the behaviour of ants at all three locations can be explained by the integration of attractive and repellent views directed towards or away from the nest, respectively. Ants are likely to acquire such views via systematic scanning movements during their learning walks. The model predicts that ants placed in a completely unfamiliar environment should behave as if at the nest, which our subsequent experiments confirmed. We conclude first, that the ants' behaviour at release sites is exclusively driven by what they currently see and not by information on expected outcomes of their behaviour. Second, that navigating ants might continuously integrate attractive and repellent visual memories. We discuss the benefits of such a procedure.

## Introduction

Navigation on a local, in contrast to a global, scale involves travelling along routes and pinpointing places, e.g. (Zeil, 2012). Much evidence has accumulated to show that ants form visual memories of how the scene looks along routes e.g. (Wehner et al., 1996; Wystrach et al., 2011a; Mangan and Webb, 2012) and that alignment matching (Zeil et al., 2003; Collett et al., 2013) between memorized and currently experienced views provides robust information on heading direction (Graham and Cheng, 2009b; Baddeley et al., 2011; Narendra et al., 2013b; Zeil et al., 2014b). Heading direction can be recovered, even from locations at some distance from familiar locations, by

detecting the minimum of the rotational image difference function resulting from a comparison between current and memorised views (rotIDF), (Zeil et al., 2003; Stürzl and Zeil, 2007; Philippides et al., 2011; Narendra et al., 2013b; Stürzl et al., 2015). This minimum provides a measure of familiarity in addition to a heading direction (Baddeley et al., 2011, 2012; Graham et al., 2010).

Before becoming foragers, ants perform a series of learning walks around the nest during which they alternate between turning to look in the nest direction (Müller and Wehner, 2010; Fleischmann et al., 2016; Fitak and Johnsen, 2017; Fleischmann et al., 2018a, 2018b) and in directions away from the nest from different compass directions (Jayatilaka et al., 2018) reviewed in (Zeil and Fleischmann, 2019). It is attractive to assume that the ants store snapshots during these turns whenever they are aligned parallel to the home vector, that is, when they are facing toward or away from the nest direction (Wehner and Muller, 2010; Graham et al., 2010; Jayatilaka et al., 2018), as this is theoretically sufficient for returning ants to align with and walk into the direction of the most familiar of nest-directed snapshots in order to pinpoint the nest (Graham et al., 2010; Baddeley et al., 2012; Wystrach et al., 2013).

Such visual ‘alignment matching’ (Collett et al., 2013) explains well how ants recover the correct direction when on their familiar route (Wystrach et al., 2011b, 2011a; Baddeley et al., 2012; Kodzhabashev and Mangan, 2015). Moreover, nest-directed views acquired during learning walks, reviewed in (Collett and Zeil, 2018), can also provide guidance from locations that are unfamiliar to ants and that can be 10-15m away from the nest in open forest habitats (Narendra et al., 2013b; Stürzl et al., 2015), although the initial movements of released ants may not be directed toward the nest (Zeil et al., 2014b), but toward a familiar route (Collett et al., 2007; Wystrach et al., 2011a).

Overall, this line of work has led to the suggestion that visually navigating insects would only need ‘procedural knowledge’ about knowing where to go rather than requiring a more sophisticated representation of their spatial environment that would allow them ‘to know where they are’ (Collett et al., 2002; Wehner et al., 2006; Cheung et al., 2014; Graham and Philippides, 2017).

To test this directly, we positioned ants that we had tethered over a trackball at different locations in their natural foraging environment, including above their nest, and recorded their intended direction and distance of movement. Ants mounted on the ball were well oriented towards the nest at both on and off route locations, but displayed a search pattern when above the nest, as if they knew they were at the nest, implying a sort of positional rather than just procedural knowledge. Using a simple agent-based-simulation we show, however, that these results can be more parsimoniously explained by alignment matching involving continuous integration of attractive and repellent visual memories, acquired when facing respectively towards and away from the nest during learning walks.

## Materials and Methods

### *Ants and experimental site*

We worked with foragers of the Australian Jack Jumper ant *Myrmecia croslandi* from a nest in the Australian National University's campus field station (-35° 16' 49.87''S and 149° 06' 43.74''E). The ants are day-active, visually navigating solitary foragers that hunt for insects on the ground at up to 4 m distance from the nest and on trees, about 12 m away from the nest where they also feed on sugar secretions of plant-sucking insects (see centre panel top row, *Fig. 26*). For details of the foraging ecology and the navigational abilities of these ants see (Jayatilaka et al., 2011, 2014; Narendra et al., 2013b; Zeil et al., 2014b). During February to March 2017 and December 2017 to March 2018, between 9:00 and 15:00, we caught foraging ants either at their foraging trees about 12 m from the nest in a 'full vector' state (FV, n=10) or at the nest in 'zero vector' state (ZV, n=18), offered them sugar water solution to feed on before immobilizing them on ice for up to 15 min and tethering them to a metal pin by their mesonotum (thorax) using Bondic liquid plastic welder (Biochem Solutions, Ellerslie, New Zealand). The ants were placed on an air-cushioned light-weight, 5 cm diameter track ball (*Fig. 25A*) on which they were free to rotate around the yaw axis but that allowed us to record their intended translational movements as described in detail by Dahmen *et al* (2017). We placed the trackball contraption with a tethered ant at each of three locations in a random order (Fig.

2, top row centre panel): 6.5 m west of the nest where none of the ants were likely to have been before (Off-route), 6.5m south of the nest, half-way along their normal foraging route towards trees (On-route), and directly above the nest (Nest).

We recorded the intended movement directions and distances on the trackball at each displacement for up to 10 minutes, before shifting the trackball contraption together with the tethered ant to the next location. Ants were carefully un-tethered and released close to the nest following the three displacements.

To demonstrate the foraging patterns of ants at this nest and the full range of learning walks, we show the paths of foraging ants, ants that performed learning walks and ants that were released after contributing to unrelated experiments that were recorded with Differential GPS over two years (*Fig. 26*, top row centre panel and *Fig. 31*); for details see (Narendra et al., 2013b). In brief, coloured flag pins were placed on the ground approximately 20 cm behind a walking ant at fairly regular intervals, carefully avoiding disturbing her progress. The resulting pin trail was subsequently followed with the rover antenna of a Differential GPS system, recording the position with an accuracy of better than 10 cm.

***Data analysis***

We recorded trackball rotations due to the intended translation of the ants at 275 fps, which reflect the direction and speed of the ants' intended movements. We present the reconstructed paths, final bearings, changes in walking direction and path lengths for the first 5 min of recordings at the three displacement locations. With the exception of one ant at the off-route location, all ants reached this criterion. We used the Matlab (MathWorks, Natick, MA, USA) circular statistics toolbox (by Philipp Berens) to perform Rayleigh's test for non-uniformity on directional data and Wilcoxon rank-sum tests on differences between displacement locations using the ranksum function in Matlab. For comparisons between all three locations we applied a Bonferroni correction with a resulting critical value for individual tests of p=0.0167.

### *Agent-based modelling*

*Reconstructed world and ant views*

We rendered panoramic views within a 3D model of the ants' natural environments that was previously reconstructed at the ANU Campus Field Station using a laser scanner and camera-based methods (Stürzl et al., 2015). We down-sampled the rendered views to 360×180 pixels, that is, 1°/pixel resolution to roughly match the resolution of the ants' compound eyes. Note that the 3D model was obtained 3 years before the treadmill experiments were conducted, so that there will be some changes to the landmark panorama, in particular involving the canopy, while all the major geometric relationships of dominant visual features such as trees will have remained the same.

*Memorised views and current familiarity*

The agent is assumed to have stored a collection of memorised views around the nest during learning walks and along their normal foraging route (*Fig. 25B*). During tests, the agent is computing a value of visual familiarity at each time step by comparing the current view to its memory bank. This is achieved by calculating the global root mean squared pixel difference (Zeil et al., 2003) between the current view and each of the views in the memory bank, and keeping the value of the lowest mismatch, as is typically done in models and studies of ant navigation (Wystrach et al., 2011b, 2011a; Baddeley et al., 2011, 2012; Philippides et al., 2011; Narendra et al., 2013b; Zeil et al., 2014b; Stürzl et al., 2015). Because high mismatch values indicate a large discrepancy between the current and a memorized view, the value indicates the current unfamiliarity score rather than a familiarity score. Note that in the insect brain, the activity of the mushroom body output neurons (MBON) also correlate with unfamiliarity rather than familiarity (Owald et al., 2015; Felsenberg et al., 2018). Importantly, views in this model are not rotated, but compared only at the facing direction of the current and memorized views. That is, the agent does not need to stop and scan because only one view is compared for each step.

*Combining attractive and repellent visual memories*

The novel aspect of this current model is that the agent is assumed to have two independent memory banks (*Fig. 25B-D*): one containing attractive views and one containing repellent views. Both memory banks contain views experienced

during learning and foraging walks; the attractive memory bank containing views that are assumed to have been memorised when the ants were oriented toward the nest and the repellent memory bank those that have been memorised while looking away from the nest. This is motivated by the very regular scanning movements of ants during their learning walks where they alternate looking towards the nest and away from the nest direction (Jayatilaka et al., 2018; Zeil and Fleischmann, 2019). For simplicity, learning walk views were assumed to have been acquired within a 1 m radius around the nest and we chose a 10 m long route, corresponding roughly with the foraging corridor of this particular nest (see *Fig. 26*, top centre panel). Both nest-directed (attractive) learning walk views and views away from the nest (repellent) were taken from positions along a spiral rather than a circle around the nest (*Fig. 25B*), to mimic the fact that successive learning walk loops reach increasing distances from the nest, e.g. (Fleischmann et al., 2016; Jayatilaka et al., 2018), and to ensure that results at the nest were not dependent on having views memorised at the exact same distance from the nest. We also included in the attractive memory bank views that foragers experience when travelling back to the nest along their normal foraging corridor (*Fig. 25B*).

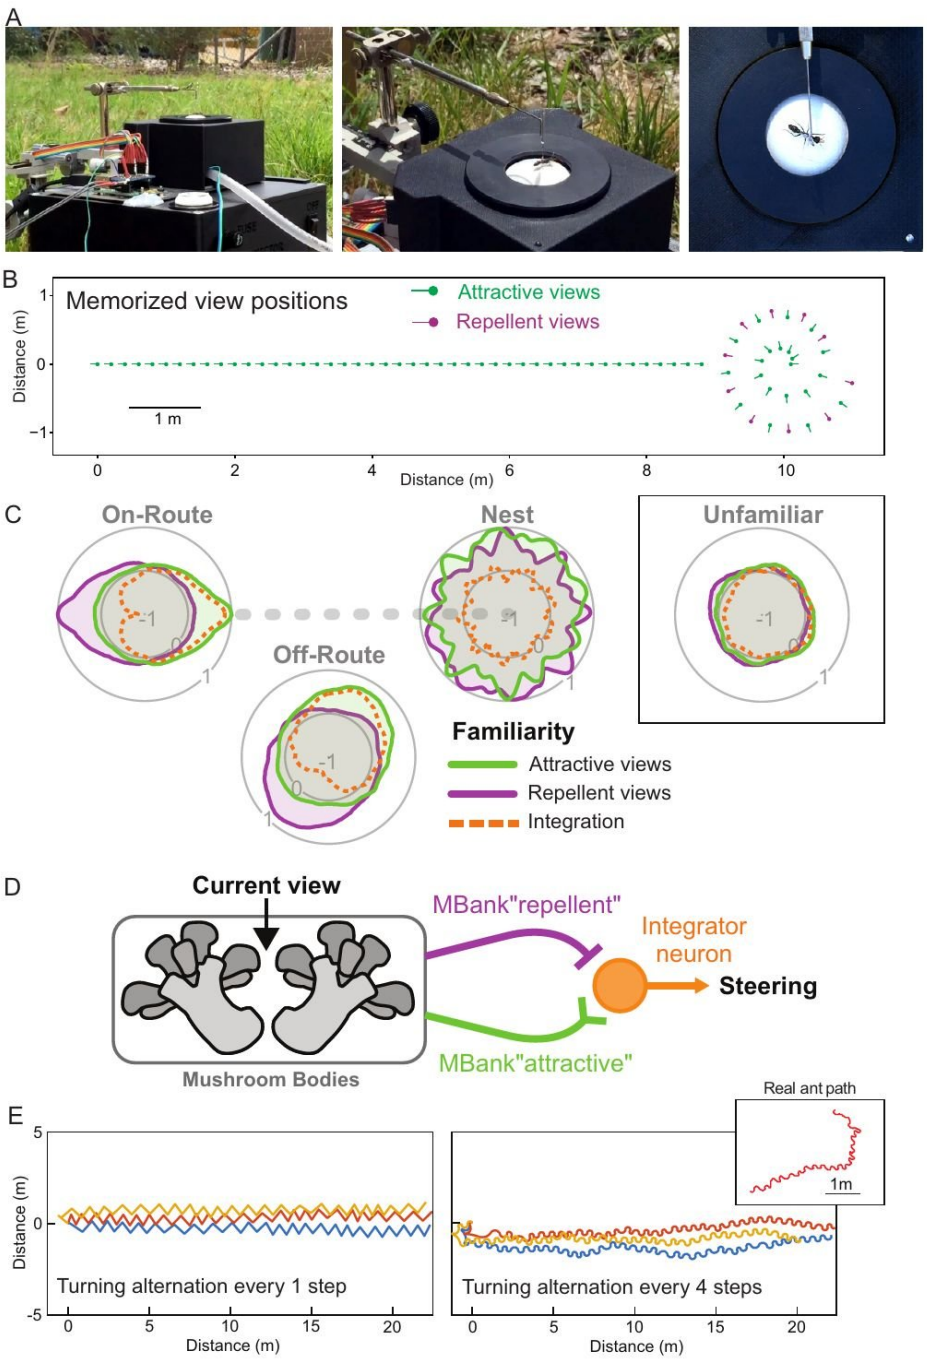


*Fig. 25: Experimental set-up and agent-based modelling*

*(A) Three views of the air-cushioned trackball contraption and the tethered ant. (B) Schematic map of the attractive (attractive) and repellent (repellent) memorized views along the foraging route and around the nest that constituted the attractive and repellent memory bank. (C) Schematic distribution of familiarity (1) and un-familiarity values (-1) for attractive and repellent views at the four release locations and the result of their integration. Note that distributions at the nest and at the completely unfamiliar site are uniform for different reasons: high familiarities for both attractive and repellent views at the nest and low familiarities for both view sets at the completely unfamiliar site. (D) A 'neuro-schematic' summary of the model comparing a current view with a repellent and a attractive view memory bank and the integration of the output providing a steering command. (E) The paths generated by the simulation reproduce the details of real ant paths better when the regular alternation of path direction is implemented at every 4th step, rather than at each successive step (as has been done in the present study).*

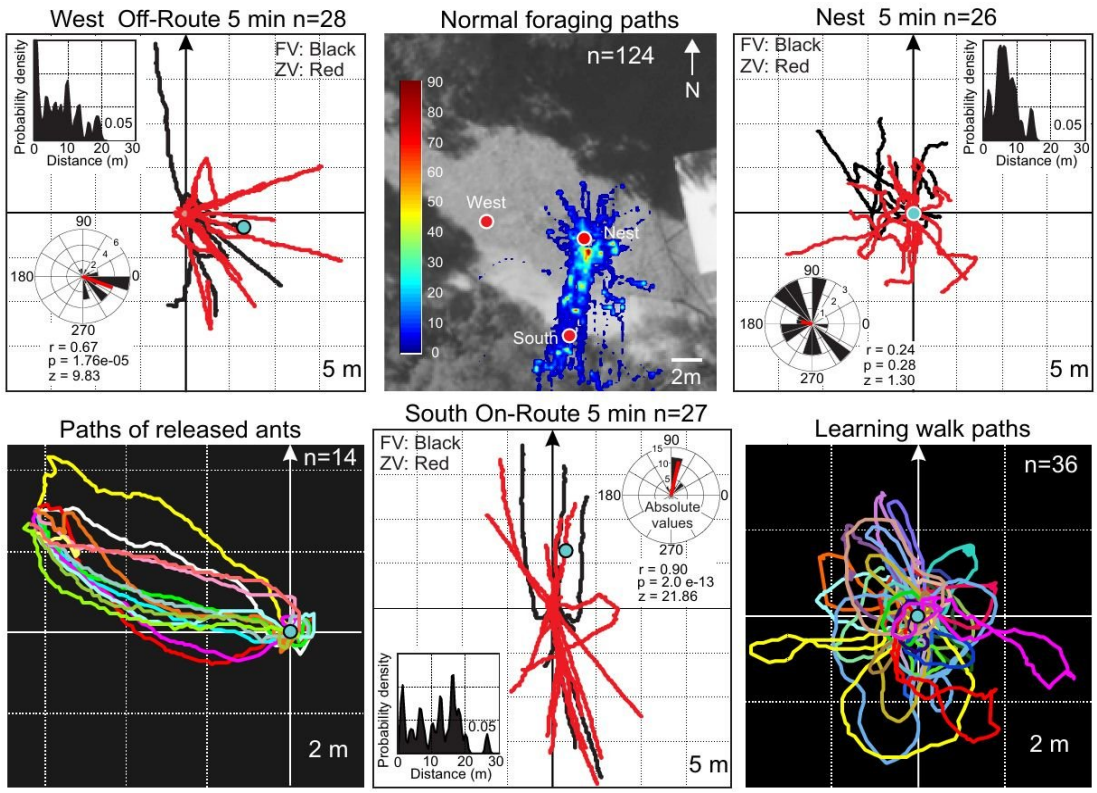


*Fig. 26: The behaviour of tethered ants at three locations in their natural foraging environment*

*Top centre panel: Aerial photograph of the nest area with Off-route, Nest and On-route locations marked by red circles. False colour-coded area shows the 2D probability density of 124 outward going paths of foraging ants that operated from this nest and were tracked with Differential GPS over a period of two years. See colour bar for scale and Fig. S1 for individual paths. Intended paths of tethered ants are shown for the Off-route location (top left panel), the Nest location (top right panel) and the On-route location (bottom centre panel), with the paths of zero-vector (ZV) ants shown in red and those of full-vector ants (FV) in black. The nest location is indicated by a blue circle. Insets show for both FV and ZV ants the probability density of virtual distances from the starting point reached after 5 minutes and circular histograms of final bearings with red line indicating the length and direction of the mean vector; r: resulting vector length; p: probability of rejecting hypothesis of uniform distribution and z: z-statistic of Rayleigh test of uniformity. Bottom left panel: Paths of 14 ants released just north of the Off-route location and tracked with Differential GPS. Bottom right panel: Learning walk paths of ants around the nest, recorded with Differential GPS.*

<u>*Modelling procedure*</u>

At each time step, the agent computes two values of unfamiliarity: one by comparing the current view to the attractive memory bank and one by comparing the same current view to the repellent memory bank (*Fig. 25C & D*). These two unfamiliarity values are assumed to have an antagonistic effect on

the agent's behaviour by turning it towards attractive and away from repellent stimuli with the balance between the two drives determining the agent's turning direction. We modelled this by a simple subtraction resulting in a raw overall drive

**Raw overall drive = (attractive unfamiliarity value – repellent unfamiliarity value) / 0.2** (1)

We normalised the value of this drive by using always the same value (0.2 in our world), corresponding roughly to the unfamiliarity score obtained between views from locations in the virtual world that are far apart, so that Raw overall drive will be contained between -0.5 and 0.5. A negative value thus indicates that 'attractive unfamiliarity' < 'repellent unfamiliarity'. A positive value indicates that 'attractive unfamiliarity' > 'repellent unfamiliarity'. We then transform Raw overall drive into an Overall drive with values ranging from 0 to 1 using a simple sigmoid function:

**Overall drive = Sigmoid (Raw overall drive)** (2)

As a result, the Overall drive tends towards 0 if 'attractive unfamiliarity' < 'repellent unfamiliarity', towards 1 if 'attractive unfamiliarity' > 'repellent unfamiliarity' and is 0.5 if 'attractive unfamiliarity' = 'repellent unfamiliarity'. In other words, a low score indicates that the current view matches a view in the attractive memory bank better than in the repellent memory bank and a high score indicates that the current view matches a view in the repellent memory bank better than in the attractive memory bank (*Fig. 25C*).

To drive the agent, we used a similar approach to Kodzhabashev and Mangan (2015). The agent is a simple dot in space (x,y) with a current heading (theta). The agent has a continuously running oscillator alternating between left mode and right mode, which controls the current turning direction. For simplicity, we modelled this by simply alternating the turning direction at each time step (Left-Right-Left-Right) as in Kodzhabashev and Mangan (2015). The resulting paths typically show sharp zigzags, however it is worth noting that alternating turning direction every 4th step produces smoother oscillations that better resemble real ant paths (*Fig. 25E*).

Turn direction is thus purely controlled by the oscillator, however, the Turn amplitude is directly dependent on the current Overall drive (see previous section), that is, on the current view familiarities.

**Turn amplitude (deg) = gain × Overall drive** (3)

We use a single parameter (gain) to convert the Overall drive (between 0 and 1) into the angular value for the turn amplitude. We simply used gain = 180, so that the turning amplitude would vary between 0 degrees (if Overall drive = 0) and 180 degrees (if Overall drive = 1), with 90 degrees if Overall drive = 0.5, that is if attractive and repellent unfamiliarity values are equal.

Across time steps (n), the agent orientation (theta) will thus alternate between left and right turns ((-1)n), with each turn varying between 0 and 180 degrees.

**Theta(n+1) = Theta(n) + (Turn amplitude × (-1)n) + noise**

To ensure that the agent is robust against the intrinsic noise of the real world, we added noise at each time step, as a random angular value drawn from a Gaussian distribution (mu=0; std=10 degrees).

Agent on a fictive tread-mill: We simulated agent behaviour on a fictive treadmill by simply preventing forward motion. That is, at each time step we assumed that the agent (1) obtains the current view and computes its Overall drive (Eqn 1 & 2); (2) turns on the spot with turn direction determined by the state of the oscillator and turn amplitude by Eqn 3 & 4. Since the location at which the agent is standing does not change, the view perceived at each time step only varies depending on the agent's current orientation. The agent on the tread mill was tested at different release locations and we recorded the resulting behaviour.

Using attractive visual memories only: We also tested the agent using the attractive memory bank only. In that case

**Raw overall drive = attractive unfamiliarity/0.2 - 0.5**.

Given that attractive unfamiliarity is always positive, we removed 0.5 during normalisation to centre the Raw overall drive on 0, ranging roughly from -0.5 to 0.5 in the same way as when combining attractive and repellent memories. We

then used the same sigmoid function to obtain an Overall drive between 0 and 1 (Eqn 2).

## Results

### Myrmecia *ants released on the tread-mill*

Irrespective of whether they were caught in a zero-vector state (ZV) or a full-vector state (FV), tethered ants behaved differently when placed 6.5 m west of the nest (off-route, Fig. 2, top left panel), 6.5 m south of the nest (on-route, *Fig. 26*, bottom middle panel) or over the nest (Nest, *Fig. 26*, top right panel).

In the off-route and on-route locations, most intended paths of both ZV and FV ants were goal directed either to the nest or to the individuals' specific foraging trees (see inset circular histograms in *Fig. 26*). This is to be expected for M. croslandi foragers, which ignore path integration information in the FV state as long as the landmark panorama provides navigational information, see (Narendra et al., 2013b; Zeil et al., 2014b). The paths tended to be straight (see *Fig. 27D*). In contrast, over the nest, ZV ants moved in random directions, while FV ants tended to move roughly along the home vector direction to the north (at 90° *Fig 26*, black tracks, top right panel, see also *Fig. 27B*). Both ZV and FV ants at the nest changed their walking direction frequently. Inset histograms show that most tethered ants over the nest ended up after 5 minutes at final virtual distances less than 10 m from the nest (median 6.07 m), while at the on-route location, most ants reached much larger virtual distances (median 12.61 m) in the same amount of time (Wilcoxson Rank Sum test: nest vs on-route distances are different: $p = 0.0045$; $z = 2.8378$. See *Fig. 32A*). The median distances reached at the off-route location are not larger than the ones at the nest (median 6.2 m), owing to a conspicuous peak at small distances contributed by ants that were lost at this location.

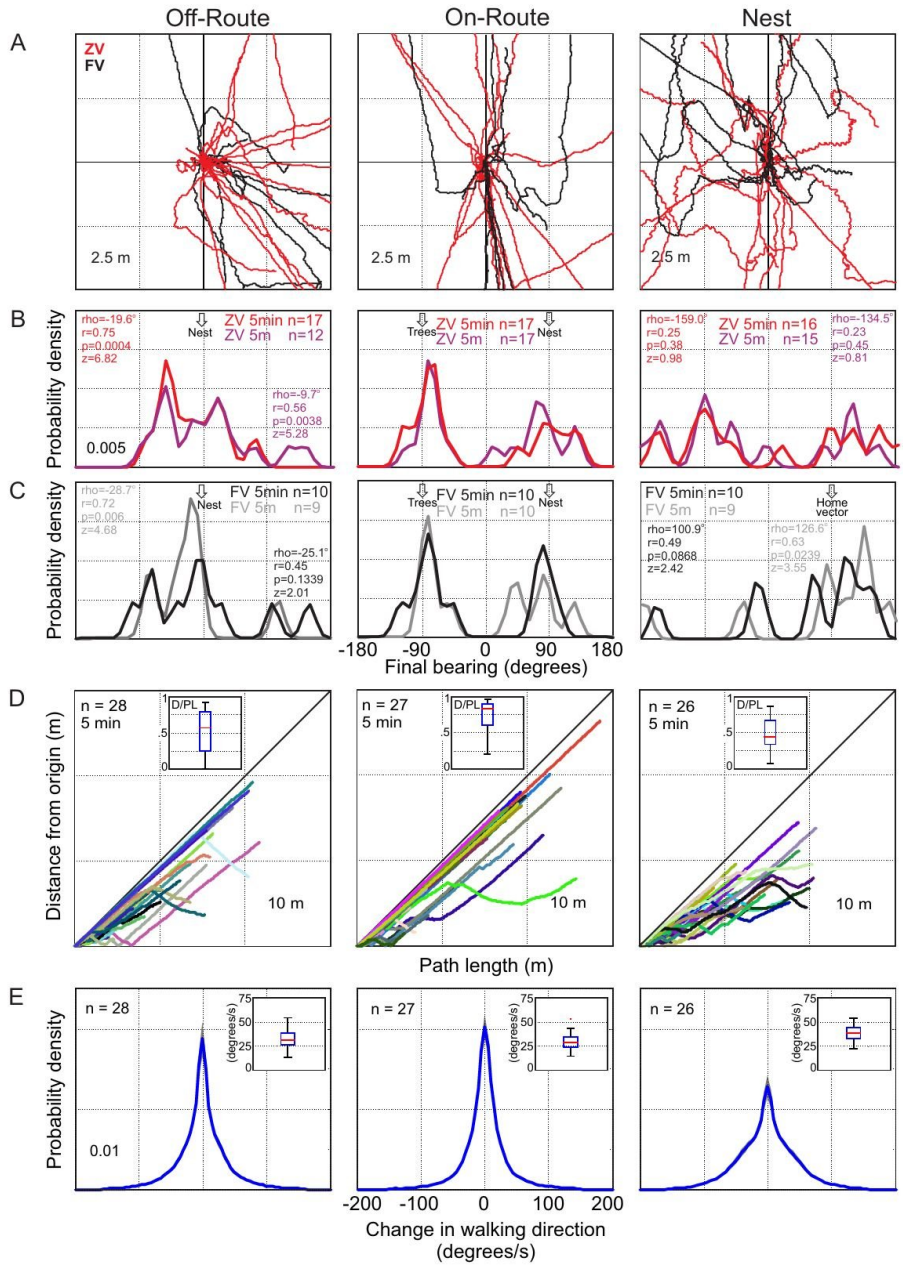


*Fig. 27: Quantitative analysis of behavioural differences at the Off-route, On-route and Nest location*

*(A) Initial intended paths of tethered ants at a finer scale. The paths of ZV ants are shown in red and those of FV ants in black. (B) Distributions of final bearings of ZV ants after 5 minutes (red) or when having reached a virtual distance of 5 m from the start (purple) at the three locations. Probability densities determined with 9o bandwidth of the kernel smoothing window; North at +90o. Inset numbers show results of circular statistics (Rayleigh test of uniformity) with rho: mean vector direction; r: mean vector length; p: probability of uniformity; z: z-statistic. Arrows mark the direction of nest and trees. (C) Distributions of final bearings for FV ants after 5 minutes (black) or when having reached a virtual distance of 5 m from the start (grey). Arrows mark the direction of nest, trees and home vector. Otherwise conventions as in (B). (D) Distance from start over path length for the first 5 minutes of paths at the three locations. Paths are randomly coloured. Insets show boxplots with median marked red for the ratios of distance over path length at the end of 5 minutes. See Fig. S2B for statistics. (E) Distributions of changes in walking direction for all 5 min paths at the three locations. Shown are the means of individual distributions (blue) and standard errors in grey (not visible). Insets show boxplots for the distributions of individual means. See Fig. S2C and D for statistics.*

The behaviour of ants at the off-route location is interesting primarily because most ants are home directed despite it being unlikely that they have ever been to this location before (see inset circular histogram, *Fig. 26*, top row left panel). A heat map of the foraging movements of 124 ants from this nest that had been DGPS-tracked on their outward foraging trips over two years shows that no ant had moved off-route of the nest for more than a few meters (*Fig. 26*, top row middle panel). Some of the tethered ants appear to have headed towards their foraging trees or the foraging corridor in south-easterly direction, however, when we track ants that were released just north of the off-route location, many initially for 2 m or so do walk in a south-easterly direction before turning east toward the nest (*Fig. 26*, bottom row, left panel). Tethered ants at the off-route location must therefore get their bearing by comparing what they currently see with nest-directed views they are likely to have gathered during their learning walks, which can extend up to 4 m from the nest (*Fig. 26*, bottom right panel), see also (Jayatilaka et al., 2018).

Both FV and ZV ants at the on-route location decided to move either back toward the nest or south toward their foraging trees (*Fig. 26*, bottom row, centre panel). Otherwise, they moved in a similar way than when at the off-route location. Most of them moved fast, straight and for distances far exceeding those needed to reach the nest or the trees.

The most conspicuous feature of paths at the nest location is the fact that the initial walking direction of ZV ants is random, while those of FV ants is in the general home vector direction (north) and that both ZV and FV ants change walking direction frequently.

We quantify these differences between locations in three ways in *Fig. 27*, considering final bearings, the relationship between path length and distance reached and changes in walking direction. *Fig. 27A* shows the initial paths of ants at the three locations in more detail to emphasize the different behaviours and to highlight the additional fact that paths are fairly smooth at the off-route and on-route locations, but show a distinct sinusoidal oscillation at the nest location. With the exception of the bearings of ZV ants at the nest (*Fig. 27C* right panel) and those of FV ants after 5 min at the off-route and the nest

location (*Fig. 27C* left and right panel) the virtual bearings of ants after 5 minutes or at 5 m distance from the start are all significantly different from uniform distributions, both for ZV (*Fig. 27B*) and FV ants (*Fig. 27C*). While the distributions are unimodal for the off-route and nest location (see insets *Fig. 27B* and C for circular statistics), they are clearly bimodal at the on-route location.

One measure of the straightness of paths is the way in which the straight-line distance from the start depends on path length (*Fig. 27D*), with straight paths without changes in direction lying close to the line of equality. After 5 minutes, the distribution of the ratios of final distance to final path length differs between the sites (see insets in *Fig. 27D* and *Fig. 32B*) with the on-route paths being significantly straighter with a median ratio of distance over path length of 0.83, compared with 0.62 at the off-route location and of 0.45 at the nest location (Wilcoxon Rank Sum test at 5% significance level: On-route vs off-route: p=0.0110; on-route vs nest: p=8.4992e-4; off-route vs nest: p = 0.6. See Fig. S2B).

Finally, the behaviour of ants at the three sites also differs on a finer scale: the distribution of changes in path direction is much broader at the nest site, compared to the off-route and on-route location (*Fig. 27E*) reflecting the conspicuous oscillations of ant paths over the nest (see right panel, *Fig. 27A*). Note that these distributions have very long tails due to spikes of very high angular velocities which may be artefacts of trackball rotations when the ants are moving very slowly (see time series in *Fig. 28*). To test whether changes in path direction are indeed systematically larger at the nest location, we calculated the means of their absolute values at 11fps over the first 5 min of walking for each ant and compared their distributions, both for angular velocities smaller than 200o/s (insets *Fig. 27E* and *Fig. 32C*) and for all angular velocities (*Fig. 32D*). Below 200o/s, nest paths are indeed wigglier compared to on-route paths (*Fig. 32C*, Wilcoxon Rank Sum test: nest vs on-route p=2.56e-4, z=-3.6563), with the difference between nest and off-route location just failing to reach significance (nest vs off-route p=0.019, z=-2.3458). Considering the whole range of angular velocities (*Fig. 32D*) there is no difference between nest

and the other locations, mainly because of high angular velocities exhibited by ants at all sites.

We note that many ants at various times during the first 5 minutes on the trackball over the nest show very regular path oscillations as documented in Fig. 3A and for three examples in *Fig. 28* (red traces). The distribution and the time course of changes in path direction over the nest are different from those exhibited by the same ants at the on-route location (shown in blue in *Fig. 28*). Regular and sustained path oscillations lead to periodicities in the auto-correlation function of changes in path direction and can be detected in 13 out of 25 cases of ants participating in all three locations (blue traces in *Fig. 33*), compared to 4/25 at the off-route location (red traces in *Fig. 33*) and 1/25 at the on-route location (green traces in *Fig. 33*). We add the caveat that the statistics of path properties are unlikely to be stationary during an experiment and that this particular aspect of ant behaviour will require future attention.

***Agent-based modelling***

To model the agent on a fictive tread-mill, we simply prevented it from stepping forward, so that views were always perceived from the same spot, and where rotated according to the agent's current facing direction. We released the agent at four locations.

When tested close to the beginning of the homing route (on-route RP), the agent oriented mostly in the correct direction, that is, along the route towards the nest (blue paths in *Fig. 29A*). This is because the overall drive is close to 0 while facing in this direction (the attractive unfamiliarity is very low and the repellent unfamiliarity is high (*Fig. 25C*) yielding very small turns (*Fig. 29B & C*). Note that if the agent happened to face in the opposite direction (due to noise), the overall drive would strongly increase and thus trigger a large turn.

When released away from the route (off-route RP), the agent also favoured one direction indicating that this direction provided a smaller overall drive (yellow paths in *Fig. 29A*). This is an indication that the view at the off-route RP and nest-directed learning walk views are most familiar because their comparison

produces a detectable minimum of the rotIDF and that the agent thus favours a direction roughly pointing towards the nest.

When released on top of the nest (nest RP), the agent produced convoluted paths with no preferred directions (red paths in *Fig. 29A*). This is due to the rather uniform distribution of visual familiarities across directions (see *Fig. 25C*). At a more local scale, the paths show much larger turn amplitudes than at the on-route or off-route RPs (*Fig. 25B & C*). This is because at the nest location, attractive and repellent memorised views provide a roughly equal match whatever the current facing direction, resulting in an overall drive around 0.5, thus yielding turns that are larger than when attractive and repellent memories match best for different directions (see *Fig. 25C*).

When released at a distant unfamiliar location (distant RP), the agent displayed equally large turn amplitudes as at the nest (marked in black in *Fig. 29A-C*) because, as for the nest location, both the attractive and the repellent memory bank provide roughly equal unfamiliarity values, thus resulting in an average overall drive around 0.5.

In contrast, when using the attractive memory bank only, turn amplitudes were large at the distant unfamiliar location (black) but comparatively low at the nest (red, right column, *Fig. 29*). This is simply because the unfamiliarity value is high in the unfamiliar location (yielding a strong directional drive and thus large turns), and low at the nest due to the good match with learning walks views (yielding a low directional drive and thus small turns).

***Testing model-predictions with* Myrmecia**

Motivated by the different simulation results when using ‘attractive only’ and ‘attractive/repellent’ memory banks as well as by the rather counter-intuitive outcome that the use of ‘attractive/repellent’ memories predicts a similar behaviour at the familiar nest location and at a completely unfamiliar location, we released *Myrmecia* ants mounted on the trackball both at the nest and at a distant location about 50 m south-west of the nest. The location was far beyond the ants’ foraging trees and thus was likely to be completely unfamiliar to the ants. Strikingly, ants at this distant release location behaved in a similar way as

at the nest, both in terms of the ratio between the distance reached after 5 minutes and the path length (see box plot insets in *Fig. 30A* centre panels, Wilcoxon Rank Sum test unfamiliar vs nest location: p=0.7984, ranksum=71) and in terms of the mean absolute changes in walking direction (see box plot insets in *Fig. 30A* right panels, Wilcoxon Rank Sum test unfamiliar vs nest location: p=0.9591, ranksum=67). The ants at both the unfamiliar and the nest site also displayed the characteristic path oscillations we observed at the nest in our previous experiments (*Fig. 30A & B*, compare with *Fig. 28*), as predicted by the attractive/repellent model.

## Discussion

Our behavioural experiments revealed three fundamental properties of visual navigation in ants that could only be uncovered using the trackball method. First, we determine that whether on-route or off-route, several metres away from the nest, ants can recover the goal direction without the need to physically move and to sample neighbouring locations. Second, we find no evidence that they 'expect' outcomes from their behaviour, such as a changing visual scene or increasing certainty about the location of the nest. M. croslandi ants show no evidence of monitoring the distance that separates them from the goal, unlike for instance ants that rely strongly on path integration (Dahmen et al., 2017). Third, ants behave differently when positioned above the nest, by following random heading directions and frequently changing their walking direction. These are the characteristics of search behaviour and thus could be interpreted as indicating that ants 'know' that they are at the nest, as if they possessed location information. However, our simulation results demonstrate that the nest-specific behaviour of ants can be parsimoniously explained by the density of attractive, nest-directed, and repellent views away from the nest that at least *M. croslandi* ants are likely to acquire in the course of systematic scanning movements during their learning walks, e.g. (Jayatilaka et al., 2018). Our simulation also confirms that the same parsimonious mechanism can recover a correct direction from on- and off-route locations, as previous modelling has indicated (Baddeley et al., 2011; Kodzhabashev and Mangan, 2015; Narendra et al., 2013b; Wystrach et al., 2013).

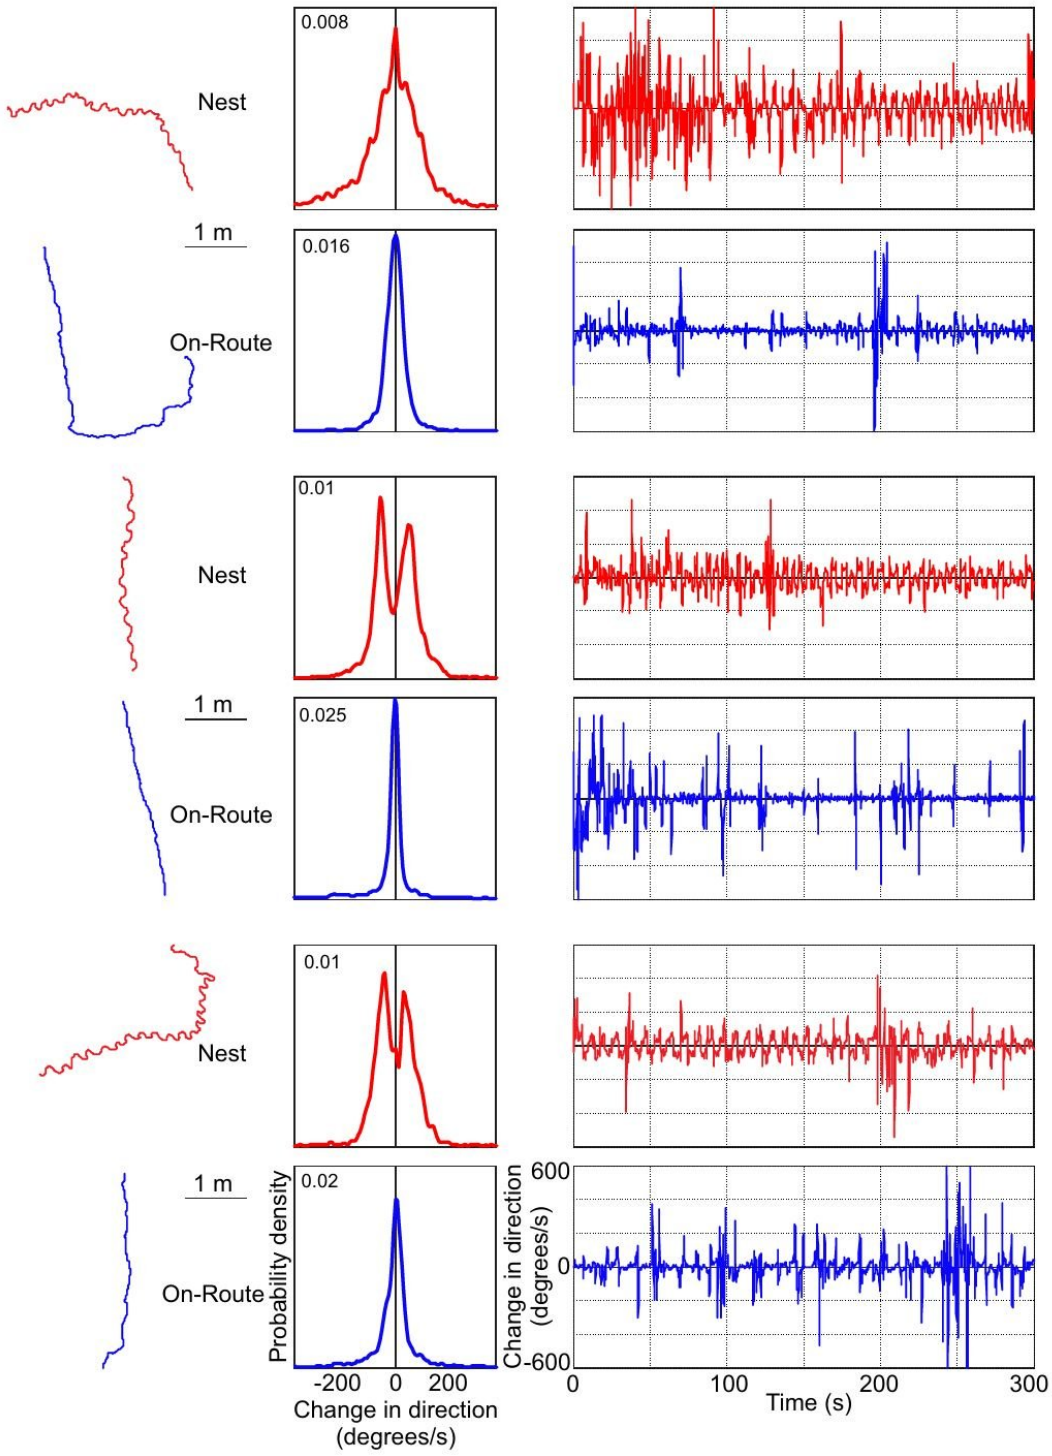


*Fig. 28: Ants walk differently at the Nest (red) and the On-route (blue) location*

*Shown are path segments on the left, the distributions of changes in walking direction during the first 5 minutes in the middle row and the time series of changes in walking direction over 5 minutes on the right for three ants (top, centre, bottom), each recorded at the nest and at the on-route location. Changes in walking direction were determined at 11fps to reduce measurement noise. See Fig. 33 for auto-correlation functions.*

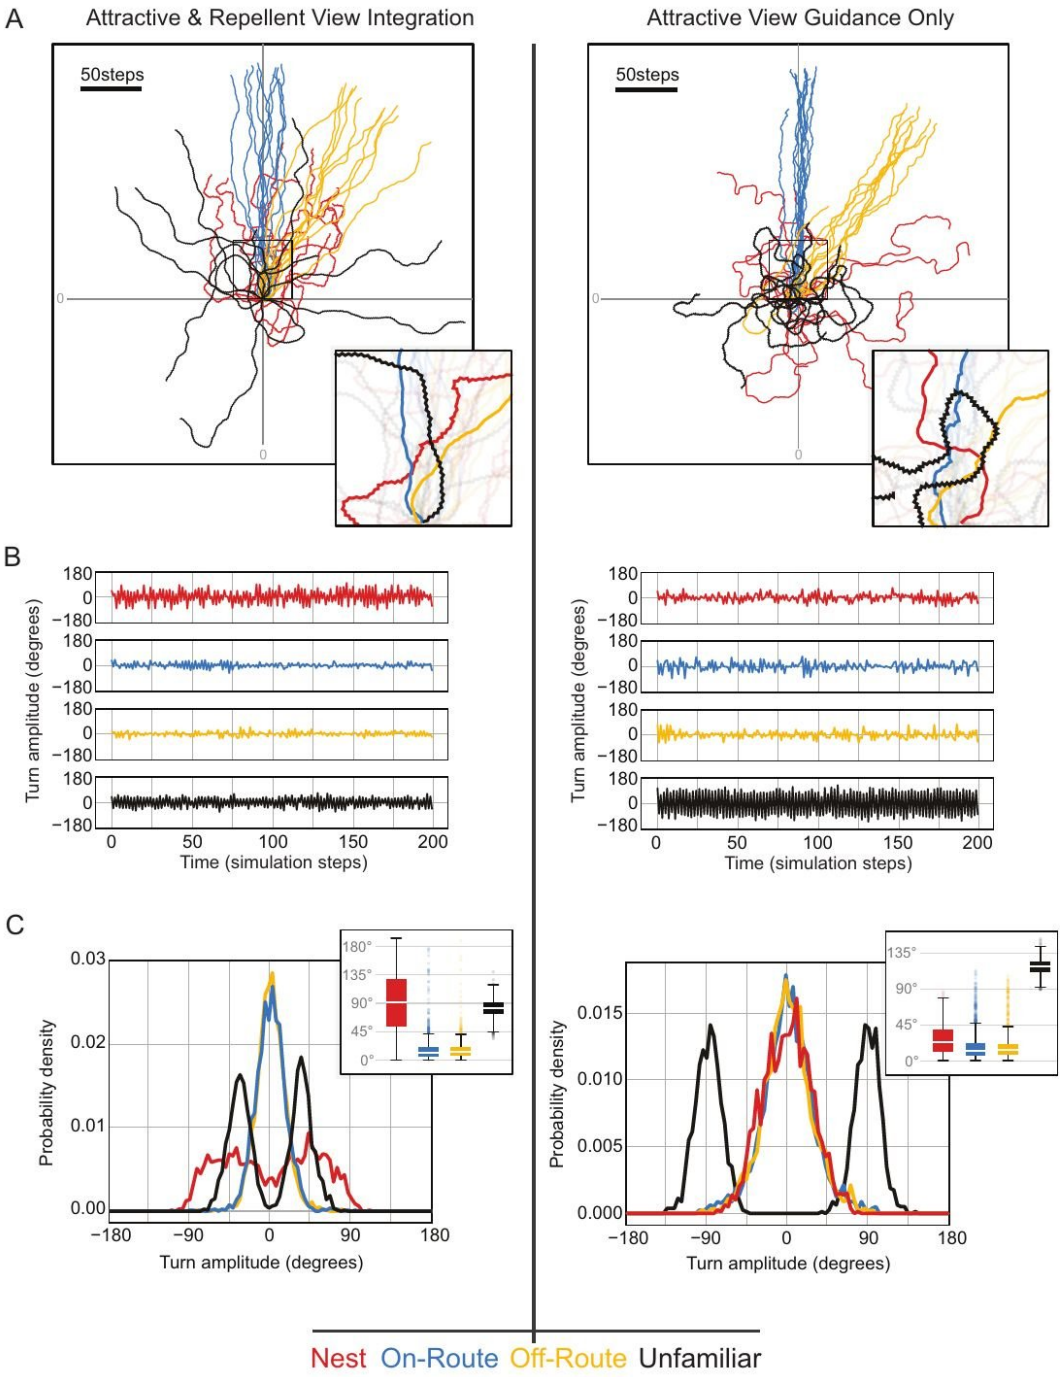


*Fig. 29: The results of agent-based visual navigation using both attractive and repellent views (left column) and attractive views only (right column).*

*We simulated ten agents walking 200 steps at each nest (red), on-route (blue), off-route (yellow) and unfamiliar (black) release locations. (A) Resulting paths. Insets show close-up details of example paths. (B) Turn amplitudes over time (simulation steps) for one example at each of the release locations. (C) Probability densities of turn amplitudes at the four release locations. Inset show box and whisker plots for the same distributions.*

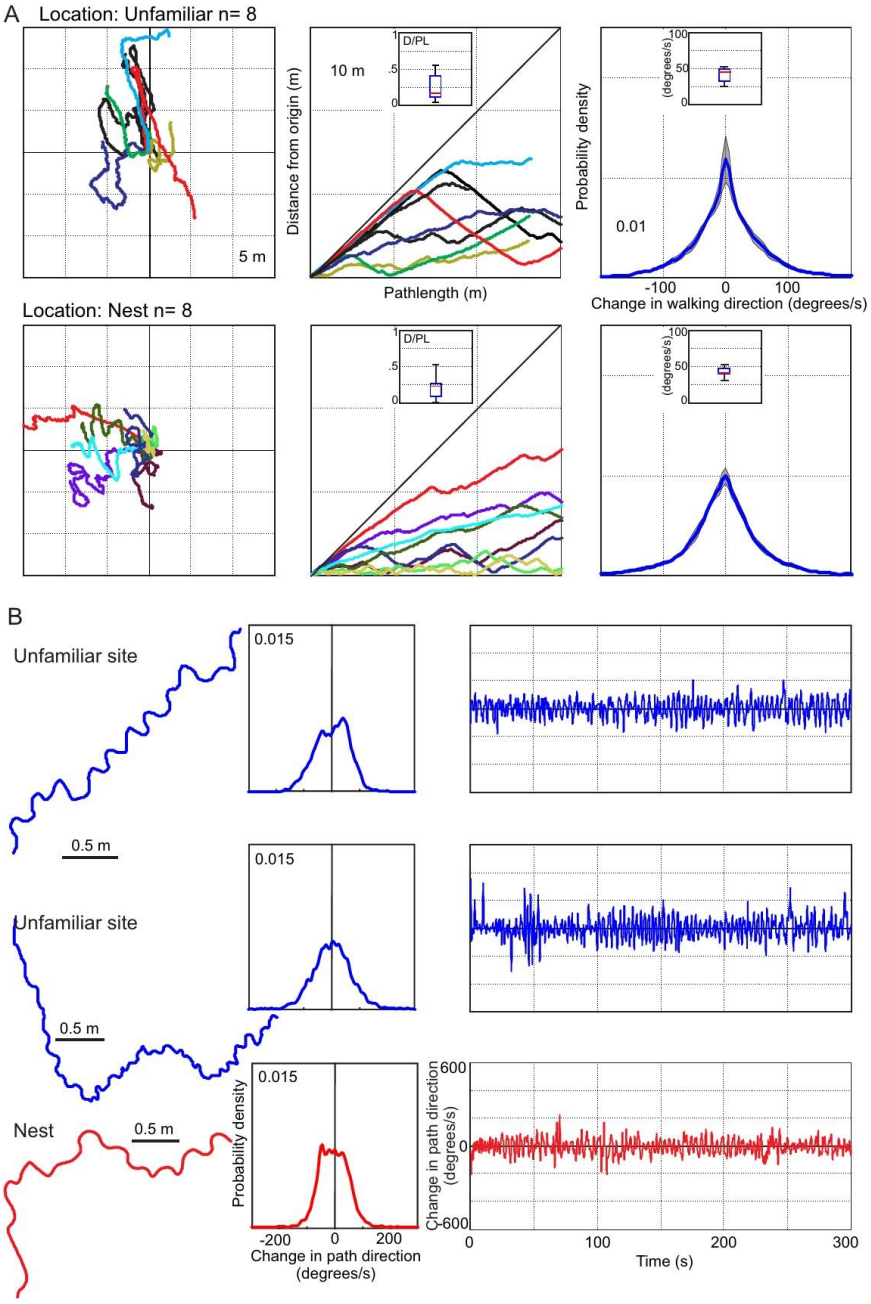


*Fig. 30: Ants behave in a similar way at a completely unfamiliar location and at the nest*

*(A) Top row: Paths (left), distance from start over path length (middle) and probability density of changes in walking direction (right) for 8 tethered ants at a completely unfamiliar location. Bottom row: Same for 8 ants at the nest location. Insets in middle panels show boxplots of final distance to path length ratios after 5 minutes, which are not different between the unfamiliar and the nest location (Wilcoxon Rank Sum test unfamiliar vs nest location: p=0.7984, ranksum=71). Insets in right panels show the boxplots of mean absolute values of changes in path direction over 5 minutes, which are not different between the unfamiliar and the nest location (Wilcoxon Rank Sum test unfamiliar vs nest location: p=0.9591, ranksum=67). (B) Example paths (left), probability density of changes in path direction (middle) and time series of changes in path direction (right) for two ants at the unfamiliar site (blue) and one ant over the nest (red).*

***Alignment matching and visual memories***

Current thinking holds that ants during their learning walks learn nest-directed views; *Cataglyphis sp*: (Fleischmann et al., 2016, 2017, 2018a, 2018b), *Ocymyrmex robustior*: (Müller and Wehner, 2010), *Melophorus bagoti*: (Wehner et al., 2004; Muser et al., 2005); and possibly both nest-directed views and views pointing away from the nest; *Myrmecia croslandi*: (Jayatilaka et al., 2018; Zeil and Fleischmann, 2019). In addition, they memorize the views they experience along routes as they go back and forth on foraging excursions (Wehner et al., 1996; Mangan and Webb, 2012; Kohler and Wehner, 2005; Wystrach et al., 2011b; Freas and Spetch, 2019).

When using their visual memories to navigate, the currently perceived panorama provides a heading direction if the comparison between memorised views and the current view generates a detectable minimum of the rotational image difference function, see (Narendra et al., 2013b). This is a basic measure of familiarity and at any location, the direction presenting the most familiar view would provide the deepest (lowest) minimum. At both the on-route and off-route location ants on the trackball were free to scan the panorama and detect the direction of any present minima. Our results show that they were successfully able to recover the goal direction by doing so (*Fig. 26*). On route, some ants headed to the nest while others aimed at their foraging trees, reflecting their motivation to home or to forage.

While the directedness of ants at the on-route site would have been supported by both learning walk views and views learnt along the route, their directedness at the off-route (west) location depends on their detecting a higher similarity with learning walk views directed at the nest from the west compared to all other nest-directed views. As shown here and before (Narendra et al., 2013b; Zeil et al., 2014b; Stürzl et al., 2015), this is possible up to 10-15m distance from the nest in the open woodland habitats of *Myrmecia* ants, provided ants have acquired such nest-directed views about 1 to 5 metres away from the nest (see *Fig. 26B*, bottom right).

When released at the nest, ants behaved differently. They walked in various directions and displayed larger turns that regularly alternated between left and right, resulting in sinusoidal paths. So are nest views special?

As far as navigational information is concerned, the situation at the nest is indeed different compared with both on- and off-route sites. During their learning walks ants will have encountered a dense set of views at different distances and compass bearings around the nest, each potentially tagged with the nest direction through path integration (Müller and Wehner, 2010; Graham et al., 2010; Baddeley et al., 2012; Fleischmann et al., 2018a; Jayatilaka et al., 2018; Zeil and Fleischmann, 2019). In contrast to other locations, tethered ants placed above the nest location thus will encounter attractive familiar views (or deep rotIDF minima) in many compass directions, which might explain why they initially walked in various directions at this location.

The high amplitude oscillation displayed by ants at the nest location, however, is puzzling. Previous models suggest that experiencing a familiar (attractive) view should inhibit turns and favour forward motion (Zeil, 2012; Möller, 2012; Baddeley et al., 2011, 2012; Wystrach et al., 2013; Kodzhabashev and Mangan, 2015; Ardin et al., 2016), which is here clearly not the case. The behaviour of tethered ants on top of the nest can be interpreted as search for the nest entrance, which in ants relying on path integration is characterized by frequent changes in path direction and a systematic pattern of increasing loops around the expected location of the goal, e.g. (Schultheiss et al., 2015). To our knowledge, however, no analysis of the fine-scale changes in orientation of searching ants – as we observed them here - has been done to date.

Previous work has suggested that the recognition of views memorised at the nest may trigger specific behaviours when subsequently released in unfamiliar locations (Wystrach et al., 2013). This interpretation may suggest positional knowledge, or at least that views close to the nest are categorised separately from route views during learning and being treated differently when recognised. In the following we discuss the results of our simulation that suggest a parsimonious and unifying explanation for view-based route guidance, pinpointing goals and the current observation of high amplitude oscillation at the

nest without the need to invoke positional knowledge or the need for a 'trigger' of search behaviour. Our agent-based modelling exhibits the same pattern of fine-scale oscillations, including overall changes in path direction, but only if we assume that the agent operates with both attractive and repellent memory banks.

***Continuously integrating attractive and repellent views***

Our model was developed quite independently to explain other recently observed phenomena, such as how ants manage to use views for guidance while walking backward and thus facing in the anti-nest direction (Schwarz et al., under review); or how ants learn to detour areas along their route associated with an aversive experience (Wystrach et al., 2020). Interestingly, this new model happens to also capture the current results remarkably well. The model is based on two assumptions: (1) that ants store both attractive and repellent views during their learning walks as suggested by Jayatilaka *et al* (2018), and (2) that guidance involves an oscillator resulting in a continuous alternation between left and right turns (Namiki and Kanzaki, 2016; Kodzhabashev and Mangan, 2015; Wystrach et al., 2016b). The model assumes no positional knowledge whatsoever, only procedural knowledge.

Several pieces of evidence suggest that insects possess an intrinsic oscillator triggering alternatively left and right body rotations, the amplitude of which can be modulated by the stimuli perceived (Namiki and Kanzaki, 2016; Lent et al., 2013; Wystrach et al., 2016b). Such a control of oscillations can provide guidance along odour plumes (Namiki and Kanzaki, 2016) and odour gradients (Wystrach et al., 2016b), support visual route following (Kodzhabashev and Mangan, 2015) and greatly facilitates the integration of different sources of stimulation (Wystrach et al., 2016b). In the case of visual route following, the amplitude of the oscillations needs to be simply modulated by the familiarity of the currently perceived view. The suggestion is that familiar views trigger small turns whereas unfamiliar views trigger large turns, and that the direction of the turn is dependent on the current state of the oscillator. Because views are assumed to be memorized while moving along the route, during route recapitulation visual familiarity is higher when facing in the correct route

direction. This model is sufficient for an agent to recapitulate a route in naturalistic environments (Kodzhabashev and Mangan, 2015). However, when released at the nest, this model does not predict large amplitude oscillations such as the ones we observed here in ants. On the contrary, because of the high familiarity experienced at the nest, which results from the collection of nest-oriented views acquired during learning walks, the model predicts an inhibition of the oscillations whatever the current facing direction (see *Fig 29*, right column).

The visual memories used by insect navigators are likely stored in the mushroom bodies (Webb and Wystrach, 2016), but current models assume only the existence of attractive memories (Möller, 2012; Baddeley et al., 2011, 2012; Wystrach et al., 2013; Kodzhabashev and Mangan, 2015; Ardin et al., 2016). Here we incorporated into the model the recent suggestion that ants store both attractive and repellent views, mimicking the so-called 'appetitive/aversive' output pathways from the insect mushroom bodies, e.g. (Owald et al., 2015; Saumweber et al., 2018) (*Fig. 25D*). Indeed, during their learning walks, many ants, not only *Myrmecia croslandi* (Jayatilaka et al., 2018) display regular head and body oscillations, facing alternatively towards and away from the nest (Zeil and Fleischmann, 2019). We assumed in our model that these views form two distinct memory banks: one holding ‘attractive’, nest-directed, views and one holding ‘repellent’ views pointing away from the nest, and that both sets are used continuously and simultaneously during homing. Our agent compares the current view to both sets of memories at each time step and thus obtains two familiarity values, one for attraction (high familiarity, inhibiting turns) and one for repulsion (high familiarity, triggering large turning amplitudes). Given that both memory pathways have antagonist outcomes, they can be simply integrated by subtracting attractive and repellent familiarity values, resulting in what we called here an ‘overall drive’ which modulates the amplitude of the oscillator (*Fig. 25C*).

Interestingly, this model closely mimics ant behaviour as documented in our behavioural experiments. If released on a fictive tread-mill (preventing the agent from translating) it displays high amplitude turns when released on top of the

nest, and much lower amplitude turns when released further along the homing route. In contrast, when using the 'attractive' memory bank only, the agent produces low amplitude turns at the nest (*Fig 29*).

The behaviour of the agent when combining attractive and repellent views is straightforward to explain (*Fig. 25C*). At the route release point, facing in the correct direction the simulation generates very small turns because only the attractive memory bank provides a good match. By integrating this with a high unfamiliarity of the repellent memory bank, we obtain a very low overall drive, and thus small turns. However, when released at the nest, whatever the direction the agent faces, there are always both attractive and repellent views that are matching the current view (*Fig. 25C*). The reason being that these views, when acquired during learning walks, are experienced in multiple compass direction at very closely spaced locations (*Fig. 25B*). Both attractive and repellent pathways signal high familiarity values and cancel each other out, resulting in large turns.

***Testing the model's prediction.***

Interestingly, the attractive/repellent memory bank model makes a rather counter-intuitive prediction, because it relies on the relative difference in familiarities between attractive and repellent pathways and not on the absolute familiarity experienced: the agent's behaviour should be similar when on top of the nest and at a completely unfamiliar location, outside the catchment area of acquired views. At the nest, both attractive and repellent memories result in high familiarity, so their signals cancel each other when integrated (attractive - repellent), resulting in large turns. In completely unfamiliar terrain, both attractive and repellent memories result in very low familiarity, and thus their signals equally cancel each other when integrated (attractive - repellent), resulting also in large turns (*Fig. 25C*).

As predicted by the model, experiments showed indeed that ants tethered at a completely unfamiliar location exhibit a very similar behaviour to when released on top of the nest: that is, they displayed regular high amplitude path oscillations (*Fig. 30*).

***Integration with path integration.***

We did not incorporate integration of path integration information and landmark panorama guidance in our model and so do not at this stage tackle the fact that full vector ants (i.e., those captured with a remaining path integration home vector) showed a small bias towards the home vector direction at the nest location (*Fig. 26 & 27*, FV vs ZV ants). In *M. croslandi* foragers, as in other ants, path integration information and scene information are integrated (Collett et al., 2001; Collett, 2012; Reid et al., 2013; Legge et al., 2014; Narendra et al., 2013b; Wystrach et al., 2015; Wehner et al., 2016) with familiar views more strongly weighted – to the degree that a current view providing information on heading direction can completely override conflicting information from path integration (Kohler and Wehner, 2005; Narendra et al., 2013b; Zeil et al., 2014b). In ants that rely heavily on path integration, this information is more strongly weighted as the length of the vector increases (Wystrach et al., 2015, 2019). The bias towards the home vector direction observed here in FV ants fits this current view, which is summarised in a recent model (Hoinville and Wehner, 2018). Also, experienced ants seem to rely less on path integration than naïve ants, and rather display a search when on unfamiliar terrain (Schwarz et al., 2017b), which may explain why path integration information is never strongly weighted in the long-lived *M. croslandi*.

## Outlook

Our results may contribute to the lingering debate about the format of spatial knowledge underlying visual navigation in insects and animals in general, see for example (Cheeseman et al., 2014a, 2014b; Cheung et al., 2014; Warren, 2019). We showed that ants released on top of the nest displayed large turns. These results were clearly at odds with the current ‘procedural’ models, stipulating that the high familiarity of views at the nest should inhibit turns. In contrast, the ants’ behaviour suggested that they could derive positional knowledge from the current views, given the interpretation that the ants searched because they recognised that they were at the nest. Previous results, such as the apparent ability of insects to make shortcuts also favoured explanations assuming ‘positional’ rather than ‘procedural’ knowledge, e.g.

(Cheeseman et al., 2014a, 2014b; Warren, 2019). However, as often in the insect literature (Cartwright and Collett, 1983; Collett et al., 2007; Cruse and Wehner, 2011; Wystrach and Graham, 2012; Narendra et al., 2013b; Cheung et al., 2014), an alternative, more parsimonious explanation can also explain our results: ants may simply combine attractive and repellent memories. Importantly, this procedural explanation did not come from actively seeking for it, but emerged from other observations, such as the way in which ants behave when learning views around the nest (Jayatilaka et al., 2018), avoid adverse situations (Wystrach et al., 2020), steer while walking backwards (Schwarz et al., 2017a, 2020)as well as how appetitive and aversive memory pathways are combined in other insects such as flies (Felsenberg et al., 2018) and fly larvae (Eichler et al., 2017).

Our simulation made the unexpected prediction that behaviour in completely unfamiliar terrain should be the same as at the very familiar nest, which we confirmed by subsequent experimentation. Purely scene familiarity-based modelling replicates these results with astonishing detail, providing support for the suggestion that ants during their learning walks, acquire both attractive, nest-directed views and repellent views when pointing away from the nest during systematic scanning movements (Jayatilaka et al., 2018; Zeil and Fleischmann, 2019). It is not clear at present, however, whether all views are memorized irrespective of gaze direction or only when the ants' head is aligned parallel to the home vector, see discussion in (Jayatilaka et al., 2018). We show here, at least, that the distinctly different behaviour of ants over the nest location can be replicated if an agent has an attractive and a repellent scene memory bank.

The most parsimonious explanation for our observations is therefore that the ants operate on 'procedural' rather than 'location' information (Collett et al., 2002; Wehner et al., 2006; Graham and Philippides, 2017): at both familiar and unfamiliar locations away from the nest they may know where to go, but they do not know where they are. Moreover, the main assumptions of our simulation - attractive and repellent view comparison driving an oscillator - can be tested by a detailed comparison of the gaze and path directions of individually identified

ants during their learning walks and during their subsequent approach to the nest, when returning from foraging excursions. Such an analysis may also reveal how ants eventually pinpoint the nest entrance, which none of the current homing models can properly explain.

## Acknowledgements

We thank Chloé Raderschall and Piyankarie Jayatilaka for tracking ants over the years, Camile Moray, Fiorella Ramirez-Esquivel and Moosarreza Zahedi for help with field work and the mechanical workshop of the Research School of Biology for constructing the housing for the trackball contraption. We are grateful to Marijke Welvaert and Teresa Neeman from the ANU Statistical Consulting Unit for their advice and to two referees for their constructive comments.

## Supplementary figures

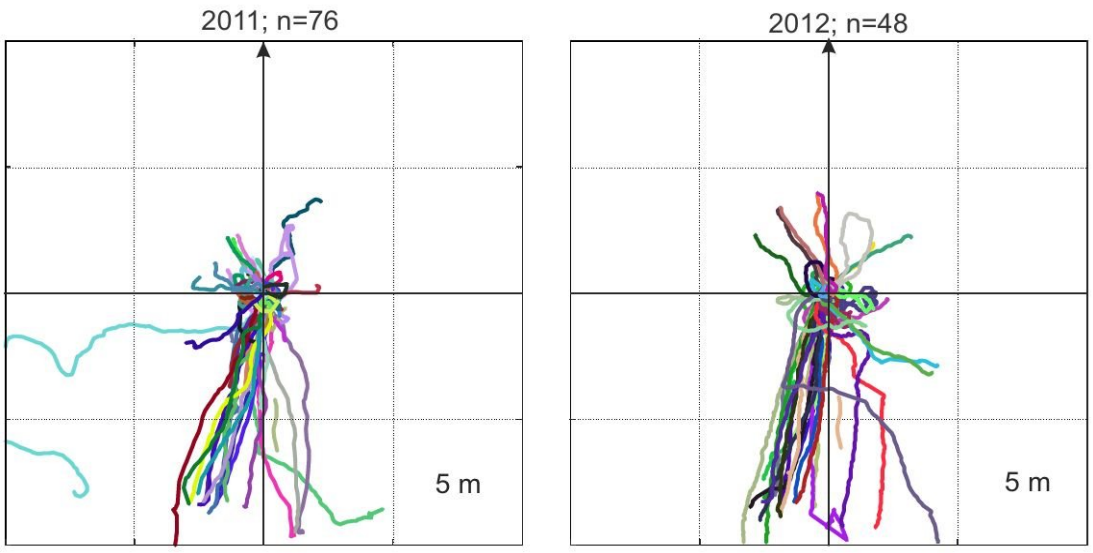


*Fig. 31: Foraging patterns at the M. croslandi nest used in this study*

*Panels show the individual paths of foragers as they have been recorded with differential GPS over a period of two years. These paths provided the original data for the 2D histogram shown in the middle panel of Fig. 25B.*

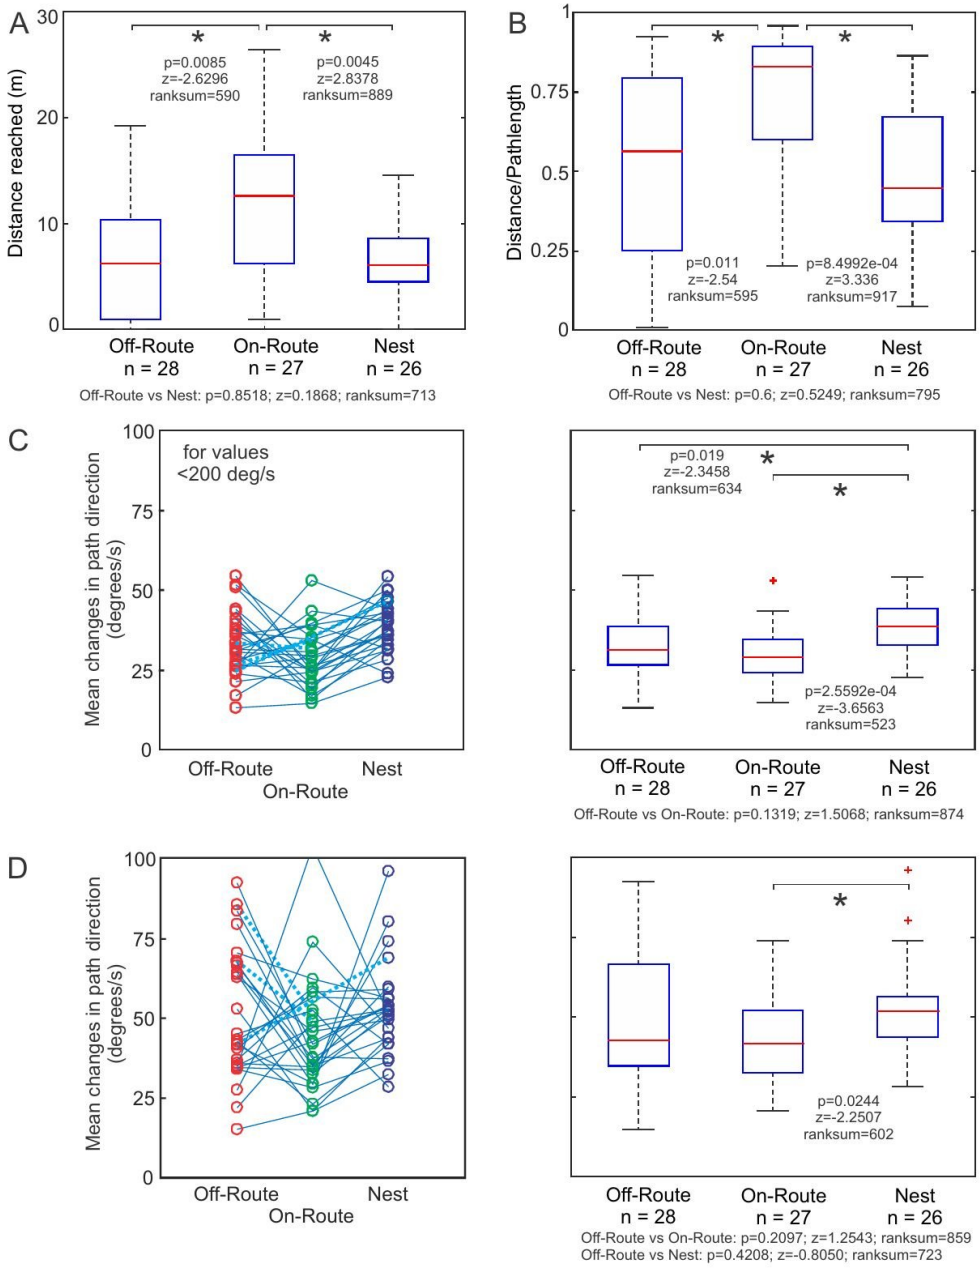


*Fig. 32: Quantitative analysis of behavioural differences between Off-route, On-route and Nest locations*

*(A) Box plots of distances reached after 5 minutes at the three locations. Significant comparisons with a Wilcoxon Rank Sum test (p<0.0167) are marked by a star and values are shown inside the panel. Values for non-significant comparisons (p>0.0167) are shown below the panel. (B) Box plots of distance over path length ratios after 5 minutes at the three locations. Otherwise conventions as in (A). (C) Left panel: Individual means of the changes in path direction (absolute values < 200°/s, determined at 11fps) for the first 5 minutes with means of individual ants connected by blue lines. Dashed lines mark cases where an ant was released at two locations only. Right panel: Boxplots of mean changes in path direction with significant differences as determined by Wilcoxon Rank Sum test marked by a star and values shown as inset. Values for non-significant comparisons are shown below the panel. (D) Same for the means and distributions of all absolute values of changes in path direction. Otherwise conventions as before.*

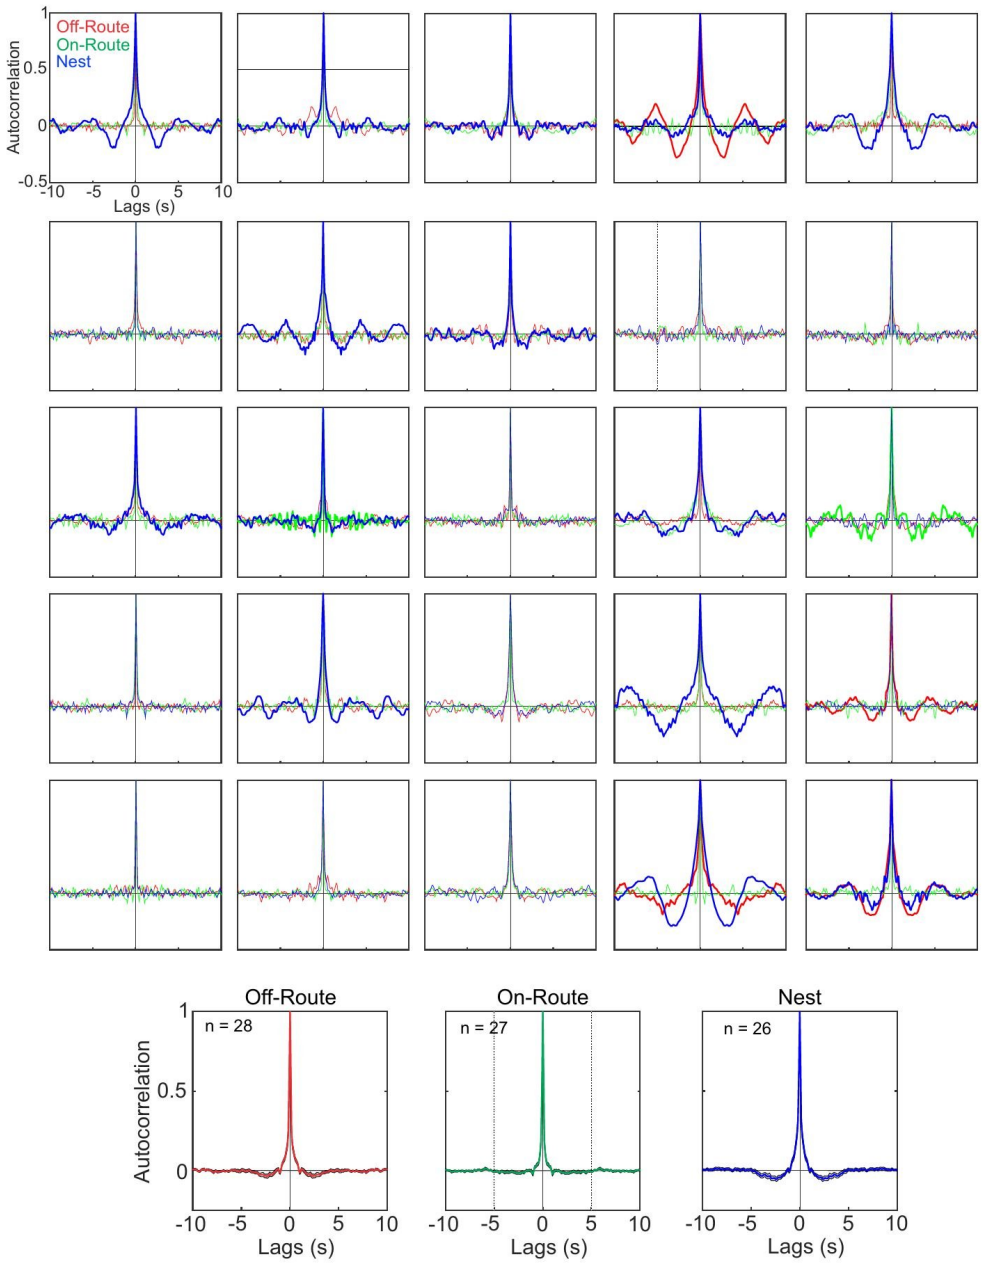


*Fig. 33: Comparison of auto-correlation functions of 5 minute time series of changes in path direction at the three locations*

*Off-route: red; On-route: green; Nest: blue for each of 25 ants that were tested at all three locations. Bottom panels show mean auto-correlations for all ants at the three locations.*

## Repeated release experiment

It was previously reported, for example in (Zeil et al., 2014b), that when *Myrmecia* ants are released at locations within their foraging environment of which they have never been to, but are still able to home, their initial heading may be quite different from the home direction. They start moving at a seemingly random direction within a roughly 90° cone centred around the home direction. They then correct their heading and turn to a nest directed route.

The question arises whether the initial heading distribution is due to the individual differences between the animals or is a result of a random process affecting all animals. To answer this question, I performed a repeated release experiment. The ants were released halfway along on their foraging corridor, but before they reached the nest, they were re-captured, taken back to the release point and released again. This procedure was repeated four times for each animal, to see whether the same individual heads off in the same initial direction or not.

Repeated release experiments have been performed previously with desert ants to study whether the animals keep a memory of the route they had taken previously (Collett, 2014), but I am not aware of repeated release experiments to assess the initial heading and subsequent path correction.

### Method

The experiment was performed in May, 2019 at the ANU Campus Field Station (-35˚16’49.87”S and 149˚06’43.74”E). After sunset, four *M. midas* were captured at the foraging tree on the outward journey. The animals were left in the laboratory overnight, with unrestricted access to food (sugar water).

In the following morning each animal was taken to a release platform located on the foraging corridor, approximately 6m from the nest. The transparent vial containing the animal was opened and placed, open side down, on the platform. A marking was made on the platform guaranteeing that the vial was always placed at the exact same place. When the animal ceased to show fight or flight reflexes after the disturbance and was on the platform, not touching the vial, the vial was slowly removed and the animal released.

The path of the ant was marked by placing small, coloured flags on the ground roughly 20 cm apart while following the animal, carefully avoiding to disturb her. When she had travelled 1.5-2 metres from the release point, she was re-captured in the same vial she had been before, then returned to the platform. After she calmed down completely, she was re-released and tracked again. Each animal was released 4 times. On the 4th release she was allowed to complete her journey and was monitored until she safely entered the nest.

The paths marked with the coloured flags were subsequently mapped using differential GPS (see (Narendra et al., 2013b) for details).

## Results

The paths can be seen in *Fig. 34*. Tracks with identical colour belong to the same animal.

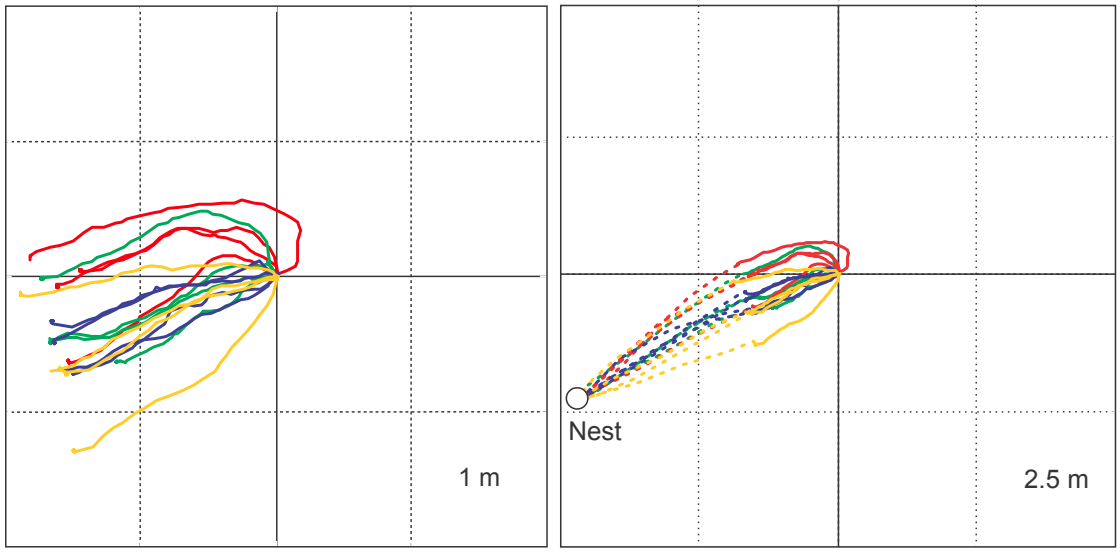


*Fig. 34: Repeated release paths*

*The figures show the paths of nest-bound animals released repeatedly at the same spot. Identical path colour indicates paths belonging to the same animal. The release point is on the foraging corridor, magnetic North is at the top. Dashed lines on the right panel connect the last path marker to the nest. Figure courtesy of Jochen Zeil.*

The figure clearly demonstrates that even the same animal shows large variations in her initial heading. It can also be seen that there is an animal dependent deviation: the red tracks, and to a lesser degree also the green

ones, show a systematic error towards North, which is subsequently corrected by the animal. How fast the heading is corrected seems to be animal dependent. The blue tracks at the initial 10-15cm show the same randomness as all the other tracks, but that animal very quickly assumed a nest-ward heading. The yellow animal, on the other hand, had diverging tracks even after 1.5m of travel.

The differential GPS has 1cm resolution and it is moved from marker flag to marker flag by the experimenter holding a stick which has the GPS receiver at its end close to the ground. The position is reported by the unit once a second. To get a finer-grained picture, the GPS positions were interpolated using piecewise cubic Catmull-Romm splines. That fine-grained path allows us to evaluate the ants' movements in more detail.

| Ant | B1 | B2 | B3 | B4 | Aver | Spread |
|---|---|---|---|---|---|---|
| Red | -66 | -22 | -2 | 69 | -5 | 135 |
| Blue | -140 | -135 | -82 | -49 | -101 | 91 |
| Green | -107 | -97 | -29 | -29 | -65 | 78 |
| Yellow | -173 | -103 | -108 | -56 | -110 | 117 |

*Table 4: Initial bearings of the animals*

*The table shows the bearings of the animals when they were 10cm from the release point. Angles measured in degrees, relative to North, clockwise (compass direction). The correct heading to the nest is -116°. The average of the four bearings is also given. The spread column is the angular difference between the largest and smallest bearing of the animal.*

With only four animals and four releases we do not have enough sample points to perform a statistical analysis, the data is presented in tabular form instead. *Table 4* shows the bearing of the animals when they were 10cm from the release point. Their initial bearings have a wide spread and, for the red and green animals even the average bearing is off target (the nest is at -116°) by a fair amount.

We can also ask how fast the animals correct their heading. *Table 5* shows the average bearings and the bearing spreads at 10, 20, 50 and 100 cm from the release point for each animal.

| Ant | B10 | B20 | B50 | B100 | S10 | S20 | S50 | S100 |
|---|---|---|---|---|---|---|---|---|
| Red | -5 | -15 | -47 | -82 | 135 | 115 | 73 | 44 |
| Blue | -101 | -101 | -104 | -107 | 91 | 68 | 26 | 19 |
| Green | -65 | -72 | -87 | -102 | 78 | 93 | 84 | 49 |
| Yellow | -110 | -113 | -110 | -113 | 117 | 92 | 70 | 57 |

*Table 5: Average bearing and bearing spread at different distances*

*The column Bn is the bearing of the animal at n centimetres from the release point, averaged over her four releases. The columns Sn display the spread of her bearings at n centimetres from the release point over her four runs. All figures are in degrees and the bearings are given as compass directions.*

All four animals assumed a roughly nest-ward bearing by the time they were 1m from the release point, consequently the spread of their bearings narrowed significantly. That is not the case at 0.5m, though. The red animal consistently started off at an incorrect direction and half a metre was not enough for her to correct for it completely. The green and yellow animals were heading at the general direction of the nest yet their bearings still show a significant spread at 0.5m.

Finally, let us examine the end of their paths before re-capture.

| Ant | H1 | H2 | H3 | H4 | Spread | AE |
|---|---|---|---|---|---|---|
| Red | -116 | -107 | -115 | -112 | 9 | 12 |
| Blue | -120 | -105 | -119 | -100 | 20 | 10 |
| Green | -112 | -94 | -96 | -110 | 18 | 18 |
| Yellow | -120 | -111 | -115 | -97 | 23 | 13 |

*Table 6: Final headings before re-capture*

*The columns H1 to H4 are the headings of the animals on the last 40cm on their paths before being re-captured. Spread shows the spread of the headings. The absolute value of the angular error between the heading and the direction needed to reach the nest from the re-capture point was averaged and shown in the column labelled AE.*

The average heading of the last 40cm of the path, just before re-capture was calculated and is shown in *Table 6*. The table also contains the spread of final headings and the averaged angular error. The angular error was calculated as the absolute value of the angular difference between the animal's heading and the direction she needed to go to the nest from the re-capture point.

It is clear from the table that the animals are heading towards the nest with a low spread in direction in subsequent runs. Their navigational confidence is

further attested by the low average error between their actual heading and the heading needed to get to the nest in a perfectly straight line.

## Conclusion

Ants with motivation to go to the nest were released at a location they were familiar with. Their path was marked and subsequently mapped using differential GPS. On their way to the nest they were captured, taken back to the release location and re-released. That procedure was repeated so that all four animals was released four times.

The number of animals and releases is not sufficient for statistical analysis, but their paths indicate that the ant's initial heading after release shows a large variation both by animal and by release. The error can also be quite large, an initial heading almost opposite to the target direction was recorded.

That initial error is then systematically corrected by the animal and the nest-ward heading becomes obvious by the time the animal travelled a metre from the release point. By 1.5-1.8 metres from the release point animals are, on average, heading towards the nest with little angular error.

These findings indicate that if we release an animal on a stationary trackball where motion parallax is not available for the animal to correct her heading, it is possible that she would assume a seemingly random heading and keep running in that direction.

# Proof of concept Antarium experiments

## Abstract

A series of experiments were performed to test whether the Antarium, a reconstructed visual reality arena that I designed, is conceptually solid. Tethered ants were put on a trackball, placed inside the arena and shown natural scenes. The ants' movements on the trackball were recorded and analysed. The analysis identified several problems with the device, the most important of which is that its LED wavelengths are a poor match to the spectral sensitivity of *Myrmecia* eyes. Nevertheless, the experiments have demonstrated that animals can and do navigate in the reconstructed visual environment presented by the Antarium, the device is conceptually solid and can be used in navigation experiments that would be very hard or impossible to perform without such an arena.

## Introduction

This chapter describes the experiments that were performed using the Antarium with the aim of proving that the animals accept the Antarium as a representation of their environment, that they navigate in it and that the Antarium can be used to ask questions about the navigational mechanisms used by ants which would be very hard or even impossible to ask without the use of reconstructed visual reality.

As I will discuss in the chapter, the Antarium suffered from a number of design flaws, most prominently the wavelengths of the LEDs in the device being a very poor fit for the spectral sensitivities of ant photoreceptors. Despite its flaws, however, the experimental results clearly indicate that the Antarium is conceptually solid because the animals are able to navigate in it.

## Methods

All experiments share common elements, which for the sake of brevity are not repeated for each individual experiment. This section gives account of those common methods and only the deviations or experiment-specific details are referred to in experiment-specific method sections.

## Animal handling

For experiments 1 and 2 the animals were mounted on the trackball the same way as was described for the outdoor experiments: they were anaesthetised by cooling, then their mesonotum (thorax) was glued to the tether using UV hardening dental glue. Flexibility of their movement was guaranteed by a roughly 1mm long single thread of dental floss between the animal's mesonotum and the metal pin of the tether.

### *Magnetic mounting*

Following advice from Antoine Wystrach, from experiment 3 onwards a different technique has been used. The animal is mechanically immobilised for a few seconds while a drop of magnetic paint is placed on her. The paint contains very fine iron particles and after drying it is paramagnetic. For the experiments Rust-Oleum Primer Magnetic Paint was used, purchased from a hardware store. The paint was applied using a toothpick.

The ant is then placed in a lid-less plastic container the side walls of which have been wiped with liquid paraffin. The ant can walk in the container, but the walls are too slippery for her to scale them. She is kept there until the magnetic paint dries, about 15-20 minutes. In the experiments rectangular plastic food containers were used and the liquid paraffin was purchased over the counter in a pharmacy.

A small but strong rare-earth magnet is attached to the tether; it is glued to the end of the dental floss thread that formerly was used to be glued to the animals. The magnets used in the experiments were purchased online from https://www.supermagnete.fr with product number S-1.5-0.5-N. These magnets have NdFeB composition, with Ni-Cu-Ni casing. They are 0.5mm thick and 1.5mm diameter discs with a mass of 0.54g. Each magnet attaches to a flat iron surface with approximately 0.33N force.

The animal is picked up from the container by placing the magnet on the tether against the droplet of magnetic paint on her back. She is then carefully lowered onto the trackball. The magnet is strong enough to keep her attached to the tether on the freely rotating ball. Since the magnet is glued to a thread, it has

flexibility and the animal remains attached even if she performs banking or pitching movements. On the other hand, it is easy to detach the animal at the end of the experiment, either using non-magnetic forceps or by offering the animal a solid surface (for example, the vial) that she can hold onto.

This technique facilitates the running of experiments in a streamlined manner. An animal can run in the Antarium while the next one's paint is drying and a third one is still feeding. Unlike the glueing technique, where the preparation time of an animal often exceeded the time she ran in the Antarium, the magnetic mount allows a session's output to be limited only by the time the animals spend in the virtual environment. In addition, it can be done without a microscope, demands less manual skill from the experimenter and does not require anaesthetising the animals.

It should be noted that Fleischmann *et al* (2018a) reported that *Cataglyphis* use the Earth's magnetic field as a reference during learning walks. It is not known whether *Myrmecia* also do so or not. Nevertheless, a strong magnet at the back of the ant will completely obliterate the geomagnetic field and very likely overwhelm the animal's magnetic sensors. However, as the paper pointed out, the animals use the magnetic compass only during learning walks, presumably to calibrate their celestial compass. Later in life they rely exclusively on the latter. Since the Antarium was used with experienced ants, the magnetic mount is very unlikely to have affected their navigational behaviour.

## Data Analysis

I developed custom code, using version 8.6 of the Tcl/Tk language[13] to process the data recorded during the experiments. Statistical analysis was performed using version 3.5 of the R package[14] invoked from the Tcl scrips. The video recordings of the animals in the Antarium were made by a Raspberry Pi V1

13 *Tcl is a free, open-source, interpreted dynamic programming language, about the same age and capabilities as Perl and Python. Tk is a graphics toolkit that was developed for Tcl but was then ported to other scripting languages, including Perl and Python.* https://www.tcl.tk/

14 *R is a free, open-source, interpreted functional programming language oriented for statistical analysis and data mining.* https://www.r-project.org/about.html

camera (1280x960 pixels, 30fps, progressive scan) mounted on the apex of the Antarium. The video recordings, after transcoding to MP4 format, were used to verify the proper operation of the analysis code and were saved for later analysis, not covered in this thesis.

***Nomenclature***

Before the Antarium experiments are described, its is warranted to clarify the meaning of terms that can be ambiguous.

Ants live in a 3-dimensional world and the Antarium projects the views of a 3-dimensional scene. However, since ants walk on the ground surface, their altitude is assumed to be constant.

The animal's *position* or *location* is thus expressed as a Cartesian (X,Y) coordinate pair. The positive Y axis points to the (magnetic) North while the positive X axis points to the East. The origin of the coordinate system depends on the experiment.

As the animal moves, her *velocity* vector is expressed in polar coordinates. The magnitude of the vector is her *speed*, the displacement over unit time; while the angle of the vector (the direction in which the ant is moving) is the *heading*. The heading is expressed as a compass direction, that is, it is measured clockwise from North (the positive Y axis). That deviates from the mathematical convention of measuring angles anticlockwise from the positive X axis.

When the animal moves from a start point, the angle of the vector from her initial location to her current one is her *bearing.* Like the heading, it is measured clockwise from North.

The trackball in the Antarium records the animal's movements as displacements relative to two orthogonal axes. These axes are defined by the physical construction of the trackball and its orientation inside the Antarium. A particular direction inside the physical Antarium device was arbitrarily denoted as North. A preprocessing step aligns the coordinate systems of the trackball and the Antarium. Positions in that coordinate system are referred to as *ball coordinates* or *real coordinates*.

The Antarium projects a real-world scene. That scene has its own North direction, namely the magnetic North of the real-world environment depicted by the projected 3D model. That North may or may not be aligned with the North in ball coordinates. Therefore, the projected visual scene has its own coordinate system, called *scene coordinates* or *virtual coordinates*. The virtual and real coordinate systems have the same metric. Thus, distance and angular difference and, consequently, speed and angular velocity are the same in both coordinate systems, but position and heading are different if the two coordinate systems are not aligned. Normally the coordinate systems are aligned at the beginning of an experiment, but during the experiment the scene might be rotated or even translated. That causes the two coordinate systems to rotate relative to each other and moves their origins apart. The processing software keeps track of the scene changes and can map the ant's path from the ball (real) coordinate system to the scene (virtual) coordinate system and *vice versa*.

Unless explicitly stated otherwise, location and heading are in the scene coordinate system. When the text refers to ball coordinates, it always qualifies this accordingly.

A further point that must be discussed is the speed of the animals. Some speed data from the experiments might seem unrealistically high. However, the data are real and there is a simple explanation for the unusual values. In their natural environment, that is, on uneven ground littered with debris and vegetation, the ants used in the experiments walk at an average speed of around a metre per minute. On the other hand, if one puts them on a flat, smooth and obstacle free surface, such as a table, they move much faster – it takes them a few seconds to run the length of a 1.5m table. The trackball is like a table: flat, smooth and with no obstacles. Consequently, the speed of the animal on the trackball is much higher than on grassy ground. When they run, they can reach peak velocities around 1m/s and when they walk determinedly, without looking around much, their average speed is around 0.1m/s; significantly faster than outdoors.

***Preprocessing***

The 5 or 10cm diameter trackball provides the position information as raw displacement numbers from its sensors; details can be found in (Dahmen et al., 2017). The data rate is around 277 samples per second, with a slight variation.

The displacement values are first converted to metres using scale factors which were experimentally determined for the trackball during calibration. The displacement values are subsequently rotated by a constant amount so that they are aligned with the real (or ball) coordinate axes. They are then accumulated to reconstruct the ant's path.

The final result of the preprocessing is a stream of absolute (X,Y) coordinates with roughly 200μm resolution. The coordinate sample points are time-stamped with 100μs precision. Consecutive samples are approximately 3.6ms apart, in a coordinate system that is fixed to the Antarium.

With the 277Hz sample rate each sample represents only a small displacement, especially when the ant is walking slowly. No useful heading or speed information can be extracted over such short distances. Also, the raw position signal is quite noisy.

To rectify these problems, the coordinate stream is first filtered with a linear phase finite impulse response (FIR) low-pass filter with a cut-off frequency of 12Hz (see *Fig. 35*) then down-sampled to 30Hz, which matches the frame rate of the video recording of the animal. The resulting (X,Y) stream is then augmented with experiment specific information such as the scene's orientation within the Antarium, movable object locations, colour transformation information, video synchronisation points, suspension of projection and so on.

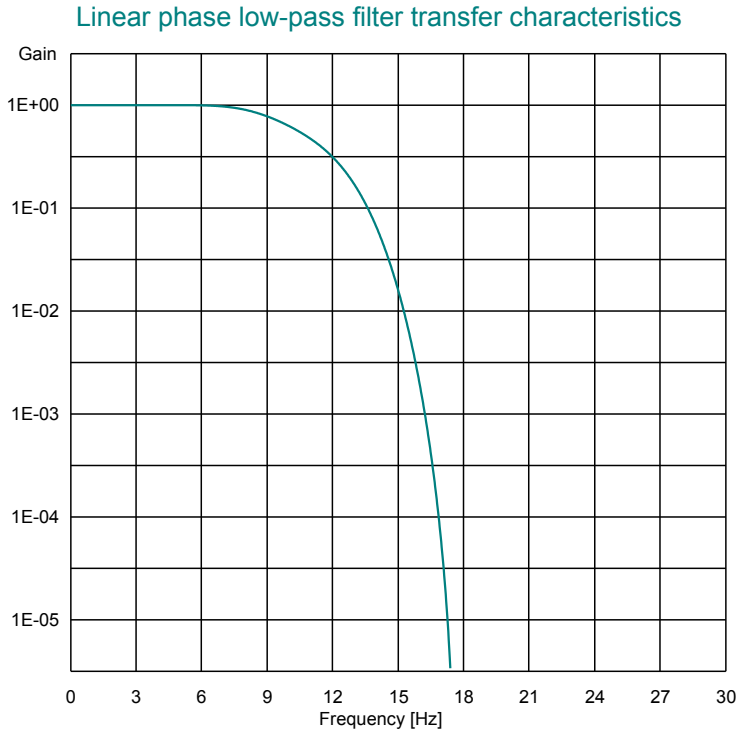


*Fig. 35: Transfer characteristics of the down-sampling filter*

***Ball to scene coordinate transformation***

Only the ball tracks are recorded, but the system keeps track of scene rotations and translations, and calculates the scene coordinates from that. The ball track is always continuous. Initially the ball and scene coordinate systems are aligned, therefore until the first scene rotation the two paths are identical.

When the scene is rotated, an “anchor” is dropped in both coordinate systems. Then, subsequent ball path coordinates are rotated relative to the anchor in the direction opposite of the scene rotation and then added to the anchor point in the scene.

The process is shown in *Fig 36*. If the animal was translated in the scene, then the scene track from that point on is simply offset by the translation vector. Naturally, in that case the scene path is not continuous.

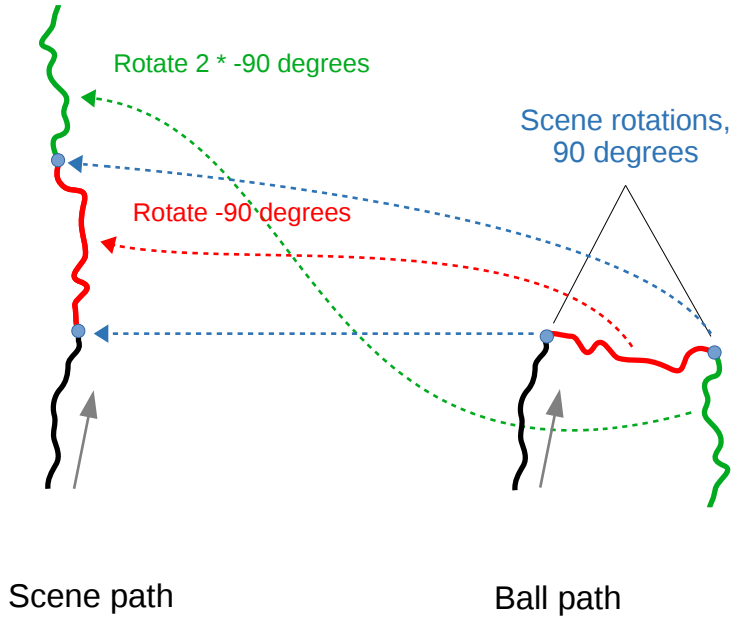


*Fig. 36: Real to virtual coordinate transformation*

***Path to vector sequence transformation***

Watching the video recordings of the ants' movements on the trackball reveals a wide range of movement patterns. Sometimes they move at a constant speed. At other times they speed up or slow down with no apparent reason. Occasionally they stop, scan the scene by turning on the spot, select a heading direction and start moving again. They can also stop for grooming, which can take 15 seconds or more.

We are interested in navigation so when the animal's path is analysed, often what we want to know is which way she was moving. However, when she is not moving at all, the concept of direction has no meaning. We need a speed-independent representation of the path.

The trackball supplies the path as samples of the animal's location with the time between samples being constant, regardless of the distance between the samples. We can re-sample the path so that consecutive samples are at equal distance from each other, regardless of how much time elapsed between them. The re-sampling process thus turns the path into a series of displacement vectors, all of equal length. The size of these vectors is an arbitrary choice. Too small vectors result in a very noisy sequence. Too long vectors lose a lot of

detail. In the experiments 2.5cm, roughly the body length of the animals, was chosen, but it was checked that varying the vector length between 1cm and 10cm does not change the analysis outcomes.

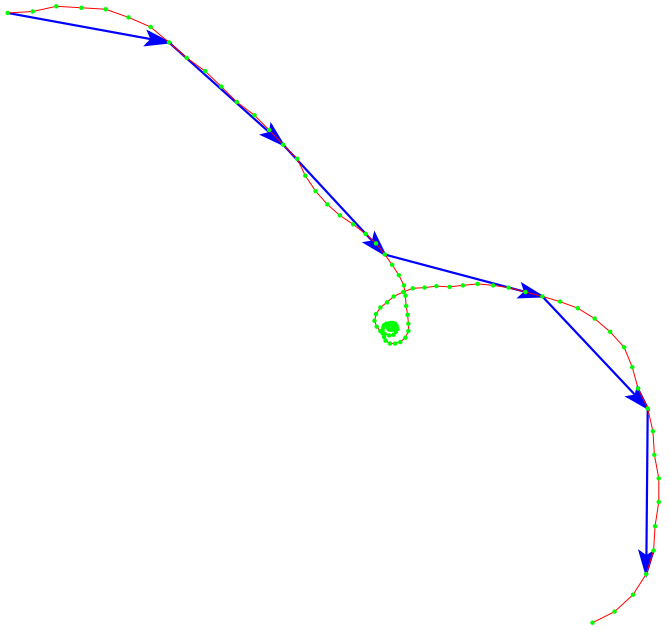

*Fig. 37: Vector sequence transformation*

*The thin red line is the path of the animal. The green dots are the 30Hz sample points from the trackball, after preprocessing. The blue arrows are the 2.5cm equal length vectors after the transformation. Initially the animal moves at around 0.13 m/s. In the middle of the path the she slows down, scans the environment with minimal translation, chooses a new direction and speeds up again. The vector sequence format replaces the entire scanning process with a single vector.*

Each vector represents a small section of the path of the animal. Since all vectors are of equal length, they differ only in their direction. They can therefore be called the average heading of the animal over the given section of the path. The transformation thus turns the path into a format that can be subjected to circular statistical analysis to determine whether the animals preferred some direction(s) more than others. The process of transforming the path into the vector sequence is illustrated in *Fig. 37*.

Care should be taken when using such vectorisation. Statistical calculations treat the vector directions as a random variable. The order of the individual vectors is ignored. Consider, for example, the left-hand side diagram of *Fig. 38*. The red and the green path are built from the same vectors, therefore the statistics will be identical for them. It is obvious, however, that the fictitious animals belonging to the two paths show very different behaviour.

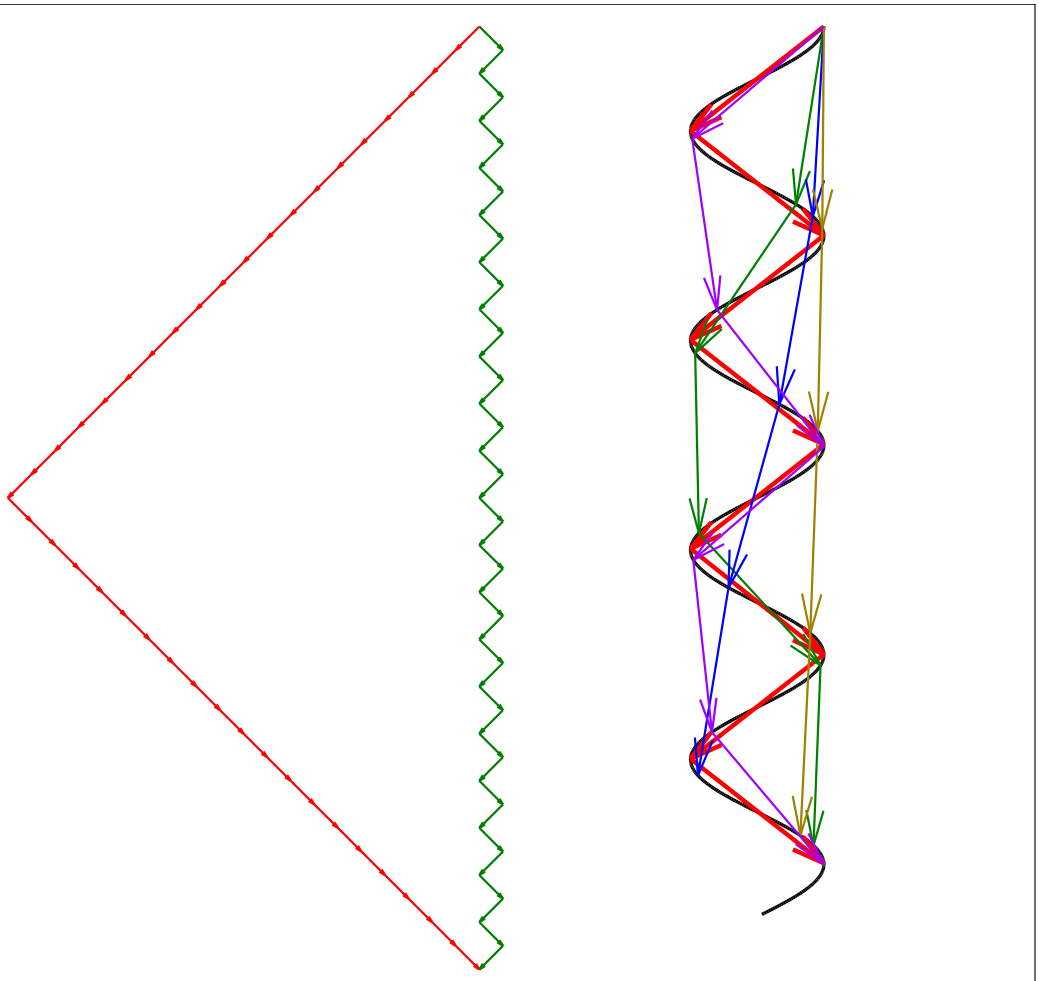

*Fig. 38: Path vectorisation artefacts*

*See the text for an explanation of what the diagrams demonstrate.*

Nevertheless, while in theory it is possible that the animal is moving in one direction while the statistics indicates a very different distribution, it is very unlikely to occur with real animals.

On the one hand, to get the completely incorrect statistics, you need a severe aliasing effect where the sampling catches a periodic behaviour. Consider the

right hand side diagram on the figure. Our hypothetical animal moves on a perfectly sinusoidal path (black). If the vectorisation happens to match exactly the distance between adjacent peaks of the sinusoid, then we will get the aliased path (red arrows) that will give us entirely incorrect information about the direction of the animal. However, even a small change of the chosen vector length disturbs that greatly: all the other colour arrows are within 10% in length of the red arrows, yet they do not show the same wildly incorrect distribution.

The other important fact is that real animals do not move with perfect periodicity or along geometrically perfect lines, circles, sinusoids and the like. They move at variable speed, with a fairly noisy heading. They tend to keep a more or less constant average heading for far more distances than the vector lengths. So the vectorisation is very unlikely to get in resonance with any quasi-periodic behaviour the animals produce.

To be on the safe side, as it was mentioned before, I checked that the statistical results did not change significantly while I varied the vector length over a wide range.

***Statistical methods***

There are three kind of data that arise from the experiments.

If the samples are particular events from a fixed set of possible outcomes, the number of occurrences of each outcome can be counted, forming a contingency table. We usually are interested in whether two contingency tables came from the same distribution or not. Fisher's exact test for count data is used to test that. The reason for choosing this particular test is that the number of samples is usually much lower than what the more commonly used chi-squared test requires.

When the experiment results are real numbers and two sets of observations need to be compared to see whether they come from the same (unknown) distribution or not, then two tests are employed. Both are non-parametric, warranted by the low number of samples available. One test is the Mann-Whitney U-test, also known as Wilcoxon rank-sum test. It is sensitive to the change of the mean of the sample sets. However, it is not sensitive to the shape

of the distribution function. The uniform distribution on the [-1,1] interval has 0 mean and 0.3333 variance. The Mann-Whitney test would not differentiate a sample set from that distribution and one from a Gaussian distribution with the same mean and variance. To mitigate that problem, the Kolmogorov-Smirnov test is also run on the data. It is less sensitive to changes of the mean but it works by comparing the empirical distribution functions of the two data sets to each other. Therefore, it does discriminate distributions with different density kernels even if their means and variances happen to be the same. When analysing experimental results, both the Mann-Whitney and the Kolmogorov-Smirnov test are run on the data and both p-values are reported.

Sample sets of angles are tested using circular statistics. Two questions are usually asked: whether a sample set is circularly uniform and whether two sample sets belong to the same distribution. The canonical test for circular uniformity is the Rayleigh test. However, that test is most sensitive to unimodal deviation from the uniform, in particular, to the *von Mises distribution*, which is the circular equivalent of the normal distribution. The Rayleigh test is especially insensitive to bi-modal data where the peaks are half a circle apart. In a 1970 paper Batschelet suggested the use of an alternative test, the non-parametric Rao spacing test for circular uniformity if multi-modal distribution is suspected (Batschelet, 1970). It has also been argued recently that the Hermans-Rasson non-parametric test should be preferred over the Rayleigh test (Landler et al., 2018, 2019). The Hermans-Rasson test is almost as sensitive for von Mises style deviation from the uniform as the Rayleigh test, but it was designed to deal with multi-modal distributions. All three tests are run on circular data and the smallest p-value, together with the test that produced it, is reported.

When two sample sets from unknown circular distributions are to be tested whether they come from the same distribution or not, the Watson-$U^2$ test is used. It is a non-parametric test, which can be used even when the sample numbers are low.

The significance level for all tests used in this analysis is 0.05. When two tests (the MW and the KS tests) were run on the same data, the Bonferroni correction is applied and the statistical significance level is decreased from 0.05 to 0.025.

A final comment: It seems that sometimes papers assume that because a test does not reject the null hypothesis, the alternate hypothesis must be false. That is not the case. A p-value above the significance level only means that the particular test method, with its particular assumptions, could not establish that the probability of a process generating the observed data while the null hypothesis being true is less than the significance level. Or, to use legal terms, not proven innocent, just could not be proven guilty beyond reasonable doubt.

# Experiment 1

## Aim

Verify whether the Antarium is operating as expected.

## Methods

### *Experimental design*

At the time of the experiment, the Antarium was operational, but several features described in Chapter 2 were introduced after, and as a consequence of, the results obtained with its first test run, described below.

The experiment was carried out at the ANU Campus Field Station, in October, 2018, in two stages. First, on a Friday, around 10:00 a *Myrmecia croslandi* was captured at the foraging tree. She was treated and mounted on the 5cm diameter trackball the same way as described for the outdoor experiments. The trackball was placed into the Antarium and the initial scene was set to a point halfway between the nest and the foraging tree, albeit without a distant background. The Antarium was set to closed-loop mode, the animal could freely explore the virtual environment by running on the trackball. While she was walking in the virtual world, occasionally the scene was rotated around her by 90 degrees. The rotations were controlled manually, without a pre-set schedule. Furthermore, she was also virtually displaced at occasions by changing the projected scene, but the displacement was such that she was always within the area that she was expected to be familiar with. The animal's progress was monitored on a computer screen that displayed the 3D model and the ant's path in it.

On completion of the experiment she was detached from the tether and released near the nest. She was monitored until she reached the nest entrance.

In the following week, a further 15 animals were tested over 3 non-consecutive days, 5 animals per day. The 15 animals were treated the same way as the first one. In addition, after the experiment in the Antarium, the whole trackball apparatus, with the still mounted ant, was taken outdoors and placed at the real location of the release point in the virtual world. The animal's movement on the

trackball outdoors was observed and the observation recorded, but the actual trackball data were not stored on a computer. All the ants then were released the same way as the first animal.

## Results

The first ant behaved perfectly. She followed scene rotations and depending on where she was displaced in the virtual environment, she was heading either to the nest or to the foraging tree. She reached the nest in the virtual world.

That was not true for the next 15 animals:

Two ants were navigating in the virtual world, but less confidently than the first animal. They followed rotations mostly but not always. In general they were walking towards the nest, but sometimes they changed their direction and started to walk away from it. Ultimately they turned back to it and reached the nest at the end.

One more animal was oriented, but not towards any identifiable goal. Still, it was obvious that she was not doing a random walk.

Five ants were walking but in a seemingly random manner. They did not show a directional preference and reacted neither to scene rotations nor displacements. They also frequently stopped for grooming.

Seven animals showed catatonic behaviour. They looked around for a few seconds then became completely motionless, including their antennae. If disturbed by putting forceps close to them, they ran away from them, but as soon as the stimulus ceased, they stopped moving again and just sat motionless until they were taken out from the Antarium.

All 15 animals, when taken outdoors, immediately started running on the trackball towards the nest or, on a few occasions, towards the foraging tree.

## Discussion

### *Antarium deficiencies*

The experiment with 16 animals resulted in an 81% failure rate. Clearly, for *M. croslandi*, the Antarium was not operating as expected. The following problems were identified:

- The LEDs in the Antarium were selected when the spectral sensitivity of *Myrmecia* eyes were not known and the LEDs selection was driven by availability and cost considerations. Recent measurements (Ogawa et al., 2015) show that *Myrmecia* ants have three types of photoreceptors with peak sensitivities around 375, 430 and 550nm. The LED spectra are a very poor match to the animals' eyes, as shown in *Fig. 39*.

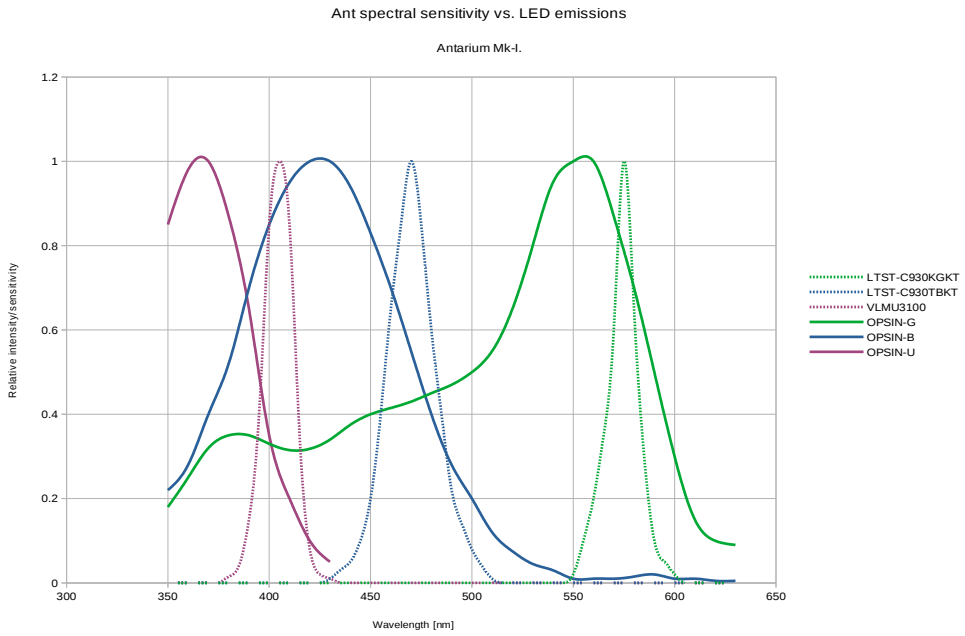


*Fig. 39: Spectral sensitivity of Myrmecia eyes and emission spectra of the Antarium LEDs.*

*The dashed lines are the LED emission curves, with the LED type in the legend. Spectral sensitivity data are from (Ogawa et al., 2015) and the LED radiation spectra are from the individual LED datasheets*

- The UV LEDs of the Antarium were turned off, thus the animals were deprived of all UV cues.

- Even if the UV LEDs had had been on, they are of incorrect wavelength, their intensity is much lower than what a daylight sky provides and they are sparsely distributed. Thus, at this stage, the Antarium does not provide reliable UV contrast.
- It is possible that the light intensity in the Antarium was too low for the exclusively day-active *M. croslandi*, the photoreceptors of which are adapted to high daylight intensities (Greiner et al., 2007; Narendra et al., 2011).
- The Antarium's circuit boards were painted black by the manufacturer, but with a high-gloss lacquer. Panels reflected a significant amount of light from the opposing panels, causing a loss of contrast.
- The frame rate was limited by the rendering software to only around 40fps. It is not known whether that is sufficient for an ant to see continuous motion or not.
- At the centre of the Antarium the viewing angle of the LEDs varies by about ±20˚ across each panel. The apparent LED intensity changes with the viewing angle which results in the panorama being modulated in intensity by a pattern that is fixed to the Antarium. A moving scene superimposed on a stationary pattern might be confusing to the animal.

***Animal behaviour***

There was no explanation for the catatonic state observed with some of the animals. The literature does not report that behaviour. However, advice was received from Eric Warrant that in laboratory conditions certain flicker frequencies in the low-hundreds Hz range were observed to trigger it (personal communication). The Antarium's flicker frequency is around 9kHz by design and it was verified by oscilloscope, thus the flicker hypothesis was discarded.

In a completely unrelated experiment, performed later, the catatonic state was observed under natural circumstances: *Myrmecia midas* captured in the evening were released early in the following morning, with the intention of tracking them to the nest. It was a cold morning and there was dense fog at the

Campus Field Station with very low visibility. From the release point (located on the foraging corridor) neither the nest nor the foraging tree (or any other tree) was visible; the visual scene was basically featureless uniform grey above and beyond the nearby ground debris level. Five animals were released and all five went motionless immediately. Each one was observed for 10 minutes and during that time none moved at all. They were then re-captured and subsequently released close to the nest entrance, from where they all went into the nest without hesitation. That observation suggests that complete deprivation of visual information can evoke the catatonic state, although it is not known what other factors (e.g. temperature) might be responsible for triggering it. Furthermore, when the animal enters the nest, it is in complete darkness and does not become catatonic. It is possible that the simultaneous presence of light and absence of visual features trigger the behaviour. Further experiments will be needed to understand this peculiar phenomenon.

***Population variation***

While most animals could not navigate in the Antarium, some could. That was also observed in subsequent Antarium experiments. There is no known explanation for that fact. It can be genetic and/or developmental variation of the opsins and/or the visual neural circuits that make those animals more sensitive to the Antarium's projections. However, currently that hypothesis has no experimental support; little appears to be known about individual variations in photoreceptor properties and absolute or spectral sensitivities. It is also possible that there is no variation in the visual capabilities of the animals, but more experienced individuals can make more use of the parsimonious visual information provided by the Antarium. This question will remain open until targeted experiments are performed.

## Antarium Improvements

In light of the first experiment, several modifications were made to the Antarium:

- The internal surface of the device was covered with low reflection black cardboard, with cut-outs for the LEDs. That reduced the reflected light to

5 to 7% through the 400-700nm wavelength range, as measured using a USB-4000 spectrometer (Ocean Optics, Dunedin, Florida, USA).

- The intensity variation due to parallax was fixed, by lowering the intensity of the LEDs closer to the optical centre of the panels. While variation disappeared that way, the price to be paid was the lower overall intensity of the projected image.
- The 3D model of the environment was refined and it was augmented with a remote static background.
- The rendering code was improved and the sustained frame rate reached the physical limit of the projector, over 190fps.
- Some bugs in the FPGA code and in the data distribution microcontroller code were fixed.
- A fan was installed at the base of the device to provide a continuous air exchange between the inside and outside of the Antarium, for details see *Thermal considerations* on page 106.
- A Raspberry PI camera was mounted on the Antarium which, through a roughly 1cm opening at the top of the device, could film the ant on the trackball during the experiment.
- The UV channel was activated by adding a uniform UV illumination to the sky of the model. The lack of polarisation, the sparse distribution of UV LEDs and their incorrect wavelength greatly decreases the usefulness of the UV channel, but, as it turned out, it still helped the animals measurably.

Despite all the other improvements, the LED wavelengths remained a very poor match for the animals' eyes. That particular issue will be addressed in detail later.

# Experiment 2

## Aim

With these changes implemented it was decided that a static image experiment should be carried out with crepuscular or nocturnal animals for which the Antarium may provide sufficient light intensities and to verify that the enhancements had actually improved the Antarium's performance, even though the LED wavelengths could not be corrected.

## Methods

### *Experimental design*

The experiment was carried out at the ANU Campus Field Station[15] (-35˚ 16' 49.87" S and 149˚ 06' 43.74" E), in December, 2018.

Just after sunset *Myrmecia midas* were captured at the foraging tree on the outbound leg of their foraging trip. They were fed, then mounted on the 10cm trackball in the Antarium the same way as described for the outdoor experiments.

The Antarium was set up in open-loop mode. It presented a static image to the animal; the animal's movements on the trackball had no effect on the projection.

Four scenes were defined. Three of these showed real locations at the ANU Campus Field Station:

- *Nest:* the view at the nest location. There were two nest entrances about 4 metres apart. Ants were observed to come out from one entrance and enter the other, so it was assumed that it was a single, very large colony where the ants should be familiar with either entrance. The entrance marked `Nest N` in *Fig. 40* was used for this view.
- *Familiar:* the view at a point on the foraging corridor, marked `Familiar` in *Fig 40*. It was about halfway between the `Nest C` entrance and the

15 *In May 2019 the ANU Campus Field Station was bulldozed and the laboratory demolished, to pave way for a commercial student accommodation venture. The satellite photo from Google Maps that is used to identify scene locations was archived in 2018.*

`Foraging tree`, or about ¼ of the way from the `Nest N` entrance to the tree.

- *Unfamiliar:* a view of the real environment, but from a place far away from the foraging corridor. Ants from that particular colony do not normally go to that location. In *Fig 40* the location is marked as `Unfamiliar`.
- *Unstructured:* a computer generated scene with a horizon line, but no other visual features. Higher elevations are bluer and lower ones are generally greener, but the azimuthal uniformity of the scene is only broken by a small amount of pseudo-random noise.


*Fig. 40: Experiment locations*

*Satellite image of the Campus Field Station with the various locations marked and labelled. The labels Nest C, Nest N, Familiar, Unfamiliar and Foraging tree mark locations described in the text. Release marks the location where the animals after the morning experiment were released and subsequently tracked to the nest. Laboratory marks the building housing the Antarium. The author's Car happened to be parked beside the building when the image was taken. Satellite photo from Google Maps, archived in 2018.*

***Scene preparation***

The 3D model was augmented with a panoramic background, resulting in a photorealistic image, as shown in *Fig. 41*.

*Fig. 41: Panoramic image of the familiar location, cylindrical projection*

*The foraging tree is marked with the yellow cross and the Northern nest entrance with the red cross. The building on the left hand side, behind the nest entrance, is the laboratory. At the bottom of the image is a metal rectangle with an orange spirit level for the panoramic camera.*

*Note that in the Antarium only a small part of the metal platform will be visible, due to the -50° elevation limit of the device. Image courtesy of Trevor Murray.*

The red channel is then turned off as the animals cannot see that part of the spectrum. Sky features (clouds) are removed as much as possible and a uniform UV illumination is added above the skyline.

Since the red component of the image was not utilised, that channel was used to represent the UV illumination. That is the cause of the magenta sky in *Fig. 42*.

*Fig. 42: Familiar location, red component and sky features removed, UV added above the skyline.*

*The magenta sky is due to the otherwise unused red channel being utilised to represent the UV illumination.*

The image is then mapped to the Antarium's 19,855 blue-green pixels and its 495 UV pixels, as shown in *Fig. 43.*

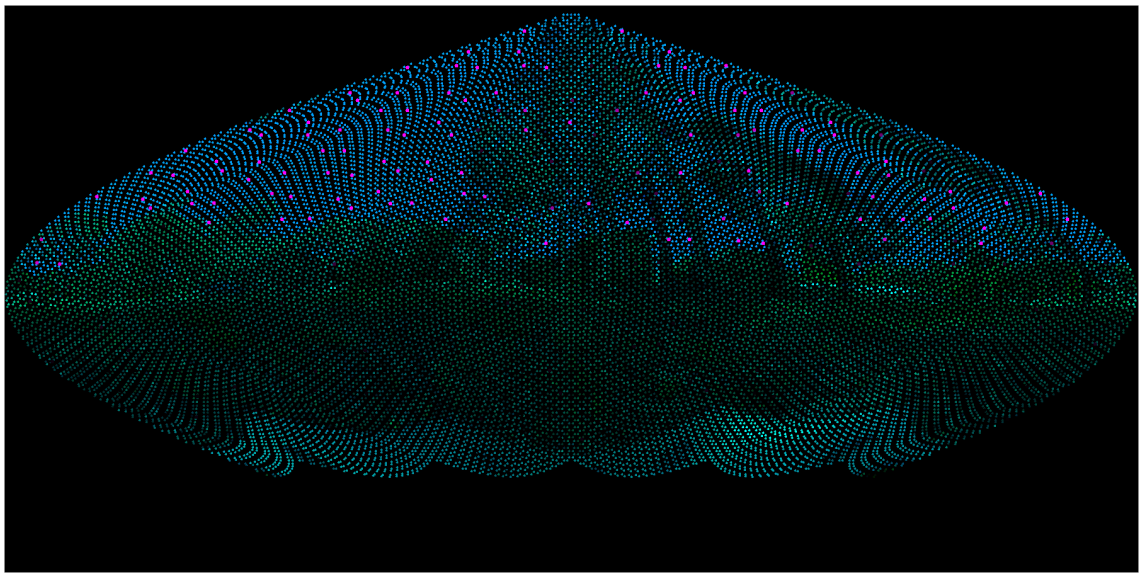

*Fig. 43: The familiar view mapped to the Antarium's LEDs*

*Sinusoidal pseudo-cylindrical projection. The UV LEDs are shown in magenta.*

To better visualise what the projected image looks like, *Fig. 44*. shows the Antarium's pixels mapped onto a sphere.

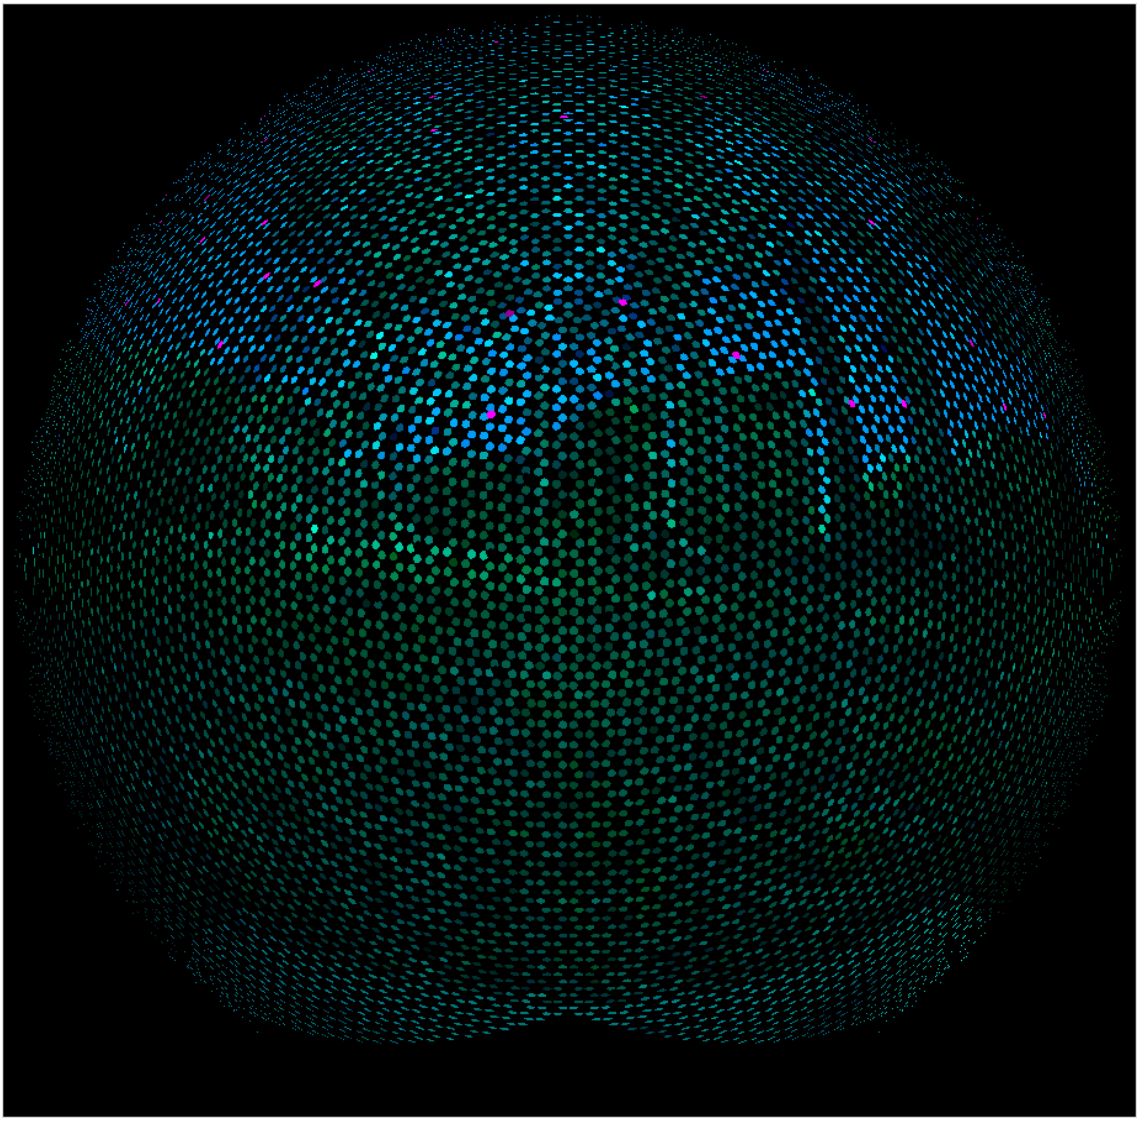

*Fig. 44: The LEDs rendered on a sphere, azimuthal projection*

The other two natural scenes used in the experiment are shown in *Fig. 45*.

*Fig. 45: Other scenes used in the experiment*

*The nest view (top) and the unfamiliar view (bottom), after colour processing. Note that the concrete bollard with the metal rod that is visible on the unfamiliar view's left hand side is not the same what marks the nest entrance on the nest view. Images courtesy of Trevor Murray.*

### *Experimental paradigm*

The experiment ran over a week, with one day skipped due to weather. 4 to 6 animals were processed each evening, 21 animals all together.

The ant was mounted on the trackball in the Antarium. The 4 scenes were shown to the animal, each for approximately 5 minutes. The presentation order was randomised by the computer. In addition, each scene was rotated 4 times,

with an approximately 1 minute delay between rotations. The rotations were always +90˚ or -90˚, randomly selected by the operator. Both the rotations and the scene changes were instantaneous. After the presentation of all four scenes, the ant was detached from the tether and put back to her vial. She was given sufficient amount of food and was kept in the laboratory overnight.

The following morning the animals were mounted on the trackball and were again confronted with the same four scenes, following the same protocol as in the evening experiment except that in several morning experiments the nest scene was omitted to save time. With the completion of the morning experiment, the ants were released on their foraging corridor and monitored until they reached the nest, or, in some cases, the foraging tree. On several mornings the released ants were tracked by placing small plastic markers behind the animal, without disturbing her, on her path. The markers were then mapped using differential GPS.

To ensure that each animal was used only once, the ants which were tested in the Antarium were marked with a dot of paint on their abdomen and only non-painted animals were used for subsequent experiments.

During the experiments the trackball data, the scene selection and rotation information were collected. In addition, the animals on the trackball were also video recorded from above, at 30fps.

***Trackball data preprocessing***

The recorded trackball paths were first preprocessed, as described in *Preprocessing* on page 162. Using the scene rotation data and the paths on the ball, the ant paths in scene coordinates were generated.

***Questions and Predictions***

To assess whether the animals can navigate in the Antarium, the following should be examined:

1. Do the animals show orientation, that is, a non-uniform distribution in their heading?
2. Do the animals react to scene rotations?

3. Do the animals differentiate between the four scenes projected to them?
4. Do the animals behave at the various scenes as they would in their natural environment?
5. Are there significant differences between animals?

For question 4 the expected behaviour at the four scenes as reported in the result section of (Murray et al., 2020) which can be found on page 131, is the following:

- At the familiar scene the animal should walk towards the nest or, possibly, towards the foraging tree.
- At the unfamiliar scene she should go towards the nest or, in case she cannot recognise the scene at all, do a random search.
- At the nest she is expected to perform a random search for the nest entrance.
- When viewing the unstructured scene she is deprived of all visual cues, so there are no particular expectations but the most likely behaviour is a random walk.

Question 5 arises from the results of Experiment 1. In that experiment some animals could navigate in the virtual environment but the majority could not. Considering the mismatch of the LED spectra and the sensitivity of the animals' eyes that could be the case in this experiment as well.

The reason for running evening and morning experiments as well was that it gave us more data. There were no predictions or expectations between the differences or similarities, except that one would assume that the same animal would show similar behaviour in both cases. In addition, it is known from outdoor experiments that in the evening the animals are more motivated to go to the foraging tree while in the morning they are inclined to head towards the nest.

## Results

### *Q1: Do the animals show orientation?*
### *Q3: Do the animals differentiate between scenes?*

In a first step to answer these questions I determined the animals' distance reached and bearing after 30s (*Fig. 46*), 2 min (*Fig. 47*) and 5 min (*Fig. 48*).

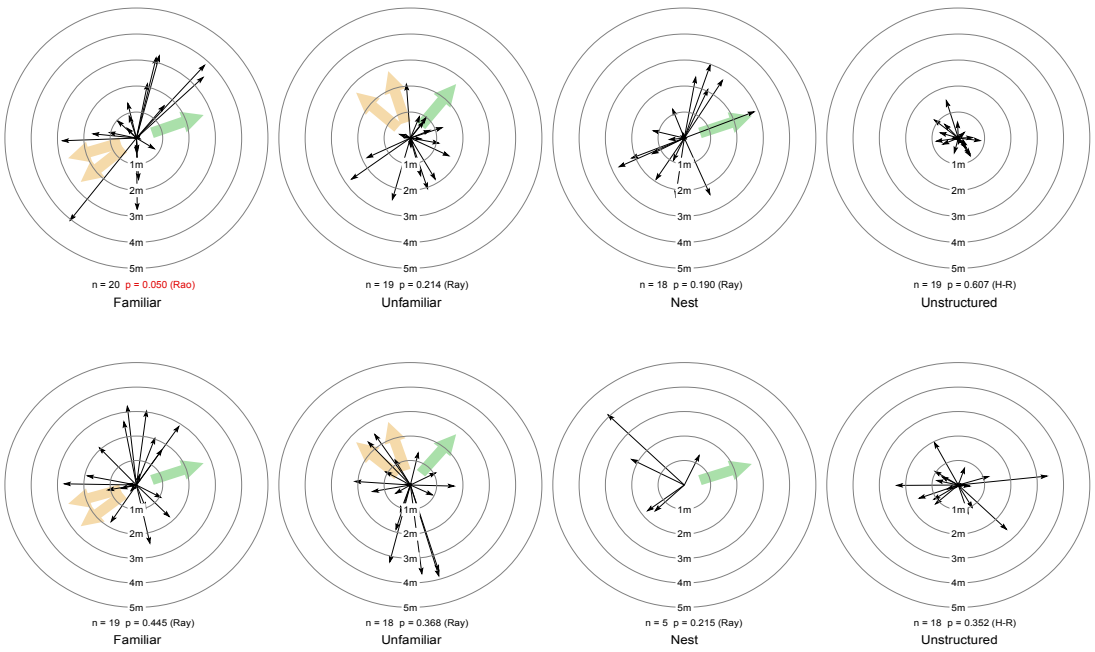


*Fig. 46: Bearing 30 seconds after release*

*The figure shows the bearing and distance of the animals 30 seconds after their release in the evening (top row) and in the morning (bottom row) experiments. The pale yellow arrows indicate the directions of the nest entrances, the pale green arrow points towards the foraging tree. The smallest p-value for circular uniformity test is noted, with the test that supplied it, Rayleigh (Ray), Rao (Rao), Hermans-Rasson (H-R). Statistically significant results are shown in red.*

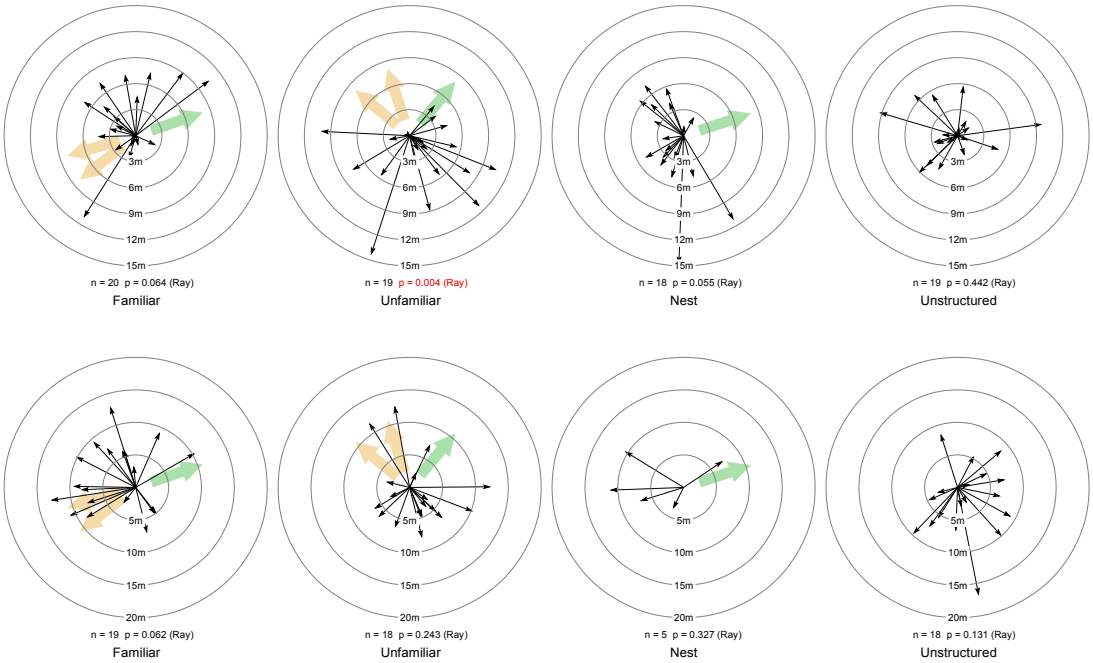


*Fig. 47: Bearing 2 mins after release*

*The black arrows indicate the bearing and distance from the release point for each animal 2 minutes after release in the evening (top) and the morning (bottom) experiments. The pale yellow arrows indicate the directions of the nest entrances, the pale green arrow points towards the foraging tree. The smallest p-value for circular uniformity test is noted, with the test that supplied it, Rayleigh (Ray), Rao (Rao), Hermans-Rasson (H-R). Statistically significant results are shown in red.*

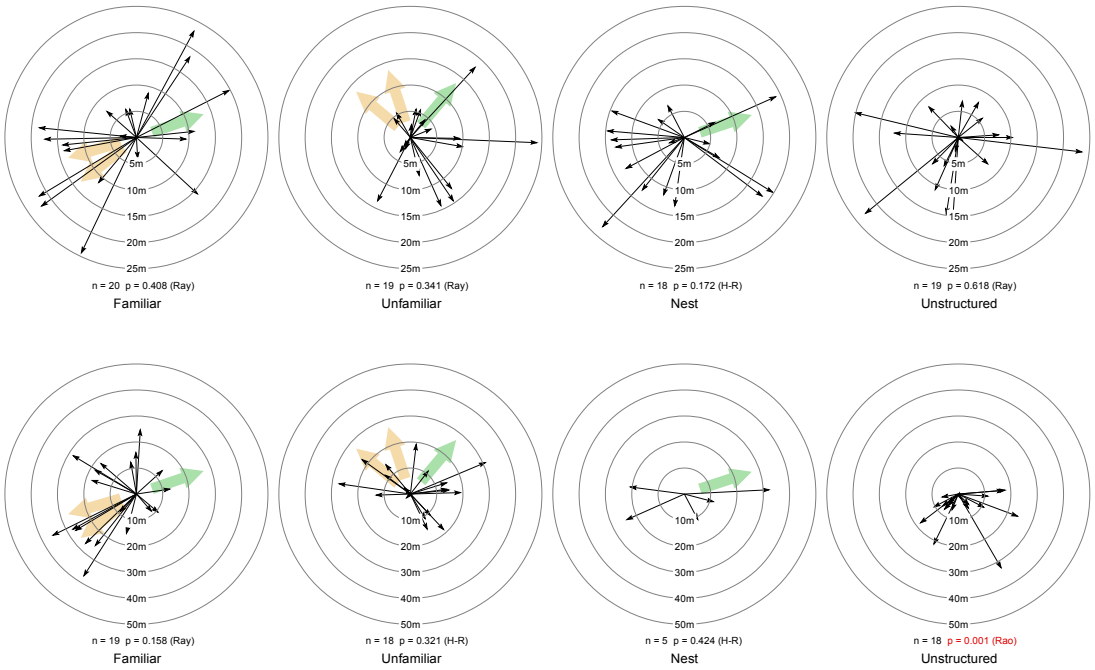


*Fig. 48: Bearing 5 mins after release*

*The black arrows indicate the bearing and distances from the release point of the animals 5 minutes after the release in the evening (top) and morning (bottom) experiments. The pale yellow arrows indicate the directions of the nest entrances, the pale green arrow points towards the foraging tree. The smallest p-value for circular uniformity test is noted, with the test that supplied it, Rayleigh (Ray), Rao (Rao), Hermans-Rasson (H-R). Statistically significant results are shown in red.*

The circular distribution statistical tests could not detect consistent orientation at either scene. It is likely that the three cases where statistically significant deviation from the uniform distribution was detected are just anomalies, especially if we consider the fact that in the morning experiment 5 minutes after release the unstructured view, where the animals have no visual clues whatsoever, have generated a non-uniform data set.

However, the Antarium experiments differ from the outdoor experiments in that in the Antarium the view is rotated several times while a scene is projected. Each rotation abruptly changes the view and, if the animal was navigating visually, forces her to re-align herself with the scene. In effect, each rotation can be treated as an independent re-release of the animal. Thus, it makes sense to analyse the bearing and distance for each rotation separately. That results in a significantly larger number of samples that can then be statistically tested.

Here we need to validate the assertion that treating the data after the abrupt scene change as an independent sample is justified. Obviously, it is the same animal therefore if the animal shows a behavioural bias, for example always moving towards the nest, the samples are not independent.

One might argue that doing so would be similar to have very few sample points, not enough for a meaningful statistical analysis and then repeating each point the same number of times to have a sufficient number of samples. We know that we can't run a test on the samples { 1, 2 and 3 } so we run a test on { 1, 1, 1, 1, 2, 2, 2, 2, 3, 3, 3, 3 } instead, claiming that 12 samples are enough for a statistical analysis. That is obviously incorrect. However, that is not what we are doing.

With the ants we do know from the re-release experiments described on page 152 that the initial heading of the same animal varies widely between releases. The Antarium projects a static image after each release, the animal receives no parallax clues to re-orient itself. Thus, that statistical uncertainty will not diminish as the animal moves on the ball. That means that even though we use the same animal, the animal is forced to pick a direction, and we do know that that direction will have a randomness similar to the randomness found between different animals. I will be cheating, but I will not be cheating much.

In addition, the analysis that will be performed by treating releases as independent samples is just one of several analysis methods and the other methods will not rely on treating rotations and independent releases of animals.

In light of the above, this is what was done:

After a rotation the first 5 seconds were ignored, to allow the animal to assess the panorama and decide on her heading. Then from the point where the animal was 5 seconds after the scene rotation to the next rotation her bearing and distance were calculated and plotted in *Fig. 49* the same way as before.

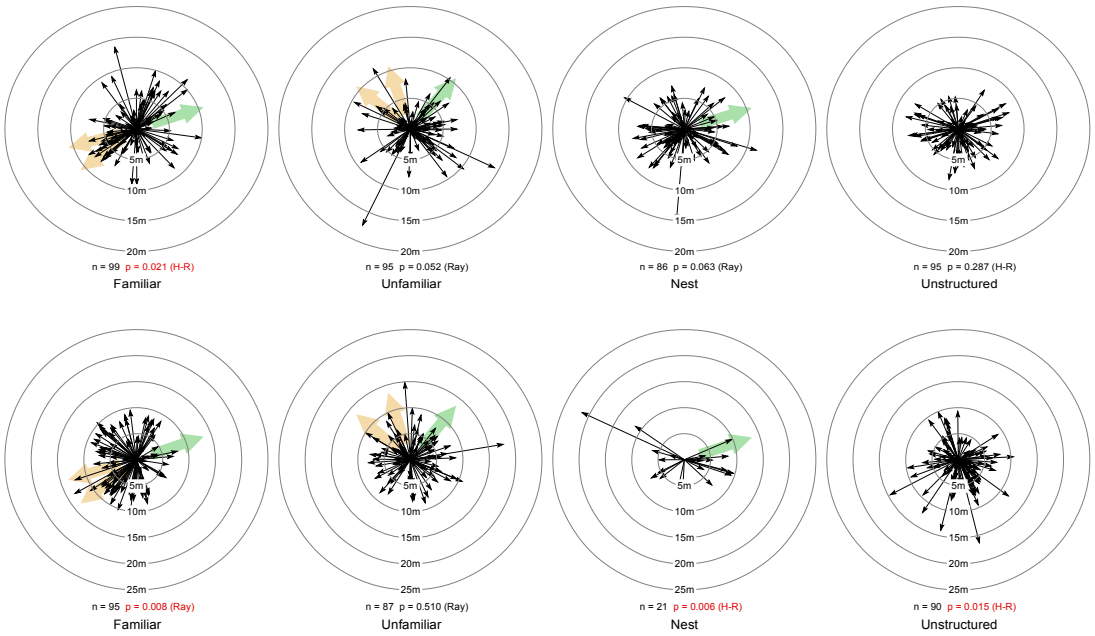


*Fig. 49: Bearings between rotations*

*The black arrows indicate the bearings and distances from 5 seconds after a rotation to the next rotation, for the evening (top) and morning (bottom) experiments. For further details see the caption for Fig. 46.*

Non-uniform distribution was detected at the familiar scene in both the evening and morning experiments but not at the unfamiliar scene. That matches the prediction that at the familiar scene the animals should show orientation while at the unfamiliar scene they can perform a random search, without any orientation.

The statistically significant result at the nest scene in the morning experiment is not indicative, because only 5 out of 19 animals were tested with that view. However, the lack of detectable deviation at the evening experiment matches the prediction of the animals performing a random search at the nest scene.

A statistically significant deviation from uniform is present at the unstructured view where the animals see a featureless image, but only in the morning experiment. That anomaly remains unexplained.

***Q2. Do the animals react to rotations?***
***Q3 Do the animals differentiate between scenes?***

To answer these questions the animals' responses for each rotation were analysed both manually and algorithmically.

The manual assessment was performed by looking at the ball and scene tracks at each rotation point. The responses to rotations were classified into 4 categories:

1. If the animal turned on the ball the same way as the scene, that is, her path remained straight in scene coordinates, she *followed* the rotation.
2. If the animal maintained her heading on the ball, resulting in a 90° turn in scene coordinates, she *ignored* the rotation.
3. If the animal turned on the ball in the direction opposite of the scene rotation, resulting in an about-face in scene coordinates, she is said to *counter* the rotation.
4. If the animal turned back on the ball, resulting in a 90° turn in scene coordinates, but in the direction opposite to what would result from ignoring the rotation, then she is said to *reverse*.

The categories are shown in *Fig. 50* for a clock-wise scene rotation.

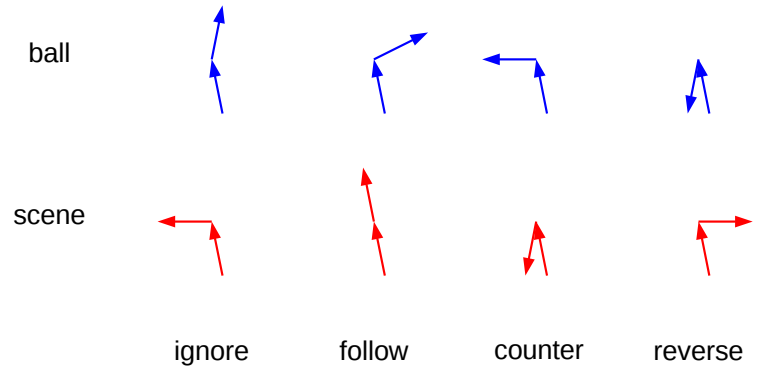


*Fig. 50: Rotation responses, clockwise rotation*

*The arrow pairs depict the rotation response scenarios. In each arrow pair the first arrow is the animal's walking direction before the rotation, the second arrow is her direction after the rotation. The scene was rotated clockwise around her. The blue arrows show her behaviour on the ball (that is, in real coordinates) while the red arrows are the ball coordinates mapped to the scene, that is, to the virtual world.*

The manually classified rotation responses were tallied for each of the four scenes. The tallied results form a contingency table for each scene. The ants and experiment days were pooled, but the evening and morning experiments were kept separate. The tables and the statistical results are shown in *Table 7*. Note that in the morning experiments very few animals were shown the nest scene.

| Evening | | | | | | | | Morning | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| occurrence | | | | loc | p-values | | | occurrence | | | | loc | p-values | | |
| I | F | C | R | | unf | nst | uns | I | F | C | R | | unf | nst | uns |
| 26 | 36 | 11 | 6 | fam | 0.437 | 0.760 | 0.000 | 18 | 47 | 7 | 4 | fam | 0.061 | 0.214 | 0.000 |
| 31 | 25 | 13 | 4 | unf | | 0.099 | 0.000 | 29 | 29 | 5 | 6 | unf | | 0.353 | 0.000 |
| 19 | 33 | 8 | 8 | nst | | | 0.000 | 4 | 7 | 2 | 3 | nst | | | 0.000 |
| 74 | 4 | 1 | 1 | uns | | | | 74 | 0 | 0 | 2 | uns | | | |

*Table 7: Contingency tables and dependency test results of rotation responses.*
*The columns I, F, C and R represent the counts of the Ignore, Follow, Counter and Reverse responses, respectively, to scene rotations. The rows represent the scenes, fam, unf, nst and uns standing for familiar, unfamiliar, nest and unstructured. The p-values occupy the upper triangle of a square matrix and show the result of the Fisher's exact test for count data of the contingency tables belonging to the scenes identified by the matrix row and column. The nest results for the morning experiment are marked with a grey background because only 5 animals were shown that scene.*

The evening and morning tables were also compared against each other, but no statistically significant difference was found.

The manual rotation response analysis has shown that the animals differentiated between the natural scenes and the unstructured view, but the limited number of available sample points could not statistically prove differentiation between the three natural scenes.

The automated rotational response analysis wants to answer the same question, but without the possible bias of the manual assessment.

For each rotation two vectors are defined, as shown in *Fig 51*. One connects the point on the animal's path, in ball coordinates, where the animal was 15 seconds before the rotation to the point on the track at the moment of the rotation. Then 5 seconds after the rotation is ignored, to give time to the animal to react to the scene change. After that the two end points of another 15 second long path segment define the second vector[16].

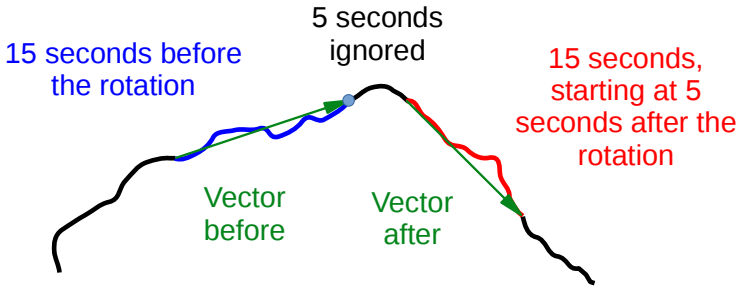


*Fig. 51: Automated rotation response extraction*

The angular difference between the two vectors is the animal's response angle. The response error is defined as the absolute value of the difference between the response angle and the scene rotation angle.

The value is 0 if the animal perfectly follows the rotation, 90° if she ignored the rotation or if she reversed and 180° if she counter-rotated.

16 *To make sure that the particular choice of values did not affect the result, the vector length time was varied between 10 and 30 seconds and the after-turn delay between 0 and 10 seconds and the results were compared to the 15s/5s case. It was verified that the statistical parameters were not significantly dependent on the particular choice of those time constants.*

For each scene the rotation error values were pooled for all experiment days and ants while keeping the morning and evening experiments separate. The pooled error values for each scene were treated as samples from a random process with an unknown distribution function. These sample sets were then pair-wise compared to see whether the unknown distributions are similar or different. Although the error values are angles, they are not really circular data. Therefore, non-circular tests can be used. The low number of samples in each set warrants the use of non-parametric tests (see Statistical methods on page 167).

As *Table 8* shows, the responses evoked by the rotations at the unstructured scene are different from any other scene. Furthermore, in the morning experiment the response difference to the familiar and unfamiliar views is also prominent.

**Evening**

| | fam | unf | nst | uns |
|---|---|---|---|---|
| fam | - | 0.052 | 0.715 | 0.000 |
| unf | 0.146 | - | 0.041 | 0.037 |
| nst | 0.302 | 0.043 | - | 0.000 |
| uns | 0.000 | 0.003 | 0.000 | - |

**Morning**

| | fam | unf | nst | uns |
|---|---|---|---|---|
| fam | - | 0.010 | 0.589 | 0.000 |
| unf | 0.025 | - | 0.298 | 0.001 |
| nst | 0.929 | 0.596 | - | 0.002 |
| uns | 0.000 | 0.000 | 0.000 | - |

*Table 8: Statistical test results for rotational responses extracted automatically The two matrices contain the p-values for the MW test (upper triangle) and the KS test (lower triangle) when the extracted rotational error values for the matrix row and column scenes are compared against each other. Red figures indicate statistically significant difference after Bonferroni correction. Blue figures indicate results that were statistically significant by the individual test, but not after the Bonferroni correction. The grey background marks results involving the nest location in the morning, where the number of data points was very low. The scenes are fam: familiar, unf: unfamiliar, nst: nest, uns: unstructured.*

The distribution of the rotation errors for the evening and morning experiments are displayed in *Fig. 52* as violin plots[17]

17 *A violin plot is a box plot that also incorporates the estimated density function. The R package that generates violin plots, together with its detailed documentation, can be found at* https://github.com/TomKellyGenetics/vioplot

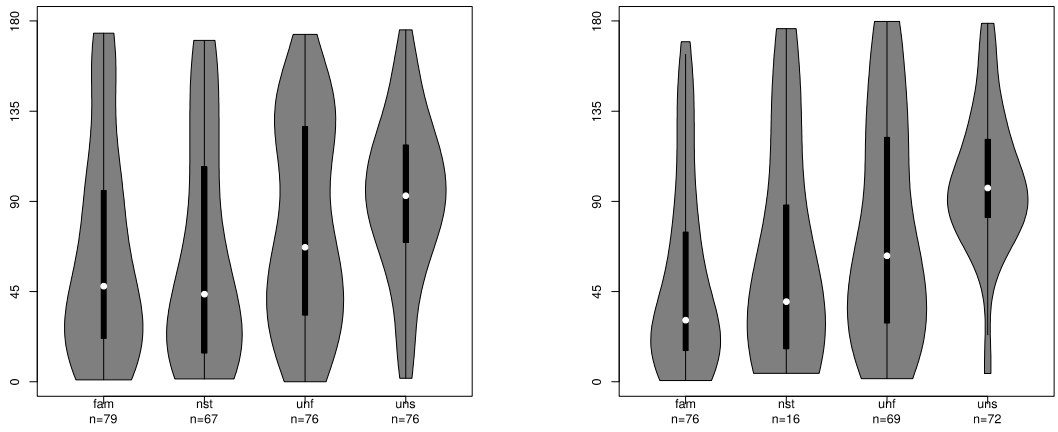


*Fig. 52: Violin plots of rotational errors, evening (left) and morning (right)*

*The white dot is the median. The black bar is the standard deviation, centred on the mean. The width of the grey area indicates the estimated probability density kernel. The vertical axis is the error in degrees. The labels fam, unf, nst and uns denote the familiar, unfamiliar, nest and unstructured scenes, respectively.*

The distributions show that at the familiar scene the animals mostly followed the rotations: the median error is relatively low and the distribution is widest at an even lower angular error. The nest scene distribution is similar. The distribution of errors in response to rotations of the unstructured scene has a median close to 90 degrees with a prominent widening of the plot around that value, indicating that the animals mostly ignored the rotations. The unfamiliar scene has the largest spread of responses, the distribution kernel is much more elongated than in any other case and the median error is higher than either for the nest or the familiar view. That indicates that the ants' responses to scene rotations at the unfamiliar scene are the least predictable.

We can also check the animal's behaviour between the rotations. Halfway between the actual scene rotation points nothing changes, the ant is shown a completely static image. However, that can be taken as a scene rotation by 0 degrees. Any rotational error thus indicates that the animal was turning on her own volition, without any external trigger.

The violin plots for the evening and morning experiments are shown in *Fig. 54* and *Fig. 53*:

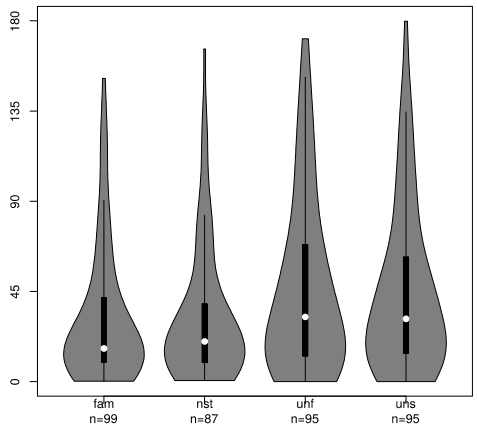


*Fig. 54: 0-degree rotation response errors, evening*

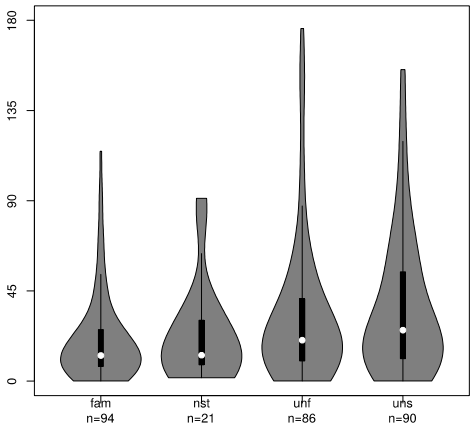


*Fig. 53: 0-degree rotation response errors, morning*

As expected, the rotation errors are low. However, the probability that the animal was on a curved path (i.e. she had a non-zero error) is highest at the unstructured scene, where she had no visual reference at all. It is also noticeable that the morning distributions show clearer differences between the four scenes; for each scene the median is lower, the standard deviation is smaller and the extent of the density function is shorter.

The statistical comparison of these distributions are shown in *Table 9*:

**Evening**

| | fam | unf | nst | uns |
|---|---|---|---|---|
| fam | - | 0.078 | 0.985 | 0.021 |
| unf | 0.082 | - | 0.080 | 0.750 |
| nst | 0.864 | 0.033 | - | 0.025 |
| uns | 0.029 | 0.959 | 0.018 | - |

**Morning**

| | fam | unf | nst | uns |
|---|---|---|---|---|
| fam | - | 0.014 | 0.808 | 0.001 |
| unf | 0.033 | - | 0.200 | 0.287 |
| nst | 0.963 | 0.377 | - | 0.065 |
| uns | 0.000 | 0.309 | 0.175 | - |

*Table 9: Statistical comparisons of 0-degree rotation response errors*

*The two matrices contain the p-values for the MW test (upper triangle) and the KS test (lower triangle) when the extracted rotational error values for the scenes of the matrix row and column are compared against each other. Red figures indicate statistically significant difference after Bonferroni correction. Blue figures indicate results that were statistically significant by the individual test, but not after the Bonferroni correction. The grey background marks results with very few samples. The scenes are fam: familiar, unf: unfamiliar, nst: nest, uns: unstructured.*

According to the table, even when nothing changes on the projected image there is a statistically significant difference between seeing the familiar scene

and an unstructured visual and also between the nest and the unstructured views (where enough data are available). No statistical difference was found between the unfamiliar scene and the unstructured image, however. As it was the case with the actual rotation responses, the familiar and unfamiliar view evoked measurably different behaviour in the morning, but not in the evening.

In conclusion, the automated analysis of the ants' responses to scene rotations, just like the manual one, found a marked behavioural difference between the animal seeing the unstructured view and when confronted with any of the natural scenes. In addition, differentiation between seeing the familiar and unfamiliar scenes was evident in the morning experiments but not in the evening ones. Furthermore, in the evening the animals responded statistically significantly differently to scene rotations of the nest and unfamiliar scenes. In the morning there were not enough samples from the nest scene to perform a meaningful statistical comparison.

***Q4: Do the animals behave at the various scenes as expected?***
***Q5: Are there significant differences between animals?***

Both the bearing analysis and the rotational response analysis are blunt instruments. The former reduces the path between rotations to a single vector. The latter reduces the path around a rotation into a single angular difference. A more detailed picture can be established with a more fine-grained approach.

*Heading density*

The path can be turned into a series of equal length vectors, as described in *Path to vector sequence transformation* on page 164. I chose 2.5cm[18] (roughly the animals' body length) for the vectors. The direction of each vector is the average heading of the animal over that 2.5cm displacement.

That process results in a large number of heading samples. They are sorted into 5° wide bins and an empirical probability density curve is then calculated using a Gaussian smoothing kernel.

18 *The 2.5cm is an arbitrary choice, but it was checked that varying the vector length between 1cm and 10cm has no significant effect on the end result.*

If the animal had a preferred direction, that will show up as a peak on the density curve. Examples of such curves are shown on *Fig. 55*.

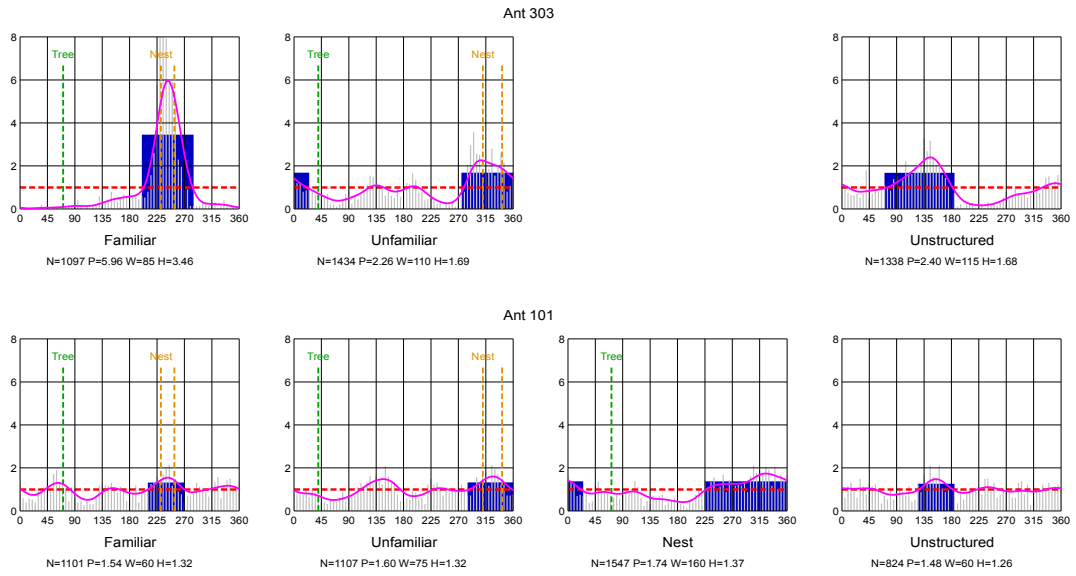


*Fig. 55: Probability densities of path vectors*

*The data belong to ant 303 (top) and ant 101 (bottom), evening experiment. For 303 the nest scene was not shown.*

*The animal's path is cut into 2.5cm long vectors. The number of vectors forming the path is given by the N value. The directions of the individual vectors are calculated then sorted into 5° wide bins. The bin counts are represented by the grey vertical lines. The X axis is the direction angle. The Y axis is the relative occurrence compared to a uniform distribution (marked by the dashed red horizontal line). The magenta line is the empirical probability density function, generated by convolving the bins with a Gaussian kernel.*

*The highest peak of the density function is located, then the points where the curve crosses the uniform (y=1) line, left and right of the peak, are determined. The blue rectangle spans that range. Its height is the mean of the curve within the range. The value of the density curve at the peak is given by P, the height of the rectangle by H and and its width, in degrees, by W.*

*The orange dashed vertical lines indicate the direction of the two nest entrances while the green line is the direction of the foraging tree.*

The probability density function at the familiar scene peaks roughly in the direction of the nest for both ant 303 and 101. However, while for ant 303 the peak is over 6 times the uniform distribution, for ant 101 it barely reaches 1.3 times the uniform. In fact, in case of ant 101, the P and H values are more or less the same for all scenes.

A detailed look at the behaviour of the two animals during the experiment can be seen on *Fig 56*.

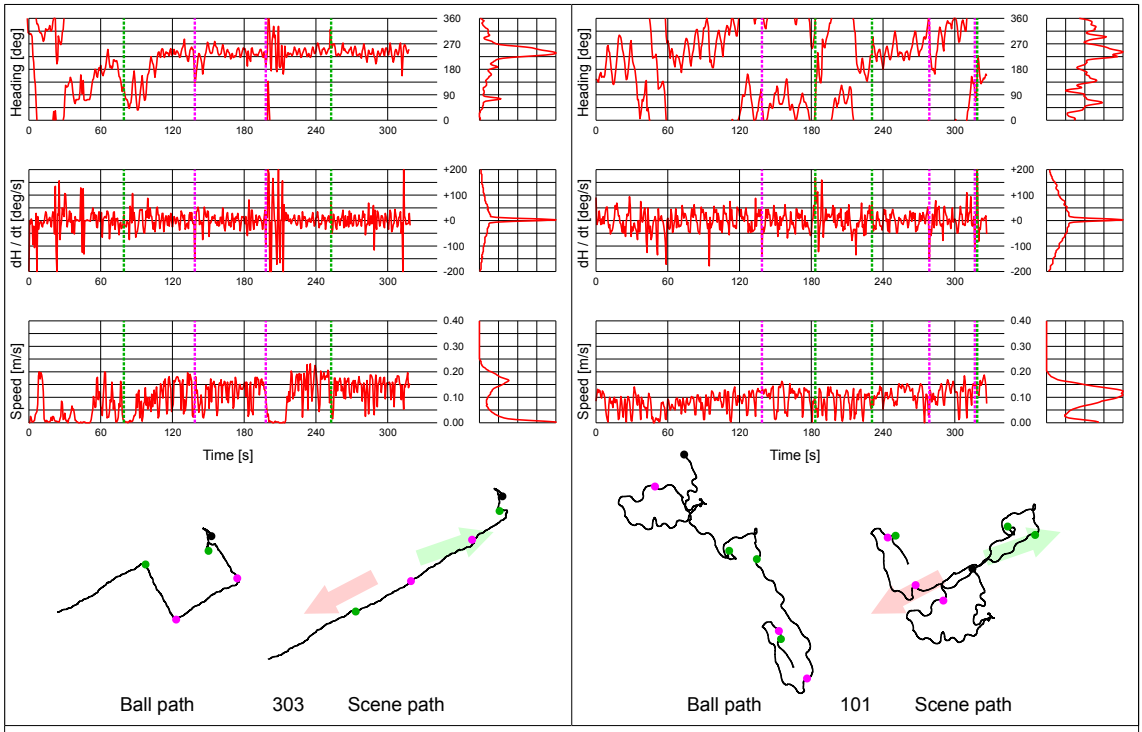


*Fig. 56: Path details of selected animals*

*The figure shows the path details of ants 303 (left) and 101 (right) at the familiar scene during the evening experiment. To the right of each time course is its probability density curve. The green and magenta vertical dashed lines on the time courses and the same colour dots on the path plots represent 90° scene rotations, magenta is clockwise and green is anticlockwise. The black dot on the path is the starting point. On the scene path plot the pale red arrow points to between the two nest entrances and the pale green arrow points to the foraging tree.*

Ant 303 moved very slowly between the start of the experiment and the first rotation. Her momentary heading was noisy and ill defined. After the first rotation, however, she quickly aligned herself with the scene and started to move in a nest-ward direction. She followed all subsequent rotations and maintained a constant heading relative to the scene (see the ball and scene tracks below the time course plots).

The speed curve shows that she followed the first clockwise rotation at 140s without stopping and without hesitation. At the second clockwise rotation at 200s, however, she stopped and scanned her environment while rotating on the

spot (her speed was close to zero but her angular velocity was high and her momentary heading oscillated widely). She spent about 15 seconds doing that and then she resumed walking, maintaining the nest-ward heading in scene coordinates. At the last (anticlockwise) rotation at 255s she again followed the rotation without stopping to realign herself.

Ant 101 did not follow the rotations. She did not maintain a constant heading, her path is characterised by frequent changes of direction. Between the first (clockwise) rotation at 140s and the subsequent anticlockwise rotation at 205s she maintains a wobbly but roughly tree-directed path and between the anticlockwise rotation at 225s and the clockwise rotation at 270s she keeps a more or less nest-ward heading, but between any other pairs of rotations her path shows no orientation at all.

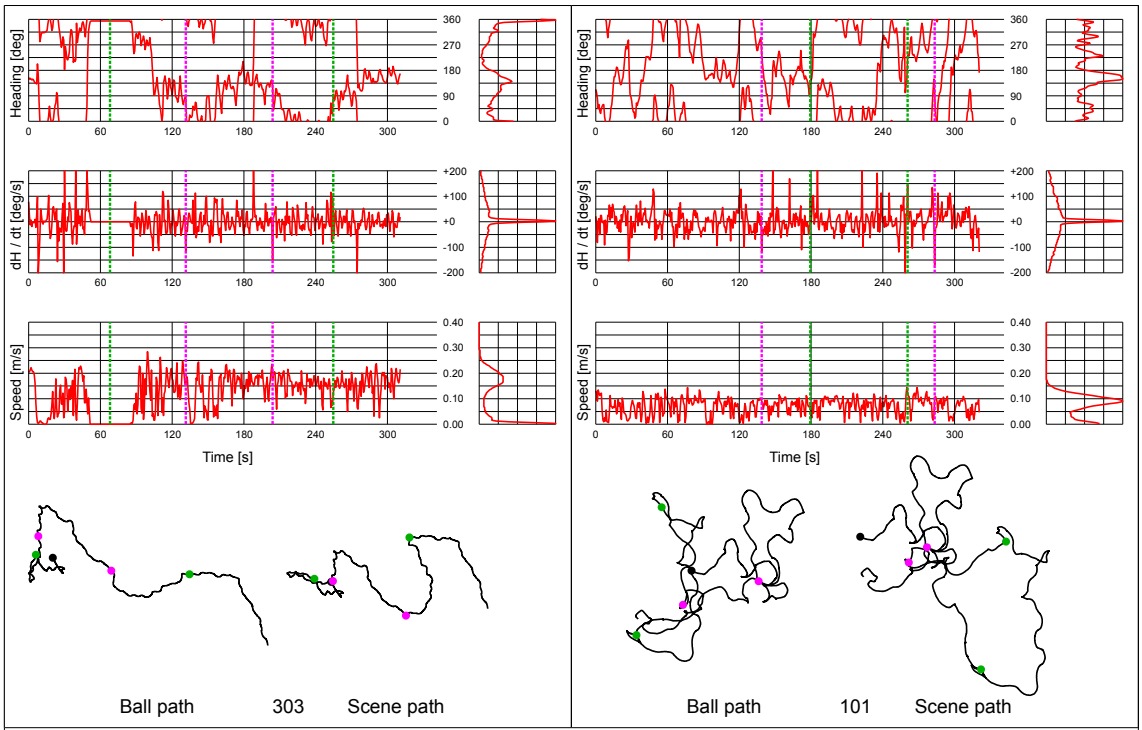


*Fig. 57: Paths of ant 303 (left) and 101 (right), unstructured view, evening experiment*

For comparison, *Fig. 57* shows the paths of the same animals when confronted with the unstructured view. Ant 101's movements can best be described as aimless wandering. Ant 303 is different. She does not follow scene rotations, which is not surprising as the scene provides no visual cues. However, after the

second rotation her ball path consists mostly of almost straight or gently curving segments.

The examples of ants 101 and 303 have already answered Question 5, “Are there significant differences between the animals?” in the affirmative. In fact, questions 1 (do the animals show orientation?), 2 (do the animals react to scene rotations?) and 3 (do the animals differentiate between the four scenes?) were answered by 303 as “yes, yes and yes” and by 101 as “no, no and maybe”.

Naturally, two animals hand-picked from a cohort of 20 is not sufficient evidence. For a more robust analysis we need to establish measures of behaviour when the animals are definitely navigating and when they definitely are not. To do that, we need reference data for the two cases.

The unstructured view can be used as reference for the animals not navigating, due to the complete lack of visual cues in that scene. It was mentioned in the *Methods* section that after the morning experiment the animals were released outdoors at a point which they were familiar with. Their tracks to the nest were marked with small flags and subsequently mapped with differential GPS. Those ants went to their nest in the real world, thus we can use their paths as references for definitely navigating animals.

The GPS samples have a much lower spatial resolution than the trackball and cannot reveal how long it took the animal to walk from one sample to the next. Since we are only interested in the location data, the lack of time (and thus speed) information is of no consequence. The resolution problem was overcome by constructing a high-resolution path from the GPS sample points using Catmull-Rom splines, a piece-wise cubic interpolation method. Spline interpolation can not re-create the small-scale details of the animals’ movement (e.g. periodic wobble, stops for scanning, etc.) but it generates an approximation of the path with sufficient detail for our purposes. In fact, only the vectorised form of the interpolated path was used to generate the directional probability density curves and the vectorisation process would eliminate a lot of the missing details anyway.

*Fig. 58* shows the reconstructed paths on the satellite image. Of the 28 animals 3 went to the foraging tree, 5 to Nest C and 20 to Nest N.


*Fig. 58: GPS tracks of the animals released outdoors*

The animals were released about 3.5 metres from the Nest N entrance. It is obvious from the paths that their initial direction varies widely. This has already been reported in (Zeil et al., 2014b). It was also shown by the repeated release experiment, described on page 152.

The bearings of paths after one metre of travel span a roughly 90° wide direction cone. The animals then turn to a more nest-ward heading but the routes they take indicate that they often correct their path. In addition, many animals miss and/or overshoot the nest and need to loop back. The miss can be as large as a metre, which for a 3.5m target distance corresponds to a 16° systemic heading error.

Therefore, the fact that the animal navigates does not mean that she goes from A to B on a straight line. Far from it, she can take a quite indirect path. Furthermore, the initial heading of an animal can be off by ±45° and she will

correct her heading as she walks, presumably by constantly re-evaluating the changing visual panorama.

It should be noted that in the current experiment the Antarium projected a static image. No matter how much the animals walked the view did not change, bar the scene rotations. The animals in the Antarium were deprived of the translation induced scene changes that could have helped them to correct their heading.

The wide variety of behaviours of actually navigating ants is demonstrated in Fig. 59 that shows the vector direction probability densities for the outdoor paths of selected animals.

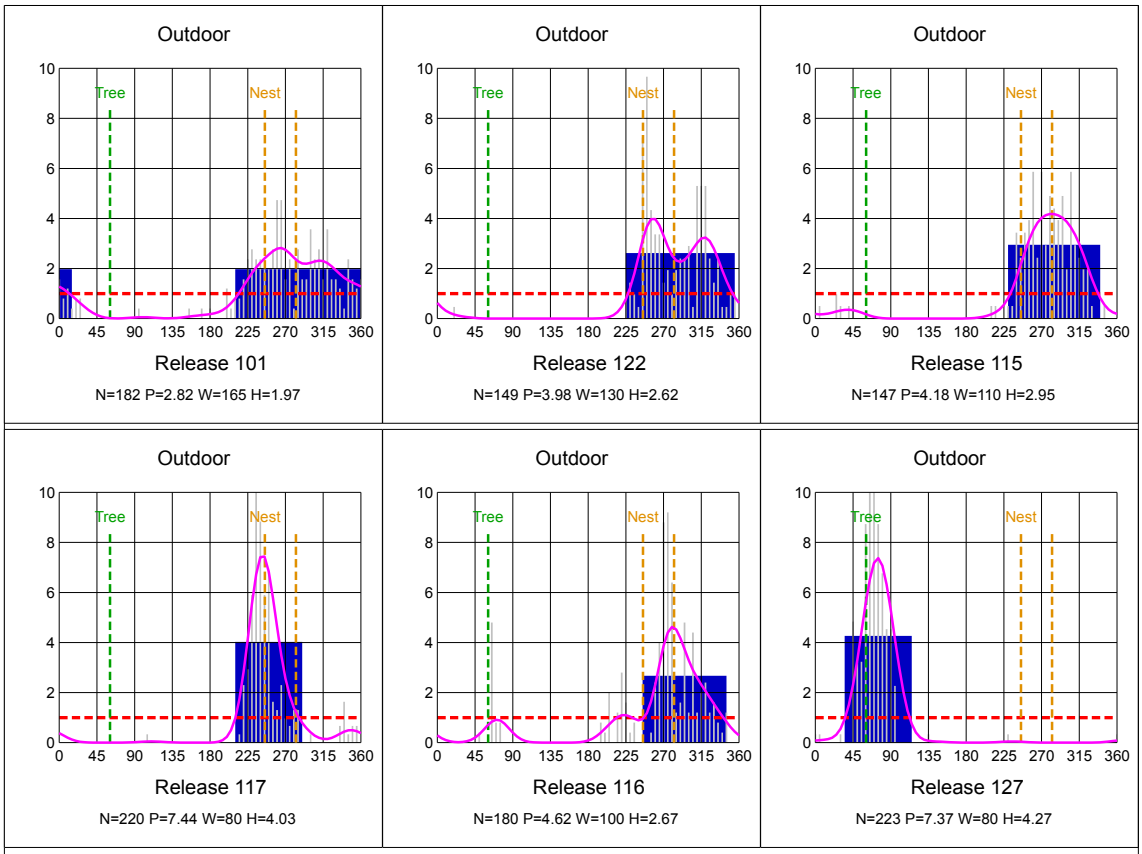


*Fig. 59: Example vector direction probability density curves, outdoor releases*

*The animal's path is cut into 2.5cm long vectors. The number of vectors forming the path is given by the N value. The directions of the individual vectors are calculated then sorted into 5° wide bins. The bin counts are represented by the grey vertical lines. The X axis is the direction angle. The Y axis is the relative occurrence compared to a uniform distribution (marked by the dashed red horizontal line). The magenta line is the empirical probability density function, generated by convolving the bins with a Gaussian kernel.*

*The highest peak of the density function is located, then the points where the curve crosses the uniform (y=1) line, left and right of the peak, are determined. The blue rectangle spans that range. Its height is the mean of the curve within the range. The value of the density curve at the peak is given by P, the height of the rectangle by H and and its width, in degrees, by W.*

*The orange dashed vertical lines indicate the direction of the two nest entrances while the green line is the direction of the foraging tree.*

We can calculate the minimum, maximum, mean and the standard deviation for the peak (P), mean (H) and angular spread (W) values over the entire cohort to get the ranges for navigating animals. We can perform the same calculation for

the unstructured view for the non-navigating animals as well. The results are shown in *Table 10*:

| | Outdoor release | | | | Unstructured view | | | |
|---|---|---|---|---|---|---|---|---|
| | Mean | Sdev | Min | Max | Mean | Sdev | Min | Max |
| **Curve peak** | 4.96 | 1.35 | 2.71 | 7.43 | 2.05 | 0.44 | 1.52 | 3.23 |
| **Box height** | 2.87 | 0.65 | 1.59 | 4.27 | 1.52 | 0.21 | 1.25 | 1.97 |
| **Box width** | 113 | 27 | 80 | 170 | 135 | 46 | 65 | 225 |

*Table 10: Vector parameters for navigating and not navigating animals*

*Fig. 60* shows the distribution of the peak values for the unstructured views (morning and evening data pooled) and from the outdoor release paths. A normal distribution was fitted to the data points.

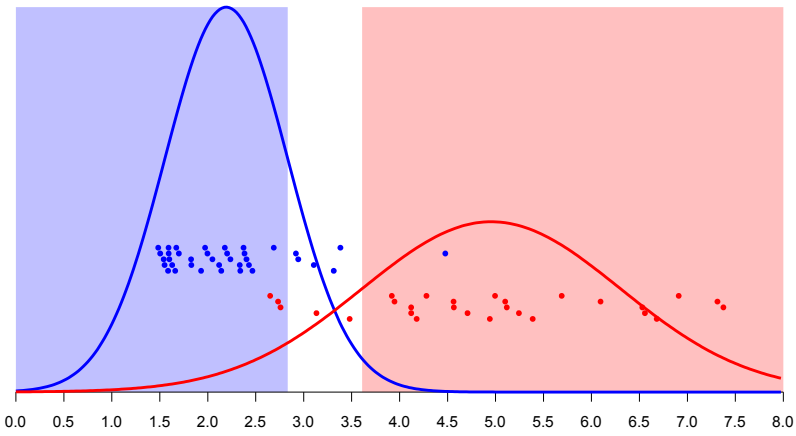


*Fig. 60: Distribution of the peaks of the vector statistics*

*The figure shows the peak values of the path vector density curves. Blue is the unstructured view (evening and morning pooled) and red is the outdoor release. The dots represent the peak values of the individual ants. The vertical spread of the dots bears no meaning, it is there only to visualise otherwise overlapping dots. The continuous curve is a normal distribution fitted on the dots. The filled areas (acceptance ranges) extend from one standard deviation from the mean in one direction to infinity (end of the graph) on the other.*

We can use this procedure for discriminating between navigating and non-navigating animals. We calculate the peak of the vector distribution of an animal's path in the Antarium. If the peak is less than the mean plus one standard deviation of the unstructured statistics, then we can conclude that the

animal was not navigating. If the peak is larger than the mean minus one standard deviation of the outdoor statistics then the animal was navigating. If the peak falls in between, then the animal is not classified either way. The shaded areas on *Fig. 60* show the decision ranges.

It is clear that the classification is not fool-proof. In *Fig. 60* there are outlier red dots on the blue background and outlier blue dots on the red one. Those represent animals that would be falsely classified as not navigating and navigating, respectively.

The probability of falsely declaring an animal a non-navigator is the blue area under the red curve while the probability of falsely classifying an animal a navigator is the red area under the blue curve, which is significantly smaller. Therefore, the classifier is more likely to classify a navigating animal as “non-navigator” than the other way around.

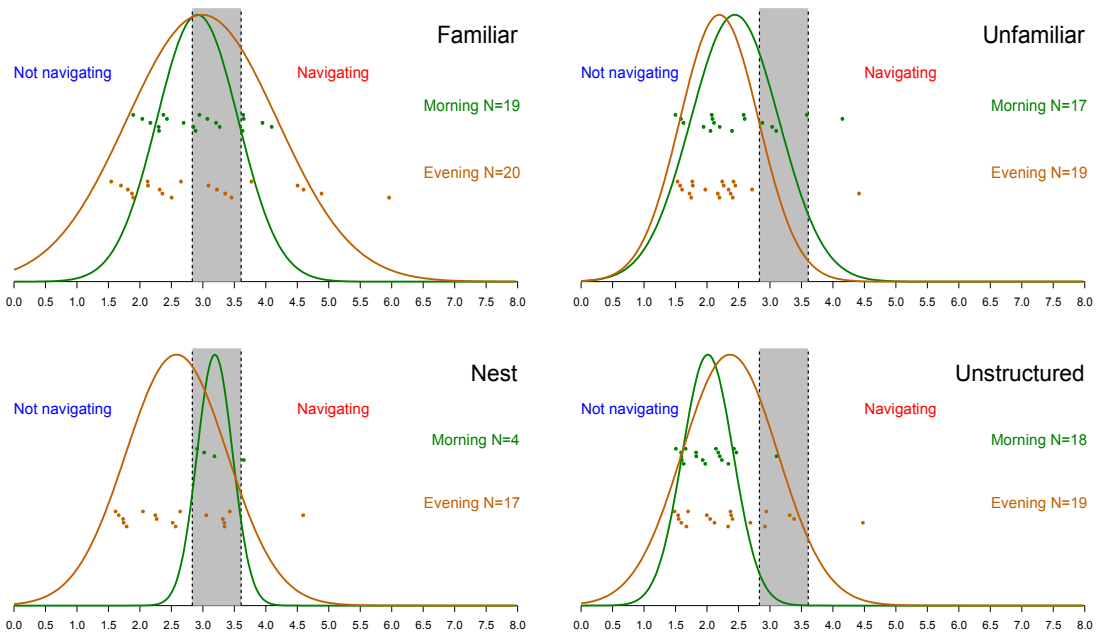


*Fig. 61: Vector direction probability density curve peak distributions*

*The figure shows the vector density curve peak values for each scene. Dots indicate the values of individual ants while the curve is a normal distribution fitted to those samples. The dashed vertical lines are the classifier's threshold values, calculated from the outdoor experiment results and the pooled unstructured scene results. N denotes the number of animals that were confronted with the given scene.*

*Fig. 61* shows the vector density curve peak values of the animals in the Antarium. The individual values for each ant and each scene are shown together with the classifier's thresholds.

The results of the classification of the vector direction density curve peaks are summarised in *Table 11*:

| | Evening | | | Morning | | |
|---|---|---|---|---|---|---|
| **Scene** | **Yes** | **No** | **?** | **Yes** | **No** | **?** |
| ***Outdoor reference*** | | | | *23* | *3* | *2* |
| **Familiar** | 5 | 11 | 4 | 5 | 8 | 6 |
| **Unfamiliar** | 1 | 18 | 0 | 1 | 12 | 4 |
| **Nest** | 1 | 11 | 5 | 1 | 0 | 3 |
| **Unstructured** | 1 | 14 | 4 | 0 | 17 | 1 |

*Table 11: Vector density curve based classification results*
*Animals were classified by the curve peak parameter derived from the direction density function of their path vector sequence. Yes is the number of animals that the algorithm found navigating, No is the number that it found not navigating and the question mark is the number of animals that it could not classify. Not all animals were shown all scenes so the numbers do not add up to the same total.*

It is *very important* to point out that in this context "navigating" means "showing orientation comparable to that of an animal going towards a target location under natural circumstances". It does not mean that she goes towards a particular target, or that she maintains a constant heading. The ant statistically preferences a certain direction, nothing more. Similarly, "not navigating" means that the animal has not shown significant preference to any direction, but it does not state anything about what she has done; let that be a random search, walking in a circle, or anything else without a highly dominant heading. In particular, at the Nest scene the animal's expected behaviour is to perform a random search for the nest entrance. Consequently, animals released outdoors near the nest would be classified as "non-navigating", despite being in their natural environment. Conversely, if an animal is placed at a location that she is not at all familiar with and she chooses to walk straight, she would be classified as a navigator even though she is very unlikely to have any usable navigational knowledge of her surroundings.

*Fig. 62* shows paths of animals that were classified as navigating.

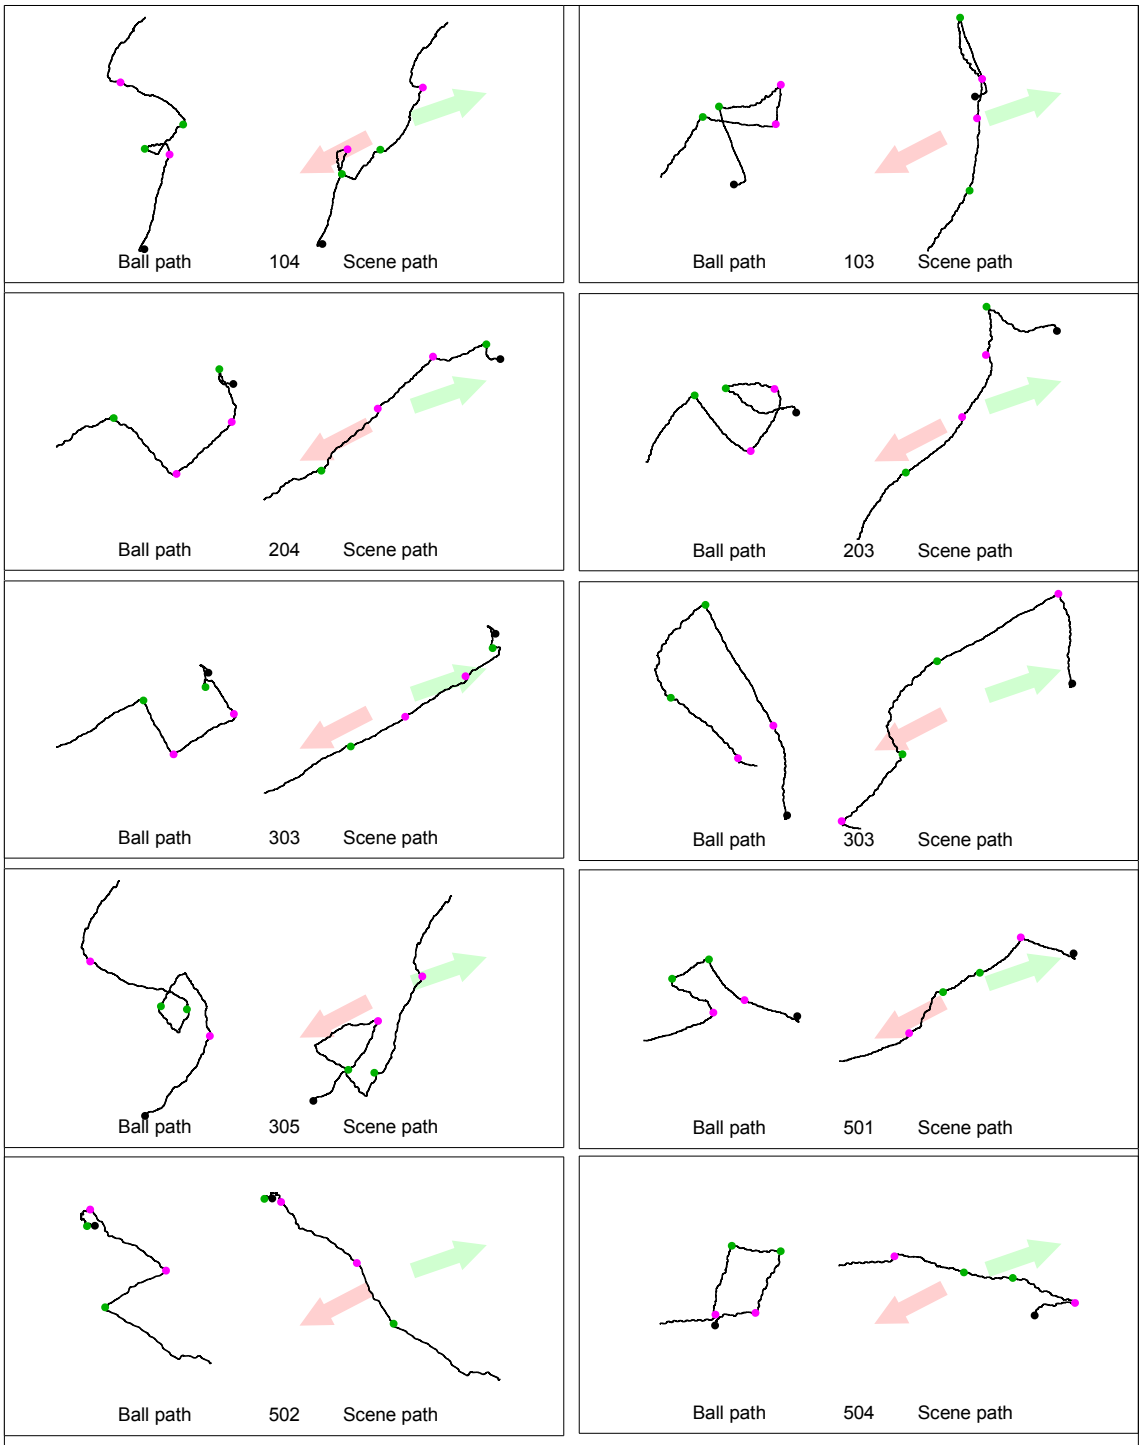


*Fig. 62: Paths of animals classified as navigating*

*The figure shows the paths of the animals which were classified as navigating by the density peak measure at the familiar scene, in the evening (left) and morning( right) experiments.*

For comparison, *Fig. 63* shows 8 tracks that were classified as not navigating at the familiar scene in the evening experiments.

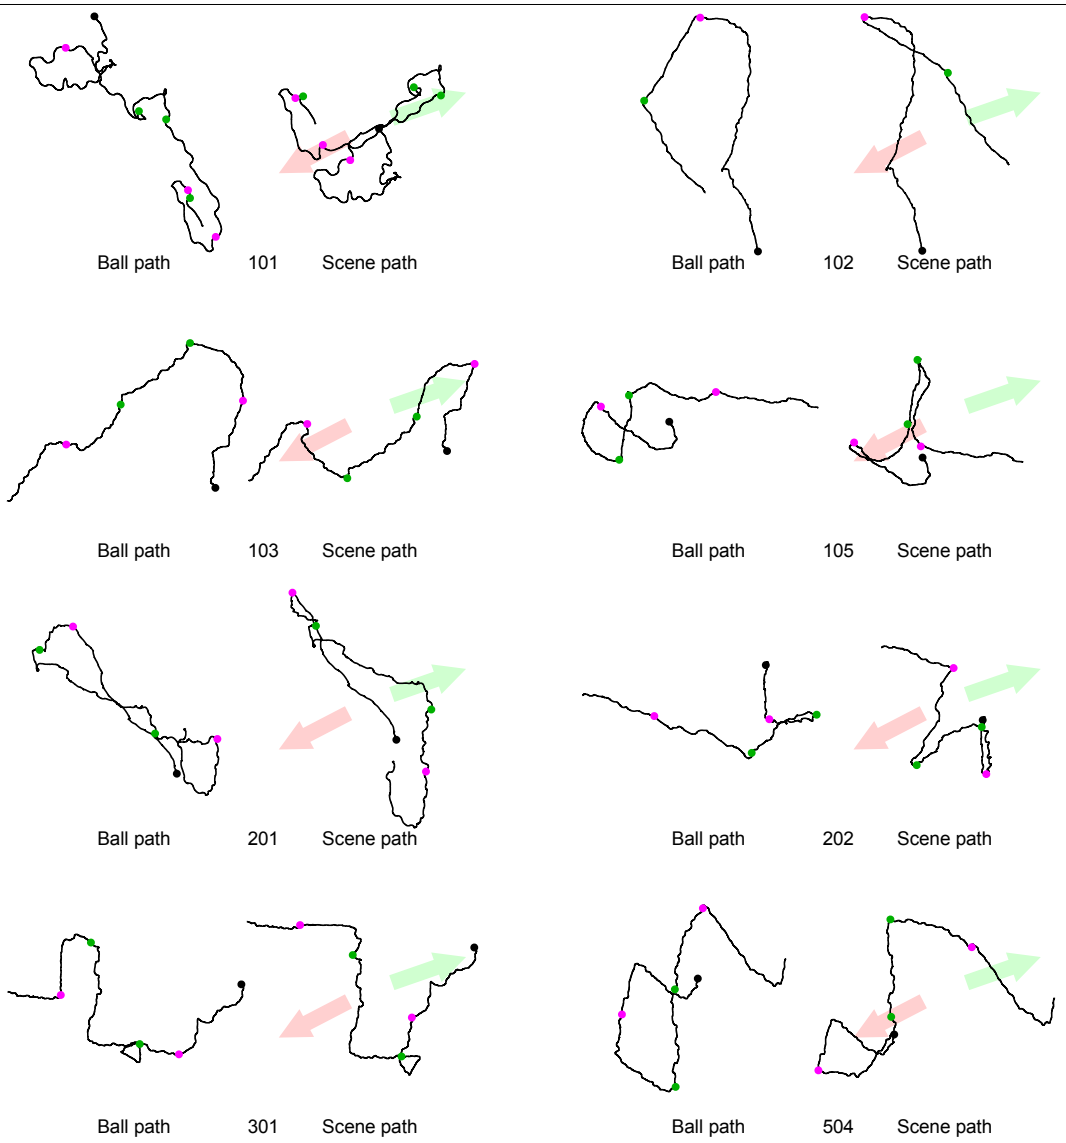


*Fig. 63: Paths of non-navigating animals*

*The figure shows eight of the evening experiment Antarium paths of animals that were classified as non-navigating.*

The paths shown in *Fig. 62* and *Fig. 63* demonstrate that the classification, though not perfect, is working. The navigating animals mostly followed rotations while the non-navigating ones more often then not ignored them. *Fig. 63* also demonstrates that even non-navigating animals show temporary orientation. Their paths are not Brownian motion, the animals maintain a heading for a while, then change direction. However, the fact that they keep their ball heading even when the scene rotates around them (see ant 103, for example) indicates

that they are not referencing the panorama, that is, they are indeed not navigating visually.

The answer to question 4 (Do animals behave at the different scenes as expected?) is not obvious. It is clear from *Table 11* on page 205 that at the familiar scene the animals were more likely to show orientation than at any other scene, which is the expected behaviour. On the other hand, there is no obvious difference between the unfamiliar, the nest and the unstructured scenes.

However, other test methods have already shown that the animals differentiate between the unstructured visual and the natural scenes (see *Table 7* on page 190 and *Table 8* on page 192). But no test method could tell the unfamiliar and the nest scenes apart. Considering that at both scenes the random search is expected behaviour, that is actually in line with the expectations.

Therefore, we can answer the question in the affirmative:

- The familiar scene evoked oriented behaviour
- The unfamiliar and nest scenes triggered random searches
- The paths of the animals confronted with the unstructured view are also characterised with random motion, but their large-scale statistics are different from that of the paths at the nest and unfamiliar scenes.

***Evening vs. morning experiments***

The average speed of the animals was calculated between rotation points and it was noticed that there was a difference between between the evening and morning experiments.

Table 12 shows that the speeds were fairly consistent between scenes. The roughly ±12% standard deviation indicates that there were no significant differences between individual animals either.

However, in the morning the animals moved approximately 20% faster at all four scenes. It should be noted that the experiments were performed in an air

conditioned laboratory, thus the speed difference between evening and morning cannot be explained by ambient temperature change.

| Scene | Evening | | | Morning | | |
|---|---|---|---|---|---|---|
| | N | mean | sd | N | mean | sd |
| familiar | 99 | 0.115 | 0.012 | 95 | 0.139 | 0.015 |
| unfamiliar | 95 | 0.104 | 0.012 | 87 | 0.131 | 0.014 |
| nest | 86 | 0.109 | 0.012 | 21 | 0.129 | 0.030 |
| unstructured | 95 | 0.110 | 0.012 | 90 | 0.134 | 0.015 |

*Table 12: Average speed between rotations*

*N is the number of segments, mean is the average speed of the animals during the segment [m/s]while sd is the standard deviation of the speed, also in [m/s]. The grey back ground indicates that the number of samples was much lower than in all other cases.*

Further experiments will be needed to confirm that the day-of-time dependent speed variation is consistently present and to find an explanation for it.

## Discussion

Animals were shown static panoramic images of natural scenes, a familiar location on the foraging corridor, the nest site and a place that they are not normally familiar with. In addition, an artificially generated featureless image was projected to them. At approximately one minute intervals the projected scene was rotated randomly clockwise or anti-clockwise by 90 degrees. The animal's path was recorded using a trackball. The path on the ball was transformed to a path in the projected scene, taking the rotations into account.

The data was then analysed using several methods: circular statistics of the animals' bearing at various time points; manual and algorithmic analysis of the animals' response to rotations; and a classification of the paths based on the behaviour of freely navigating animals in their natural environment.

Five questions were asked and answered, summarised here:

Do the animals show orientation?

The comparison with the freely navigating animals indicated that not all, but several animals have shown clear orientation at the familiar scene.

Do the animals react to scene rotations?

The rotation response analysis, both manual and automated, has shown that the animals reacted to scene rotations. It was most prevalent when confronted with the familiar scene, much less so at the nest and unfamiliar scenes and not at all the unstructured scene.

Do the animals differentiate between the scenes?

Differentiation between the unstructured view and the natural scenes was demonstrated by the rotation response analysis. Differentiation between the familiar scene and the other two natural scenes was shown by the path classification analysis: at the familiar scene the vector direction probability density curve peaks resulted in a distribution with significantly higher mean, and in the evening also larger standard deviation, than for the nest or unfamiliar scenes. The nest and unfamiliar scenes could not be separated.

Do the animals behave at the various scenes as expected?

Yes, the animal behaviour at the various scenes was consistent with the expectations: oriented movement at the familiar scene, movement consistent with a random search at the nest and unfamiliar scenes and random movement at the unstructured view.

Are there significant differences between animals?

The classification method has indicated that at the familiar scene approximately 20-25% of the animals were clearly oriented, about 40-50% clearly not oriented and the rest somewhere in between.

That is better than the results of the first experiment (19% navigating, 75% not navigating). The improvement is in line with expectations. Various shortcomings of the Antarium were eliminated after the first experiment. Furthermore, the crepuscular/nocturnal *M. midas* was used instead of the diurnal *M. croslandi*. While those changes improve the result, the main problem, the spectral mismatch between the Antarium's LEDs and the photoreceptor sensitivity of the animals could not be fixed.

In addition to answering the questions that the experiment was designed to ask, a further discovery was made, namely that in the morning the animals moved measurably faster. The speed difference has no explanation as the animals were treated the same way and were in the same environmental conditions during the evening and morning experiments.

Furthermore, it was observed that when the animals were shown the unstructured scene, which offered no visual cues whatsoever, some still had long, relatively straight runs and sudden changes of direction between such segments. It might be that menotaxis is part of the animal's toolbox when it gets lost, although then the question arises, why not all animals have shown that behaviour? There are animals that can make sense of the Antarium's projections and others cannot, but at the unstructured scene there was nothing to make sense of, all animals were completely deprived of any and all visual cues.

It should also be pointed out that the Antarium experiments differ from the outdoor experiments in two major ways. One, scenes are rotated while the animal is otherwise undisturbed. That has never happened to animals running on the ball outdoors. Second, the animals are presented a view that, although depicting a natural scene, for their eyes looks unnatural. The only colour component that matches the expectations is the green one, but even that is disturbed by the mismatching blue projection. The UV component is missing as the Antarium's UV LEDs excite their blue receptors more than the UV ones and the density of the UV LEDs is much lower than the green/blue ones.

# Experiment 3

## Aim

The goal of this experiment was identical to that of Experiment 2, testing the usability of the Antarium and the improvements made to it. However, in Experiment 3 *M. pyriformis* was used instead of *M. midas* and a new mounting technique, described in *Magnetic mounting* on page 158 was implemented. For external reasons the experiment had to be aborted after the first evening and only 9 animals were processed. The experiment ran for one day, in January, 2019.

## Method

The same method was used as in Experiment 2, that is, the ants were captured at the foraging tree on the outbound leg of their foraging trip. They were fed then mounted on the trackball in the Antarium.

They were presented static images of four scenes each. The order of the scenes was randomised. Each scene was shown for about 4 minutes. During that time, the scene was rotated in the Antarium randomly 90° clockwise or anticlockwise at approximately one minute intervals. Rotations were instantaneous. Scene changes first dimmed down the Antarium to black in 2 seconds, then brightened the new scene up in 2 seconds.

The trackball data, the scene order and the rotations were recorded and the ant was filmed from above.

The four scenes (see *Fig. 64*) were the view from the nest entrance, from halfway between the nest and the foraging tree, from a location that the ants were not familiar with, and lastly a computer generated image containing a uniform blue sky and green below the horizon ground, with a small amount of pseudo-random noise added to the image.

When the animal finished her run in the Antarium, she was released close to the nest entrance and was monitored until she successfully entered the nest.


*Fig. 64: Experiment 3 locations*

*The location marked Corridor is on the foraging corridor, aproximately half-way between the nest and the foraging tree. That location was used as the familiar scene in the experiments. Satellite image archived from Google Earth in 2018.*

## Results

Nine animals were processed. One of them was not cooperating and had to be excused from the experiment. With only eight animals the statistical analysis could not bear meaningful results.

However, with so few animals, it is feasible to plot all their paths and do qualitative assessment of their behaviour.

*Fig. 65* shows the paths at the familiar location. Animals 102 and 106 display orientation: after the first rotation they keep a relatively constant heading in scene coordinates by following subsequent rotations. 101 follows one rotation but ignores the last one. 104, 105 and 110 are clearly not oriented in the scene. 103 looks as if she maintained a scene heading, but looking at the ball path

reveals that it is accidental; she did not actually follow rotations. It would be hard to classify 108.

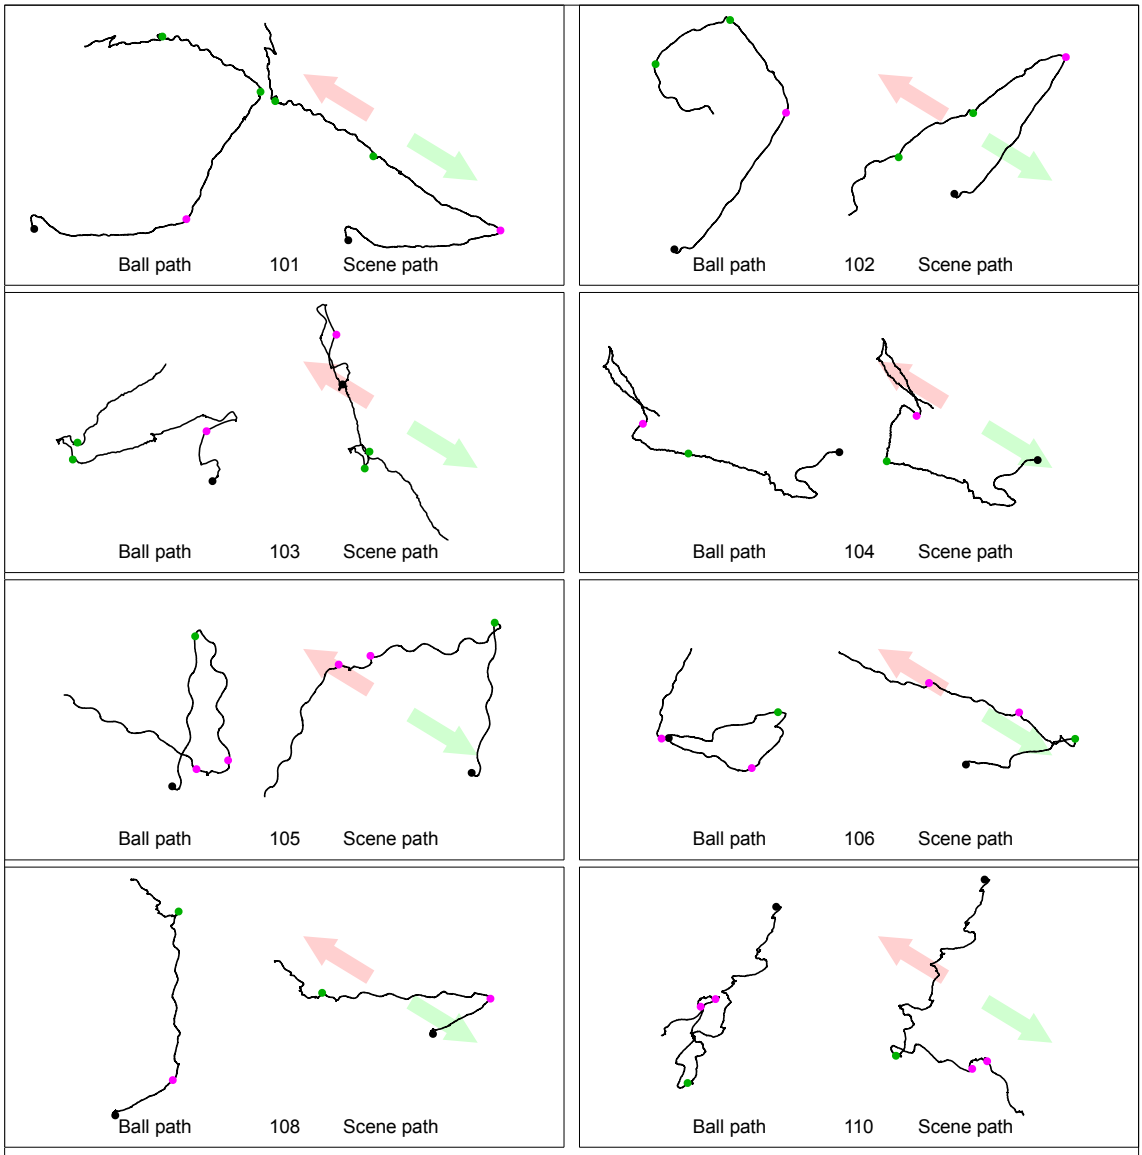


*Fig. 65: Paths at the familiar scene*

*The pale red arrow on the scene path points towards the nest and the pale green arrow points towards the foraging tree. The animal's identifier number is shown in the middle of each path diagram. The black dot indicates the start point of the path, a purple dot is a clockwise 90° rotation while the green dot is a 90° anticlockwise rotation.*

Fig. 66 shows the paths at the unfamiliar scene.

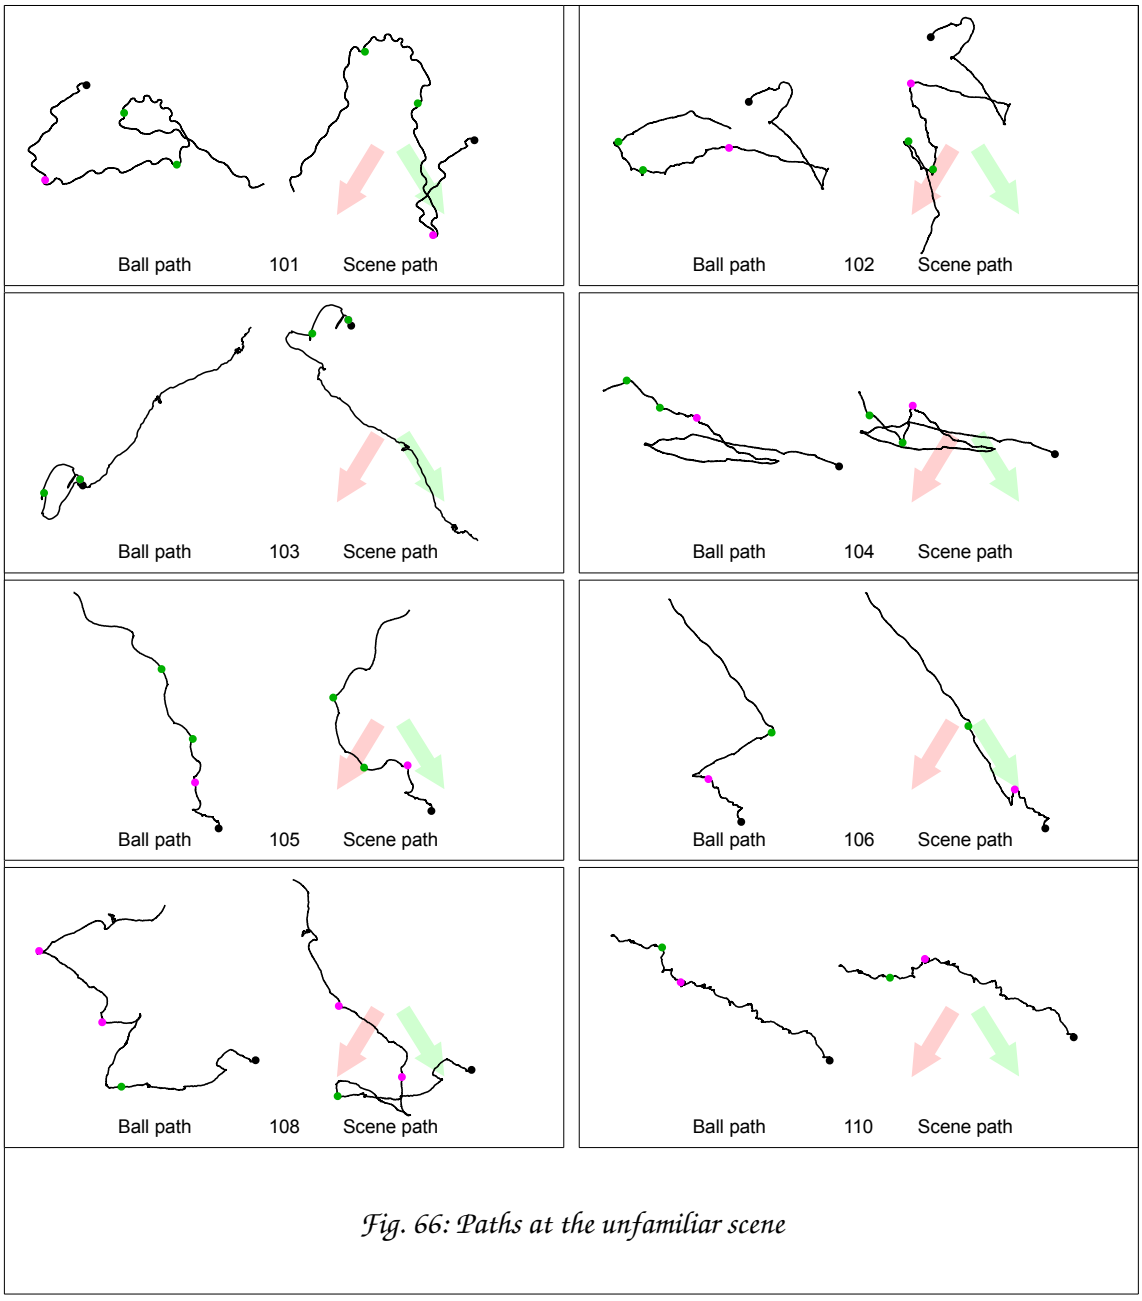


*Fig. 66: Paths at the unfamiliar scene*

It is quite obvious that animals 101 to 105 and 110 are not maintaining a direction relative to the scene; they do not follow rotations. On the other hand, ant 106 clearly does and 108, after what could be described as initial confusion also follows a relatively stable scene direction and reacts to rotations to maintain it.

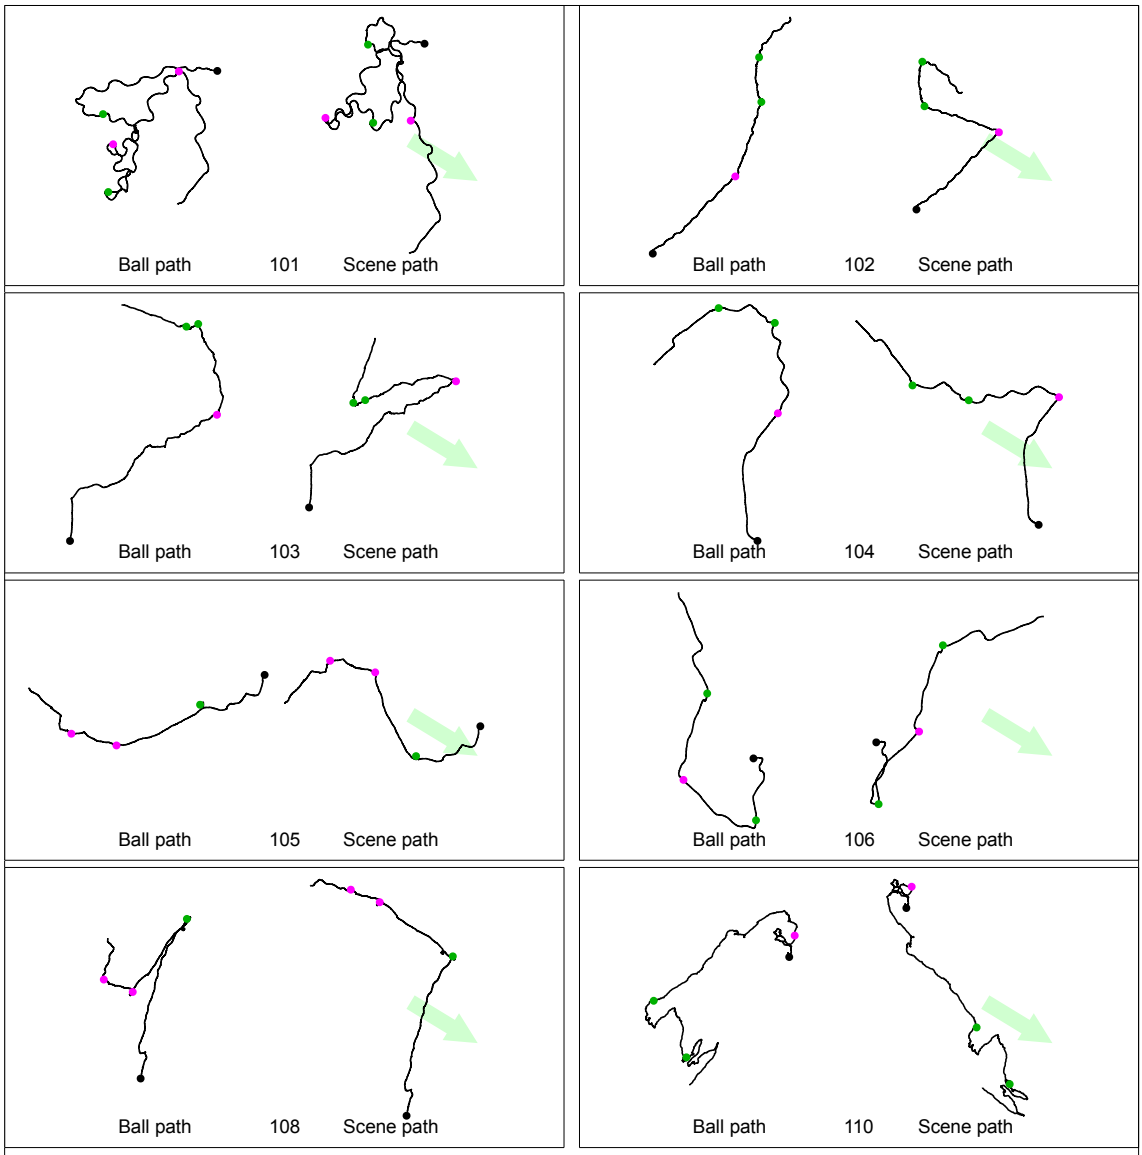


*Fig. 67: Paths at the nest scene*

The nest scene paths are shown in *Fig. 67*. Animal 108 follows the last two rotations but all other animals are ignoring them. Two animals, 102 and 105, maintain a ball path which, though not straight, is without loops and frequent random direction changes. That, however, does not translate to orientation in scene coordinates, as they do not follow rotations. The rest of the animals do not seem to be directed in either coordinate system.

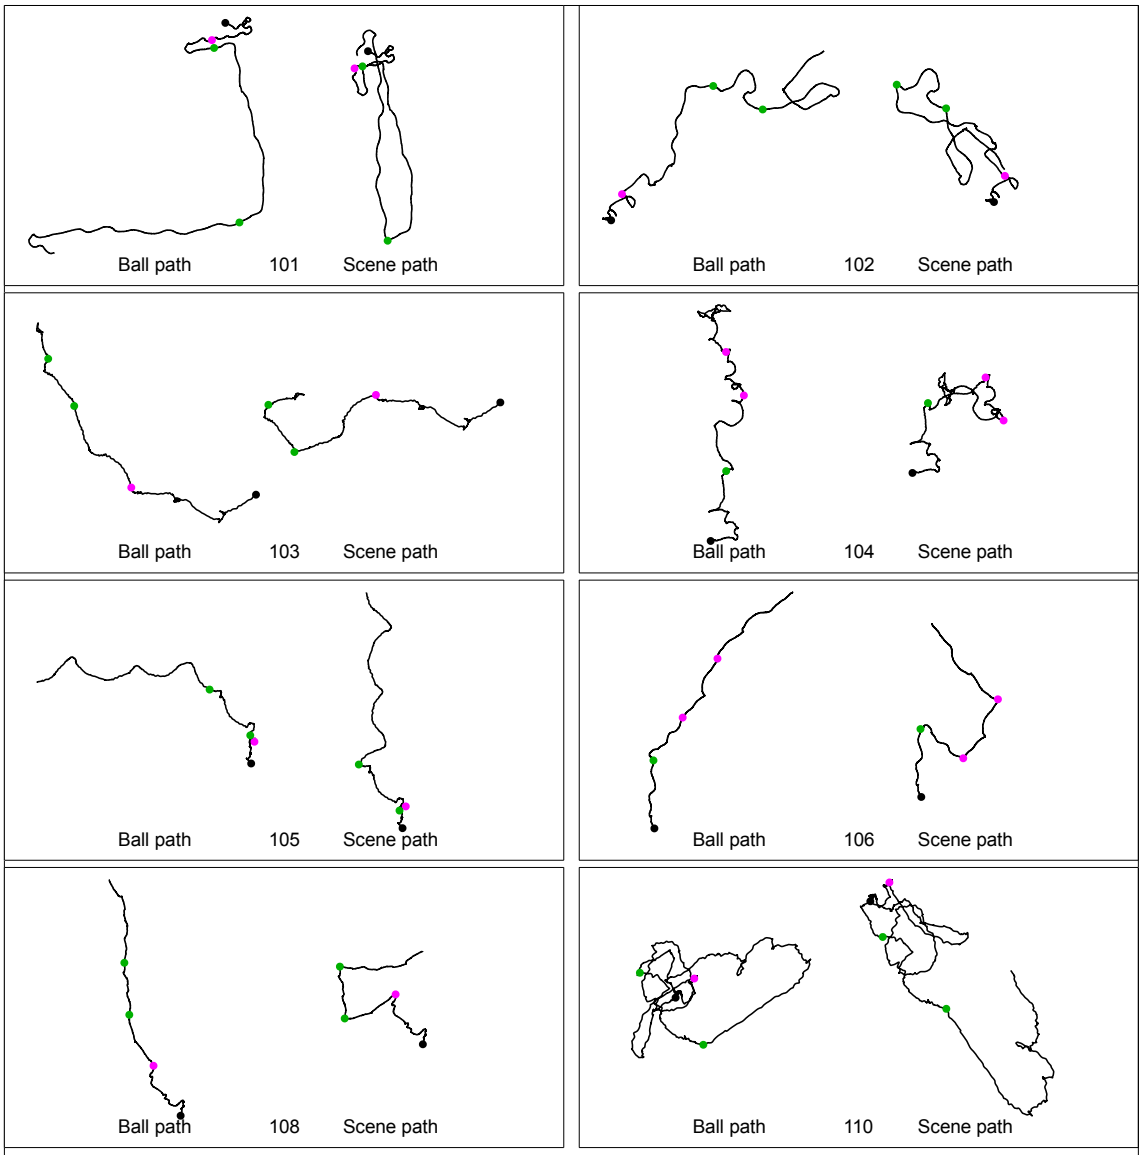


*Fig. 68: Paths at the unstructured scene*

In *Fig 68* the paths at the unstructured scene show no orientation in scene coordinates at all. Some animals more or less maintain a direction on the ball (103, 106 and 108), while 101 has long, straight stretches between direction changes. The rest of the animals show no orientation in either coordinate system.

## Discussion

In the previous two experiments it was hypothesised that about 20-25% of the animals can make use of the Antarium's spectrally mismatched visuals and navigate successfully. The path analysis results of this experiment corroborate that, 2 of the 8 animals (25%) have shown clear navigation at the familiar scene.

As it was the case in Experiment 2 (see Fig 62 on page 206), animals that show orientation at the familiar scene often start in a random heading but the first scene rotation triggers them to re-evaluate the panorama and pick a new direction, which they then subsequently maintain.

There was one animal (106) that maintained her heading in scene coordinates at the unfamiliar scene (see *Fig. 66* on page 216), moving away from the nest. That behaviour has also been observed with displaced animals during outdoor experiments (Jochen Zeil, personal communication). Ant 106 also was one of the two animals which have shown clear orientation at the familiar scene (*Fig. 65* on page 215) and, in fact, she was heading towards the nest.

No orientation was observed at the nest scene, which is in accordance to the predicted random search behaviour. As expected, when confronted with the unstructured scene no animal has shown any indication of visual navigation. It Surprisingly, though, some animals maintained a constant heading for quite long segments. It cannot be completely excluded that they used some part of the tether that held them in place as a "landmark" and kept their heading relative to that.

Thus, even though the number of animals was not sufficient to perform a quantitative analysis as in Experiment 2, the qualitative analysis in this experiment was still useful in corroborating the conclusions drawn from Experiment 2.

## Conclusion

The experiment results have demonstrated that the Antarium is a viable experimental apparatus in studying the navigation of ants. Although in Experiment 2 and 3 the device operated in open loop mode, the rotation of the panorama and the instantaneous displacement of the animal from one location to another are experimental paradigms that would be very hard to do in the natural environment. In fact, with a new version of the device with LEDs spectrally matching the photoreceptors of the animals, experiments designed to study the animals' reaction to instantaneous displacements could furnish insights into their navigation mechanisms.

A shortcoming of the current Antarium device is the emission spectra of its LEDs, which are a poor fit for the photoreceptor sensitivities of *Myrmecia* ants. However, some animals (about 20-25% of the cohort) can clearly derive navigational guidance from the projected visual scene, despite the spectral mismatch.

Therefore, even in its current form the Antarium can be used to perform navigation experiments that would not be possible in a natural environment, assuming that one keeps in mind that only 1 in 4 or 5 animals will be clearly guided by the scenery projected by the device. Furthermore, it is hoped that a new device with the proper LED wavelengths and active polarisation will elicit responses from the ants close to what one would expect in a natural setting.

The experiments also raised two questions that require further study: why the same animals under the same environmental conditions and seeing the same panorama move faster in the morning than in the evening; and why some animals are better than others at interpreting the spectrally mismatched imagery that the device supplies to them.

# Altered reality experiments

## Abstract

Two experiments were performed in the Antarium where we changed the projected visual world.

In the first, we moved the foraging tree to test whether the animals go to the *tree* using the tree itself as a beacon, or they go to its *expected location* using the whole panorama. The results indicate that the animals use the whole panorama, which is in agreement with published studies demonstrating the effect of physical tree removal on bull ant navigation. In addition, I analysed the rotational image difference function model of navigation (see page 71) and found that it indeed explained the observed behaviour.

In the second, we re-coloured the world to see how colour and colour contrast changes affect the navigational ability of the animals. Due to the deficiencies of the Antarium, we did not get an answer to our question, but we discovered that even a very small amount of UV contrast measurably improves the navigation performance of ants.

## Introduction

The aim of these experiments were to use the Antarium to ask questions that, without a virtual reality setting, would be very hard to ask and answer.

From the proof of concept experiments it was known that the Antarium's level of fidelity in reconstructing the visual scene is low. However, the experiments also have shown that about one quarter to one third of the animals can, and do, interpret the imagery provided by the device.

Therefore, if we can establish an objective measure to classify the animals into two groups, one that clearly responds to the virtual world and one that does not, the former group can provide useful information despite the deficiencies of the current Antarium.

Two experiments were performed. One probed a long-standing question: do animals navigate using individual landmark objects or by the whole panorama. It

is a non-trivial question to ask because modifying the natural panorama or manipulating the natural objects that the animals can use as landmarks is often simply not possible: hills, trees or buildings are not easy to move around. In a virtual world, however, those are just 3D models that can easily be manipulated.

The second question we investigated was the importance of colour. In the Antarium we can arbitrarily change the hue so that we can assess the animals' reaction to falsely coloured scenes. To experimentally test the role of colour vision in navigation is very hard to do without a virtual reality device.

# Experiment 4

## Introduction

This experiment uses the Antarium to ask whether ants heading toward a foraging tree use the tree itself as a beacon (landmark based navigation), or they choose a heading relative to the whole landmark panorama (panorama based navigation). In a feature rich environment both strategies are sound. An overview of that topic can be found in (Wystrach et al., 2011c), with the conclusion that the ants probably use either or both, depending on what method provides them with the most salient navigational cues.

There are two reports investigating how *M. pyriformis* respond to scene changes. Narenda and Ramirez-Esquivel found that felling dead trees close to the foraging corridor made the animals less certain in their navigation, but after a few days they still found the foraging tree (Narendra and Ramirez-Esquivel, 2017). Immediately after the removal of the trees the majority of the animals were lost. However, one day later only about 40% of the animals had to abandon the foraging trip and two days later all animals reached the tree. The uncertainty of the animals was measured by the straightness of their paths to the foraging tree and by their speed. After the removal of the tree the animals slowed down and took more tortuous paths, but their uncertainty decayed over time. It was also discovered that before the tree removal most animals left the nest and went to the tree without looking back towards the nest. Immediately after the removal almost half of the animals performed a short learning walk before heading to the tree, but the number of animals re-learning the visual environment also quickly went back to normal.

Reid *et al* experimented with blocking parts of the panorama by means of a screen (Reid et al., 2011). At three separate *M. pyriformis* nests a screen was erected about 1.5m from the nest, blocking ¼ of the panorama. The screen could be positioned to block the view in front of the animal, behind her, left or right of her, relative to the foraging direction. The initial heading, the time it took the ant to get 30cm from the nest and the percentage of the animals that crossed a 90° arc 1.2m from the nest in the foraging direction were measured and statistically analysed. The results indicated that the animals were most

confused when the screen was in front of them when heading towards the foraging tree. However, the level of confusion (as attested by the three parameters measured for the animals) varied greatly between the sites. The screen only blocked the panorama up to 45˚ elevation. The presence of the screen, even when it was in the direction of the foraging tree, caused only a minor disruption where the canopy outline above the screen had salient navigational features. However, at a nest where the high-elevation view lacked salient cues the ants became very confused. The authors concluded that the animals use the entire panorama for navigation and as long as abundant navigation information is provided by the visible part of the scene, they will not be significantly affected even if the foraging target itself is obscured.

In addition to the literature, I made a personal observation as well. While the Antarium was built, a tree close to the laboratory had to be removed for safety reasons. Incidentally, that tree was a foraging tree for the same *M. pyriformis* nest that was used in this experiment. I observed that for several days animals still attended the location of the felled tree, searching around the tree stump and the cut-up woodpile beside it, then going back to the nest. The number of animals attempting to forage at the tree waned every day. About a week after the removal of the tree no *M. pyriformis* could be observed heading in that direction any more. That observation corroborates the conclusion of the two papers, namely that *M. pyriformis* primarily determine heading direction attending to the whole visual panorama rather than to individual landmarks.

The aim of Experiment 4 was to test that statement.

## Methods

The experiment was performed in 3 evenings (non-consecutive) in January, 2019 and it included 28 animals.


*Fig. 69: Moving tree experiment locations*

*The satellite image is from Google Maps, archived in 2018.*

### *3D model preparation*

Before the commencement of the experiment, the 3D model of the animals' natural habitat was pre-processed. The foraging tree was extracted from the 3D model of the scene. It thus became an independent 3D object that could be placed anywhere in the scene. The scene itself was rendered from the 3D model augmented with a panoramic photograph taken at the `Release` location.

Four *scenes* were defined. Each scene shows the environment from the location marked `Release` in *Fig. 69*. They differ in that the foraging tree can be at its natural position, removed from the scene, rotated around the `Release` point by 64° anticlockwise or 117° clockwise; the latter two positions are marked `Shadow tree 1` and `Shadow tree 2` in *Fig. 69*, respectively. *Fig. 70* shows the four scenes.

The scene coordinates are either aligned with the ball coordinates, rotated 90° clockwise or rotated 90° anticlockwise. That results in 12 possible scene and

alignment combinations. The two combinations where both the scene and the tree within the scene are rotated in the same direction are omitted, resulting in 10 *views*.

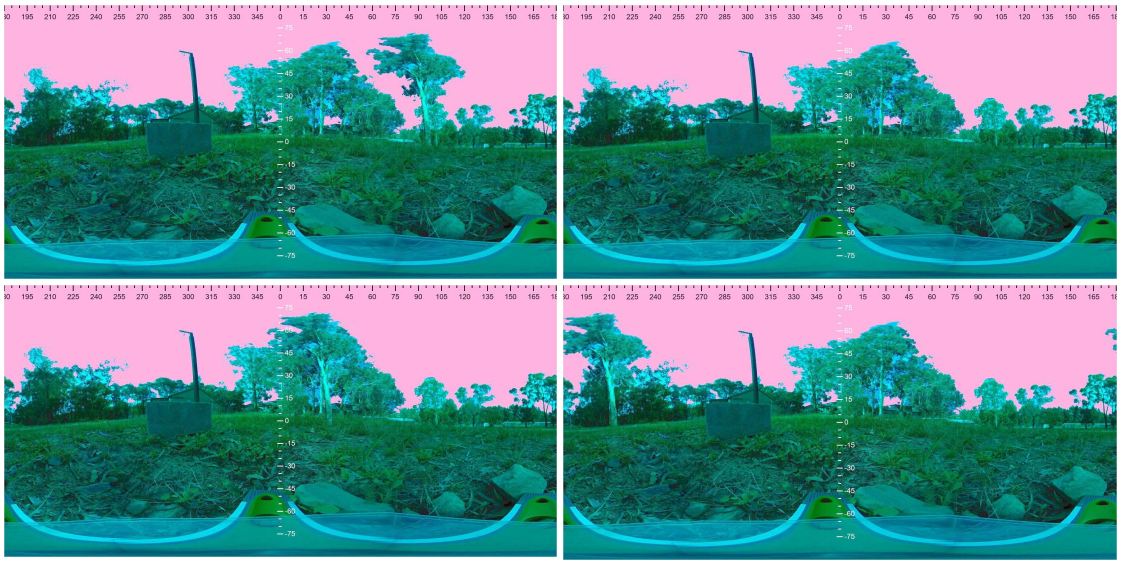

*Fig. 70: Moving tree experiment scenes*

*On all panoramic images the white scale at the centre is the elevation angle. The black scale on the top is the azimuth, 0˚ is magnetic North. The nest entrance is at 276˚ azimuth and -20˚ elevation, at the bottom-left corner of the concrete bollard with the vertical metal rod. The building in the background, behind the bollard, is the laboratory. Below -50˚ elevation, not projected in the Antarium, the levelling platform of the panoramic camera can be seen. The red channel of the images represent the UV added to the scene above the skyline, hence the purple sky.*

*Top left: natural scene, the foraging tree is at 96˚ azimuth. Top right: foraging tree removed. Bottom left: foraging tree rotated anticlockwise, it is at 32˚, in front of a group of distant trees. Bottom right: foraging tree rotated clockwise, it is on the left at 213˚. Images courtesy of Trevor Murray.*

### *Animal handling*

*M. pyriformis* were captured immediately after sundown on the outbound leg of the foraging trip, approximately 1m from the nest on route to the foraging tree, close to the place marked `Release` on *Fig. 69*.

The animals were then prepared for their session in the Antarium as described in *Magnetic mounting* on page 158. In order to keep their motivation to go to the foraging tree, they were not fed before the session.

Every animal was shown all 10 views, in random order, each for approximately 40 seconds. The change between consecutive views was instantaneous. The Antarium was set up for open-loop mode, the trackball movements had no effect on the projected scenery.

During the session the trackball data and the view changes were recorded for later analysis. Furthermore, the animals were filmed using the Antarium's overhead camera for later analysis of their gaze direction.

When an animal completed her Antarium session she was fed and then released close to the nest. She was monitored until she entered the nest.

## Results

The animal was shown 10 views, each of which has projected one of the 4 scenes (normal, no tree, tree rotated clockwise or anticlockwise) with the scene coordinate system being rotated by -90°, 0° or +90° relative to the ball coordinate system.

For each view the animal's position at the end of the view, relative to her position at the beginning of the view, was calculated in scene coordinates. These displacement vectors were sorted by the scene they projected and then all animals were pooled.

The resulting vector ensemble is shown in *Fig. 71*.

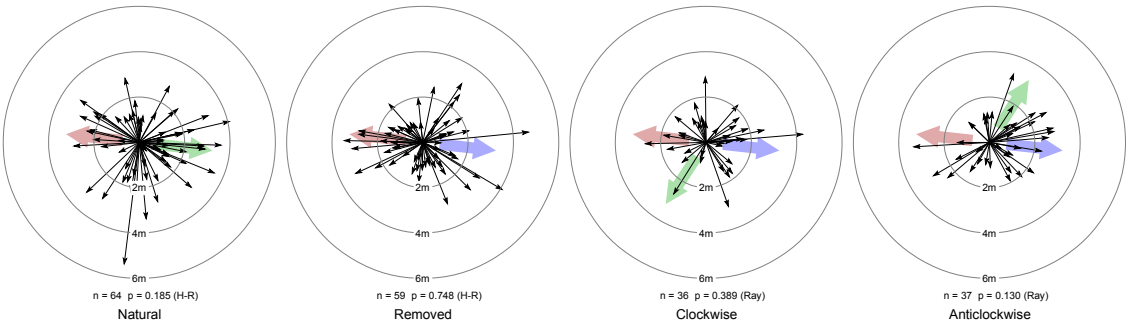


*Fig. 71: Final positions at the end of the view*

*Under each plot n is the number of position vectors plotted while p is the smallest p-value of 3 statistical tests, the Hermans-Rasson (H-R), the Rayleigh (Ray) and the Rao (Rao) test for circular uniformity. The tests use only the bearings and do not take the vector length into account. As the p-values show, the animals' bearings at any scene cannot be statistically differentiated from a circularly uniform random distribution.*

*The pale red arrow indicates the direction of the nest, the pale green arrow shows the direction where the tree was projected and the pale blue arrow points towards the original (natural) location of the tree.*

Not much can be inferred from the displacement vectors themselves, the statistical tests did not reveal any apparent orientation of the animals. A more fine grained test is to decompose the animals' paths into small, equal length (2.5cm) vectors and calculate the direction of each such vector, then generate the empirical probability density function for the directions, as it was described in Heading density on page 195. Doing that and pooling all animals together results in the distribution shown in *Fig. 72*.

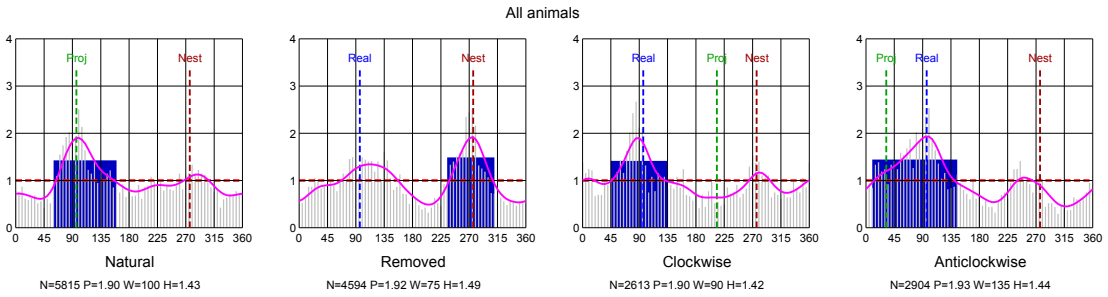


*Fig. 72: Empirical vector direction probability density functions*

*The grey vertical lines represent the 5° bins into which the vectors were sorted. The vertical axis is the frequency, relative to a perfectly circular distribution (marked with the red horizontal dashed line). The magenta curve is the empirical density function created by applying a Gaussian smoothing kernel on the bin counts.*

*The vertical dashed red line marked Nest indicates the direction towards the nest, the green Proj line, if present, shows the direction towards the projected tree while the blue Real line, if present, indicates the direction to the foraging tree in real life.*

*N is the number of vectors sorted into the bins, P is the height of the largest peak, W is the angular distance between the points where the curve crosses the uniform line left and right of the peak and H is the mean of the curve over that interval. The blue rectangle is W wide and H high.*

The figure shows that when the natural scene was projected, the vector directions peak towards the foraging tree, with a much smaller secondary peak in the direction of the nest. When the foraging tree is removed from the scene the peak shifts to the nest and the secondary peak is towards the direction where the tree should be. In case of the rotated tree, regardless of the direction of the rotation, the highest peak is towards where the tree should be and a secondary peak is in the direction of the nest.

However, in the previous experiments it was established that peak values below around 2.8 can occur for non-navigating animals while good navigating animals produce peaks above 3.6 (see page 203). None of the peaks in *Fig. 72* are above 2. It was also established that only about a quarter of the animals can extract navigation information from the Antarium's spectrally mismatched projection. Therefore, a method needs to be found that can select the animals that are "good navigators" in the Antarium.

In the previous experiments an animal that has shown orientation was defined as a good navigator. In this experiment we are looking for the animal's response to the displacement of the foraging tree. The answer to our question lies not in whether the animal was maintaining a heading but in the actual heading. Therefore, the selection criteria have to be modified.

Because we are asking a question about behaviour when the position of the foraging tree has been changed, we can only use the unmodified, natural scene parameters in the selection criteria. When an ant is presented the natural scene, she is expected to go to the foraging tree, or possibly to the nest.

In light of the above, the following selection process was established:

- The first 5 seconds of the animal's path for each view (scene and rotation combination) is discarded, to give the animal time to respond to the scene and/or rotation change.
- The remaining 35 seconds of her path is vectorised. If the number of 2.5cm vectors is less than half the number of direction bins, the path is discarded. Note that since there are 72 bins, the path gets discarded if the length of the path was shorter than 90cm. In 35s that corresponds to the speed of 0.026m/s while *M. pyriformis* normally walks on the ball at around 0.1m/s. Thus, a path is eliminated if the animal was not moving (at least not translating) about ¾ of the time the given view was projected.
- If the animal has less than 2 paths remaining at the natural scene, or if at any other scene all her paths were discarded, she is eliminated.
- Otherwise, the paths belonging to the natural scene are pooled for the animal. The vectors are then sorted into 5° bins and the empirical vector direction probability density curve is calculated using a Gaussian smoothing kernel. It is then normalised relative to a circularly uniform distribution.
- If the curve has a peak that is at least 3.8 high (the limit established in the previous experiment) and is located within ±45° of the foraging tree direction or the nest direction, the animal is selected as a good navigator.

- Otherwise, if the curve has two peaks, both larger than 1.0, one within ±45° of the foraging tree direction while the other within ±45° of the nest direction and the sum of the two peaks is at least 3.8, the animal is selected as a good navigator.
- Otherwise the animal is eliminated.

The above process selected 8 animals as good navigators while 20 animals were eliminated.

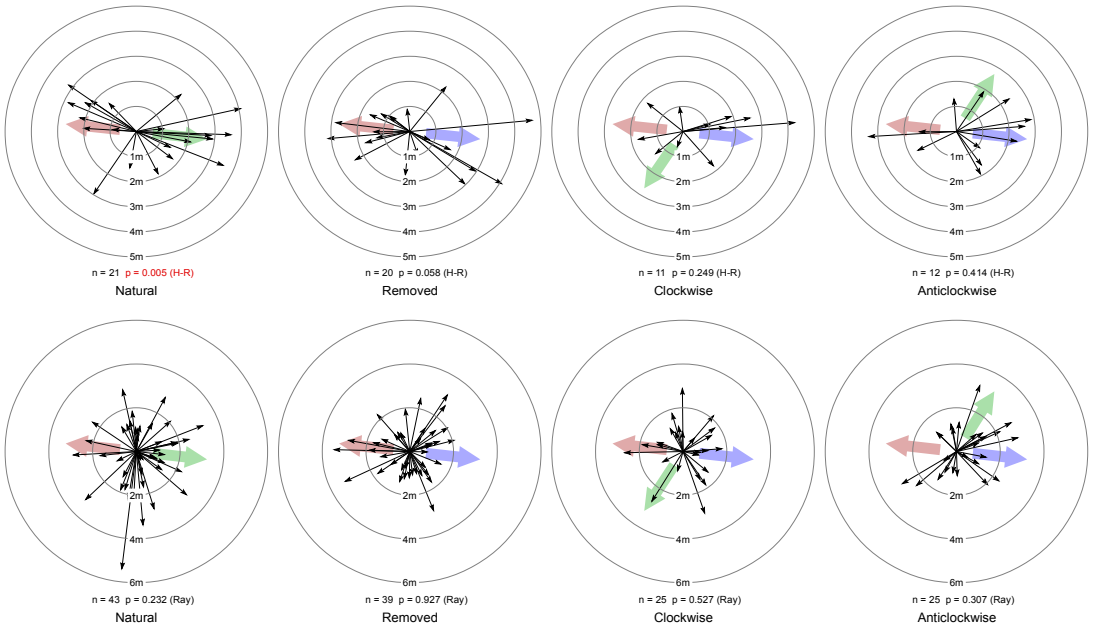


*Fig. 73: Bearings of the animal groups*

*The upper panel shows the bearing of the 8 animals that were classified as good navigators and the bottom panel is the bearings of the rest of the cohort.*

*Fig. 73* shows the bearings at the end of each view of the good navigators and that of the rest of the cohort. The very low p-value at the natural scene for the good navigators is not surprising, because moving in specific directions was part of the selection criteria.

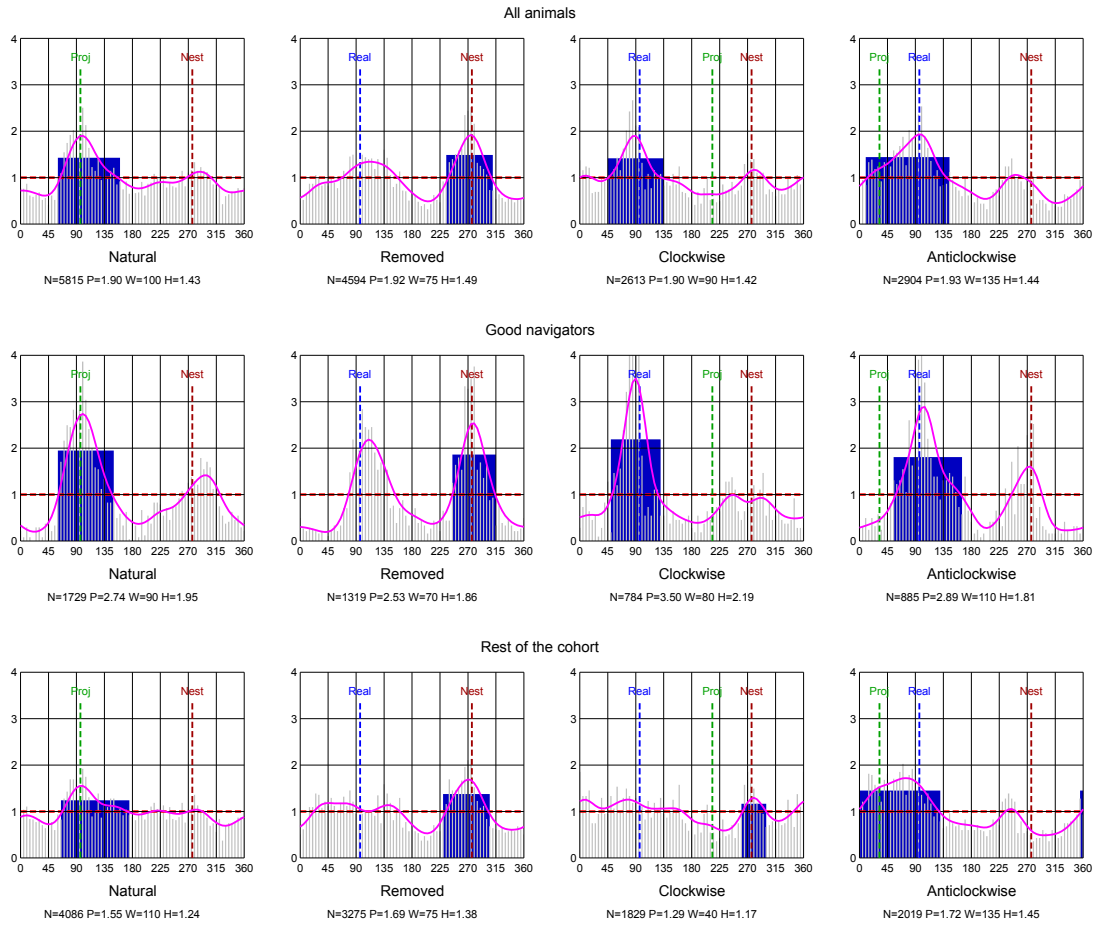


*Fig. 74: Vector direction probability density functions for animal groups*

*The grey vertical lines represent the 5° bins into which the vectors were sorted. The vertical axis is the frequency, relative a uniform circular distribution (marked with the red horizontal dashed line). The magenta curve is the empirical density function created by applying a Gaussian smoothing kernel on the bin counts.*

*The vertical dashed red line marked Nest indicates the direction towards the nest, the green Proj line, if present, shows the direction towards the projected tree while the blue Real line, if present, indicates the direction to the foraging tree in real life.*

*N is the number of vectors sorted into the bins, P is the height of the largest peak, W is the angular distance between the points where the curve crosses the uniform line left and right of the peak and H is the mean of the curve over that interval. The blue rectangle is W wide and H high.*

*Fig. 74* shows the empirical vector direction distribution density functions for the good navigators and for the rest of the cohort. The function for all animals, already presented in *Fig. 72*, is repeated for convenience.

It is clear that when the good navigators are removed from the cohort, the remaining animals as a group show little orientation. Individually, some of the animals have definite peaks, but those are in random directions and cancel out when the animals are pooled.

The good navigators, when presented the natural scene, head towards the tree and to a lesser degree, towards the nest. When the tree is completely missing from the scene, the probability of nest-ward movement becomes elevated, but still, heading in the direction where the tree should be remains prominent. If the tree is present but at an incorrect location, the animals were most likely to go towards the place where the tree should have been and have not shown preference towards the direction of the actually projected tree.

***Rotational image difference function analysis***

The scenes that the animals were shown have been analysed in order to find an explanation for the observed behaviour, both outdoors and in the Antarium. The following will detail the processing steps and the findings.

Ommatidial model

An ommatidial map was created. The eyes of the animals are composed of ommatidia which form a slightly distorted hexagonal grid, mapped on the surface of two flattened hemispheres, as described on page 24. The output of the two eyes of are combined, thus both sides of the animal's brain receive the output from both optical lobes as it was explained in Higher level brain centres on page 42.

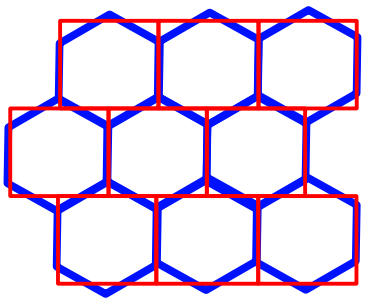

*Fig. 75: Mapping between a hexagonal and a rectangular grid*

In a simplified model, instead of a hexagonal grid, a rectangular grid was used, where each grid row is shifted by half a rectangle. That is topologically equivalent to the hexagonal grid, but is much easier to handle in calculations. The equivalence is demonstrated in *Fig 75* where the two grids are shown overlapping.

I ascertained that the change from hexagonal to rectangular grids made no perceivable difference in the results of subsequent image processing.

A further simplification is that the varying local radius of the eyes (which results in a non-uniform angular resolution), the overlap of the anterior visual fields, and the loss of the posterior and ventral visual fields occluded by the body are not taken into account. The visual field of the eyes is assumed to be spherical with a uniform angular resolution.

In the model each ommatidium (i.e. a rectangle) has an angular acceptance function exactly matching the projection of that rectangle to a sphere. Furthermore, the optical gain within the angular acceptance function is assumed to be uniform.

The distribution of the UV, blue and green sensitive photoreceptors over the *Myrmecia* eye is not known. Therefore, in the model it was assumed that each ommatidium contains all three receptors, that is, the animal has perfect trichromatic vision.

Image mapping

The panoramic images are mapped to the ommatidial grid by averaging all the pixels that fall into the acceptance area of the rectangle using a spherical projection. For this model the angular separation of adjacent ommatidia was chosen as 2.4°, resulting in just over 7 thousand ommatidia on the sphere all together. Those figures are reasonably aligned to the actual figures for *M. pyriformis* (Narendra et al., 2011).

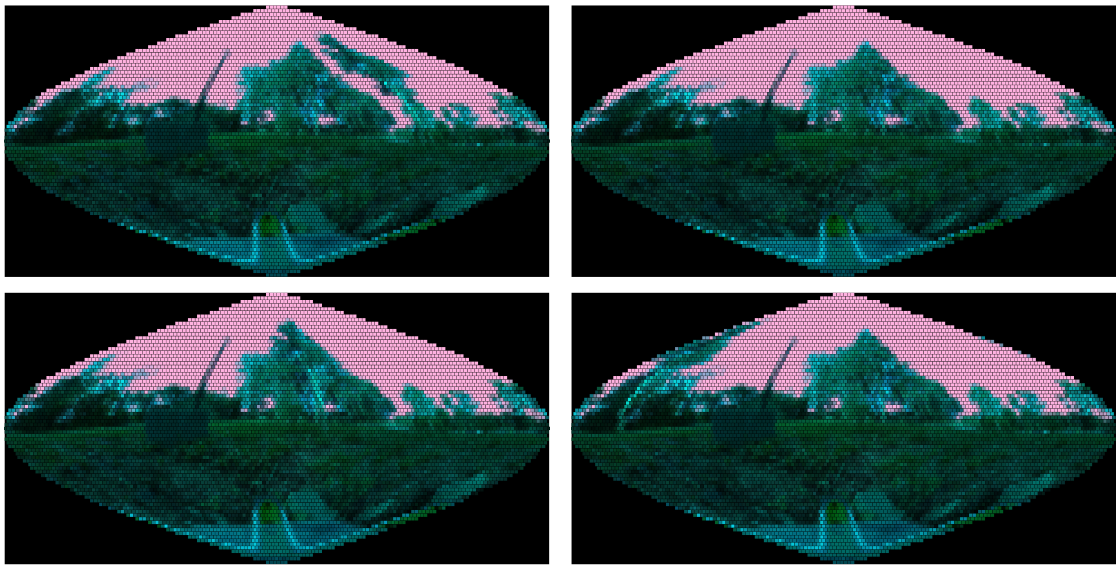

*Fig. 76: Panoramic images mapped to the ommatidial array*

*The foraging tree at the normal position(top left), removed (top right), rotated anticlockwise (bottom left) and rotated clockwise (bottom right). The red channel represents the artificially added UV component, which is a uniform illumination above the skyline and zero below.*

The mapping of the four panoramic images to the ommatidial grid are demonstrated in *Fig. 76*:

Contrast normalisation

To compare like to like, the images are transformed to equalise their contrast. There are several methods, see for example (Stürzl and Zeil, 2007). The method used in this experiment is called *intensity histogram equalisation*[19]. It calculates the empirical probability density function (PDF) of pixel luminosity values and applies a transformation on the values so that the PDF of the resulting image becomes flat. That process increases contrast. It also increases noise, but the gain by contrast enhancement usually outweighs the negative effect of the more pronounced noise. The method was chosen because there seems to be evidence for it being used by insects as well. Laughin reported that the large monopolar cells in the blowfly lamina seem to be responsible for an intensity transformation which, as indicated by intracellular recordings, is

19 *The exact details of the algorithm and its detailed analysis are beyond the scope of this thesis, but Wikipedia provides a good introduction and links to articles for the interested Reader: https://en.wikipedia.org/wiki/Histogram_equalization*

equivalent to the histogram equalisation procedure on computers (Laughlin, 1981).

The normalised images are shown in *Fig. 77*:

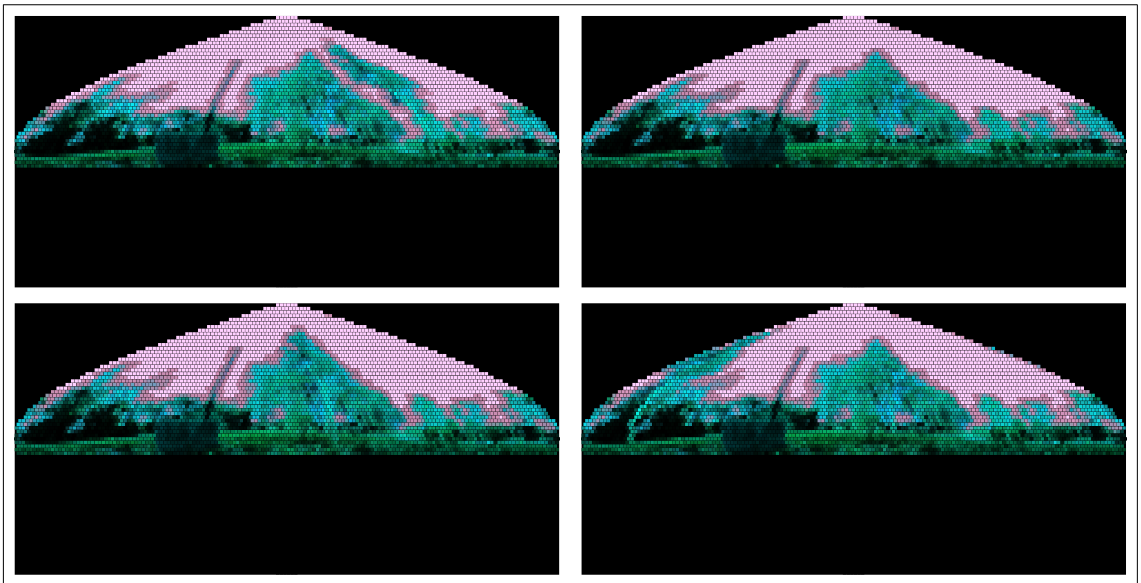

*Fig. 77: Ommatidial grids of the outdoor panorama after histogram normalisation*

*The bottom part of the images, below -12° elevation, have been blacked out because nearby ground features are unlikely to play any role in the long-distance navigation of the animal. The images, clockwise from top-left, are: foraging tree at its normal position, removed, rotated clockwise and rotated anticlockwise.*

Rotational image difference function

The rotational image difference function (rotIDF) was explained in detail in the Introduction on page 71. The calculation is performed by first creating an ommatidial map of the reference image, which in our case is the natural scene:

$$\boldsymbol{M} = T(\boldsymbol{I}(\phi, \theta))$$

where $\boldsymbol{I}$ is a panoramic image defined by the $\phi$ azimuth and $\Theta$ elevation angles, $\boldsymbol{M}$ is the ommatidial map and $T$ is the transformation operation. We then define the rotational image difference function of two images, $\boldsymbol{I_1}$ and $\boldsymbol{I_2}$ as:

$$rotIDF(\alpha) = \sqrt{\frac{1}{N} \cdot \sum_{N} [T(\boldsymbol{I_1}(\phi, \theta)) - T(\boldsymbol{I_2}(\phi + \alpha, \theta))]^2}$$

where $N$ is the number of ommatidia in the map and $\alpha$ is the rotation angle. In effect, the rotIDF is the root mean square of the ommatidium by ommatidium difference between the two mapped images, of which one was circularly shifted (i.e. rotated) before the calculation. The intensity values are normalised to the maximum so that the rotIDF() function returns 0 if and only if the two images (after mapping) are identical and returns 1 if and only if all ommatidia of both images are completely saturated (that is, either 0 or 1 with no intermediate values) and each and every ommatidium has opposing values for the two images.

The calculation is performed on the green, blue and UV channels separately. Since the animals are expected to have memorised the natural image, that will always be the reference image. *Fig. 78* shows the resulting rotIDF functions for the four scenes.

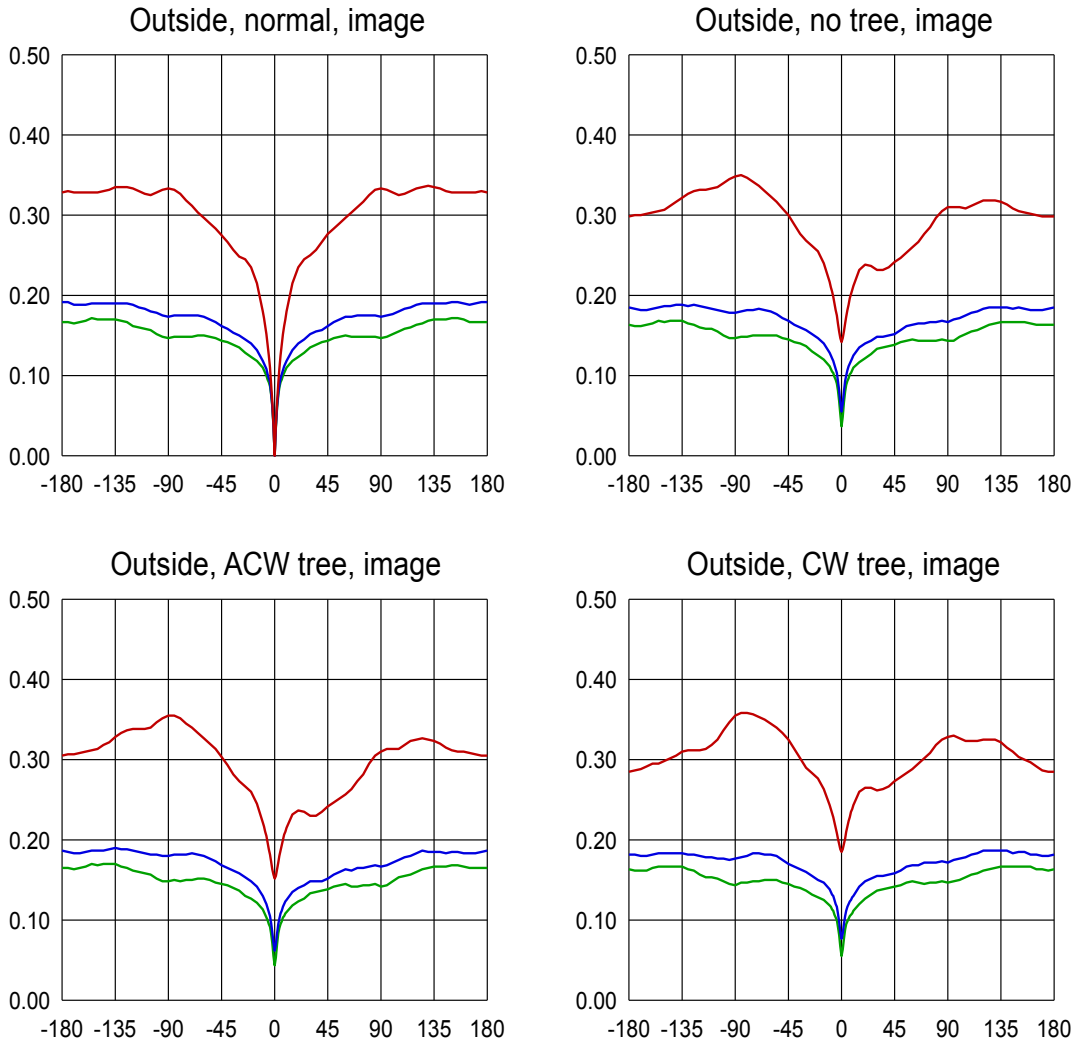


*Fig. 78: Rotational image difference function, original reconstructed views*

*The reference image is the unmodified (denoted normal) image in all cases. The red curve represents the UV component, which was artificially added to the panoramic images as a uniform illumination above the skyline and constant zero below it. The blue and green lines are the blue and green channels, respectively. The image difference function calculation was limited to ommatidia in the +90° to -12° elevation range to filter out nearby ground features that are unlikely to play any role in long-range navigation.*

The global minimum in all four cases, and on all colour channels, is at zero rotation. There is a perceivable second local minimum at 180° on the modified images, but that minimum is significantly higher than the global one. It should

also be noted that the largest difference between the maximum and minimum values of the rotIDF is in the UV channel.

The existence of a definite global minimum offers an explanation for the observations mentioned in the *Introduction* section on page 223, namely that in real life even after removing several trees in their environment the animals will still find the foraging tree and that after felling the foraging tree itself the animals will keep going to its former location for a while. A single tree is a small part of the overall image and its presence or absence does not change the fact that the best match is when everything else is aligned between the reference and the momentary view.

However, what we are interested in is the animals' behaviour in the Antarium. Therefore, we need to create an ommatidial map of what the animals saw in the device. It is a multi-step process.

First, the panoramic image is mapped to the Antarium's pixels, as it was demonstrated on page 179.

Then the pixels of the Antarium (modelled as small lit circles in front of a completely black background) are mapped to the ommatidial grid. Because the Antarium's LED wavelengths do not match the photoreceptor sensitivity of the animals, when the Antarium's projection is transformed to an ommatidial grid, a colour transformation is also performed. The transformation is the following:

$$\begin{pmatrix} G' \\ B' \\ U' \end{pmatrix} = \begin{bmatrix} 0.46 & 0.34 & 0 \\ 0 & 0.43 & 0.64 \\ 0 & 0 & 0.20 \end{bmatrix} \begin{pmatrix} G \\ B \\ U \end{pmatrix}$$

where G, B, U are the green, blue and UV LED intensities and G', B' and U' are the excitation levels of the photoreceptors of the same colour. The transformation matrix is based on the measured sensitivity curves and actual LED emission spectra, see *Fig. 39.* on page 172 and also the relative intensities of the LEDs. The images that result are shown in *Fig. 79*:

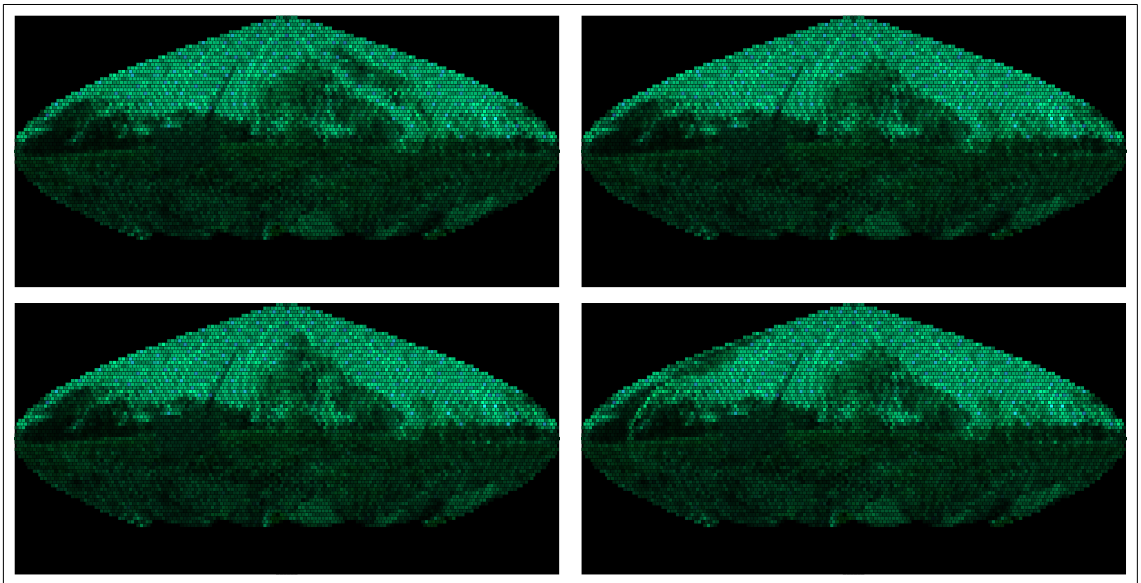

*Fig. 79: Scenes projected in the Antarium mapped to ommatidia*

*The scenes, clockwise from top-left are: tree at its natural position, removed, rotated clockwise and rotated anticlockwise. Note the lack of UV component and the green-tinged sky. The noise observable over the sky area is due to the moiré pattern between the ommatidial grid and the Antarium's LEDs.*

The images are unnaturally green, because the green receptors are excited by both the green and the blue LEDs. The blue receptors are excited by the blue LEDs, but only with a diminished efficiency. The UV LEDs excite mostly the blue receptors and to a very small degree the UV ones, but there is only 1 UV LED for each 40 blue/green LED pair. The moiré pattern between the ommatidial grid and the Antarium's LED grid introduces noise to the image, which is most prominent in areas with otherwise homogenous intensity and hue, that is, the sky.

The result of applying the same contrast equalisation procedure to these images that was applied to the outdoor images are shown in *Fig. 80*:

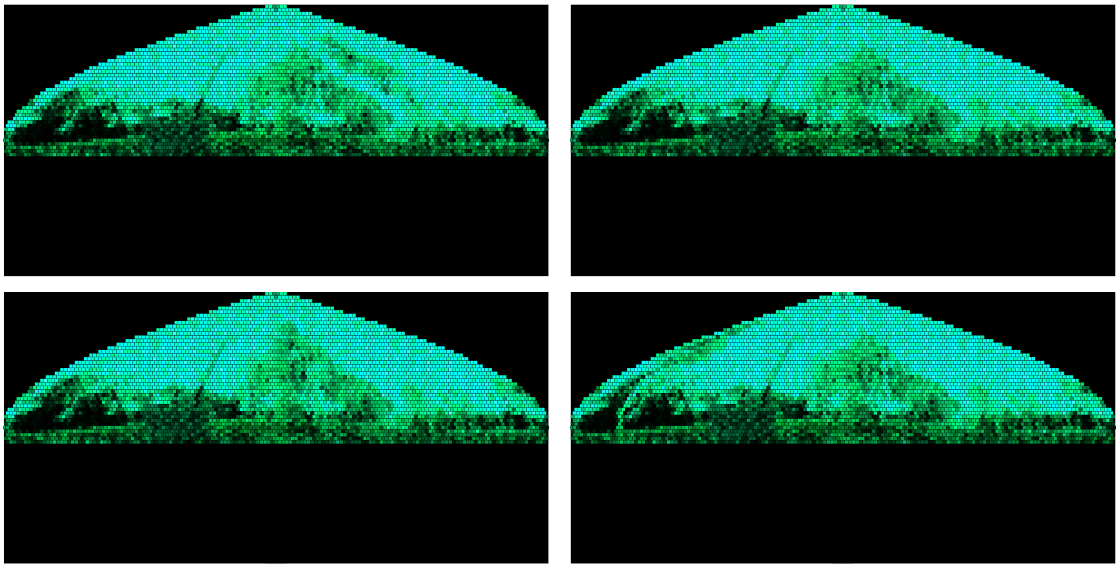

*Fig. 80: Scenes projected by the Antarium, contrast equalised*

*The scenes, clockwise from top-left are: tree at its natural position, removed, rotated clockwise and rotated anticlockwise. The bottom part of the images, below -12° elevation, have been blacked out because nearby ground features are unlikely to play any role in the long-distance navigation of the animal.*

After the contrast equalisation the sky seems to have shifted from green towards cyan, but that is actually an optical illusion due to the colour representation on computer monitors and in the printing process, and the peculiarities of human colour perception. The contrast equalisation process only changes the intensity of the image and leaves the hue and saturation levels untouched.

To work out what navigation information the animal could extract from these images, the rotIDF can be calculated between them and a reference image. The animals learned their environment outdoors, therefore the reference image is the natural scene, as seen outdoors, contrast equalised; that is, the top-left panel of *Fig. 77*. However, when calculating the rotIDF, the rotation can be applied at two ways: either we rotate the panoramic image, map it to the Antarium's LEDs then map that to the ommatidia; or we map the panoramic image to the Antarium, rotate the result and then map it to the ommatidia. Since in reality it is the animal that rotates relative to the Antarium, the second method is the correct one and that is what was used.

The resulting rotIDF curves are shown in *Fig. 81*:

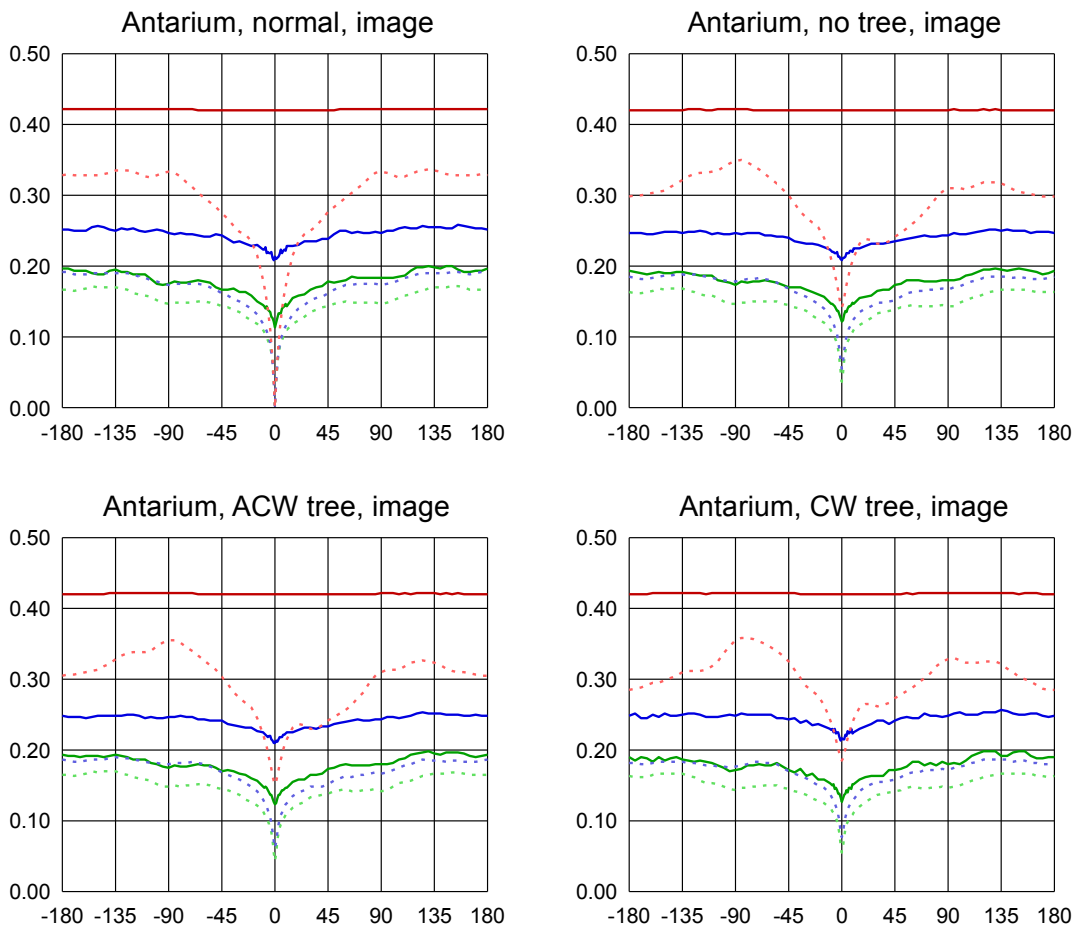


*Fig. 81: rotIDF functions for the scenes projected by the Antarium*

*Solid lines represent the rotIDF in the Antarium while the pale, dashed curves are repeated from Fig. 78 on page 238 as a comparison for what the animal could expect outdoors. The red line represents the UV component. Due to the lack of UV, on that channel there is no usable information at all, in fact it tells the animal that the projected image is completely different from what she memorised. The blue and green channels both have a global minimum at 0° rotation, but the minimum, especially on the blue channel, is very shallow at all of the four scenes.*

In the Antarium the UV channel is flat, at a level that indicates high difference. While the blue and green channels retain a global minimum at the no rotation point, this minimum is very shallow compared to what the animal would expect in her natural habitat.

Edge extraction

In personal communication A. Wystrach and F. Noël from Toulouse indicated that when they modelled the image storage capacity of the mushroom body (MB), they found that storing raw images would very quickly saturate the Kenyon cell synapses. However, if they applied a differential operator (that is, an edge detector) on the images, the modelled MB was able to store and recognise a large number of such images, significantly more than what an ant needs for foraging and homing. The Sobel filter is a very simple yet well performing edge detector. It is based on two 3 by 3 pixel kernels:

$$\boldsymbol{H}=\frac{1}{4}\cdot\boldsymbol{I}*\begin{bmatrix}+1 & +2 & +1\\ 0 & 0 & 0\\ -1 & -2 & -1\end{bmatrix}\quad \boldsymbol{V}=\frac{1}{4}\cdot\boldsymbol{I}*\begin{bmatrix}+1 & 0 & -1\\ +2 & 0 & -2\\ +1 & 0 & -1\end{bmatrix}$$

where I is the original image, * denotes convolution, while H and V are the horizontal and vertical components of the intensity gradient. If the direction of the gradient vector is not important and only its absolute value is sought, then it can be calculated as:

$$\boldsymbol{E}_{x,y}=\sqrt{\boldsymbol{H}_{x,y}^{2}+\boldsymbol{V}_{x,y}^{2}}$$

Due to the half-rectangle shift of each row of the ommatidial model, the convolution kernels were modified as follows:

$$\boldsymbol{V}=\frac{1}{8}\cdot\boldsymbol{I}*\begin{bmatrix}+1 & +1 & -1 & -1\\ & +4 & 0 & -4 & \\ +1 & +1 & -1 & -1\end{bmatrix}$$

$$\boldsymbol{H}=\frac{1}{4}\cdot\boldsymbol{I}*\begin{bmatrix}+1 & +1 & +1 & +1\\ & 0 & 0 & 0 & \\ -1 & -1 & -1 & -1\end{bmatrix}$$

then the absolute value calculation was performed as normal.

The edge detector was applied on the UV, blue and green channels separately. The results for the outdoor images are shown in *Fig. 82* and *Fig. 83*.

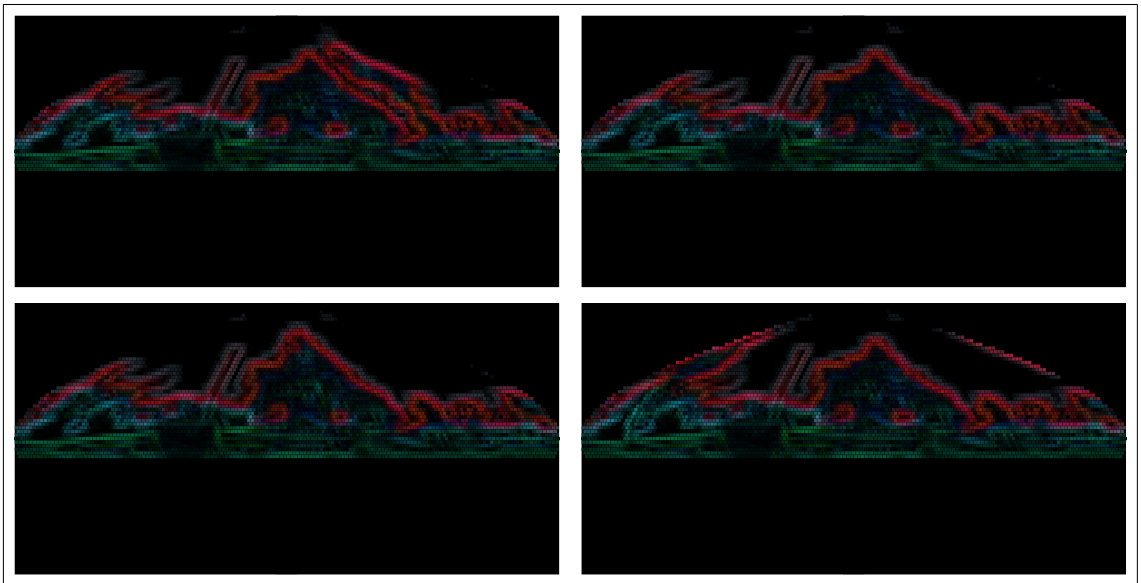

*Fig. 82: Edge detector result, outdoor scenes*

*The panels, clockwise from top-left are: tree at its normal position, tree removed, rotated clockwise and rotated anticlockwise.*

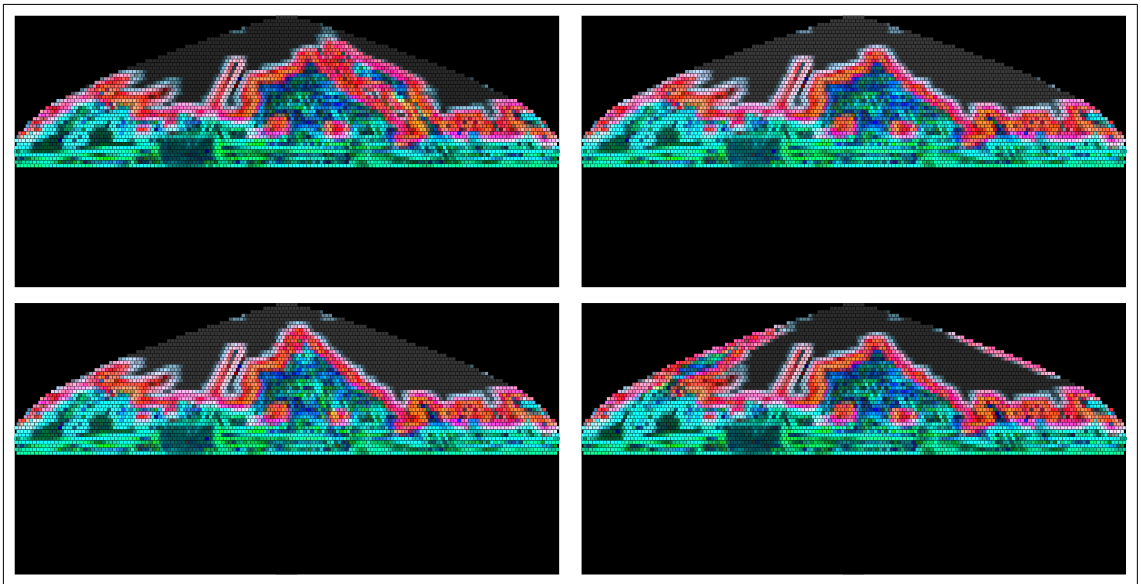

*Fig. 83: Contrast equalised edge detector output, outdoor scenes*

*These are the same panels as on Fig. 82, but contrast equalised for better visibility.*

Since the contrast equalised images, where the image was blacked out below -12˚, were fed to the edge detector, the edges are detected at and above the horizon and nearby ground features have no effect on the result.

The rotational image difference function can be calculated on the edge detector's output. The results are shown in *Fig. 84*.

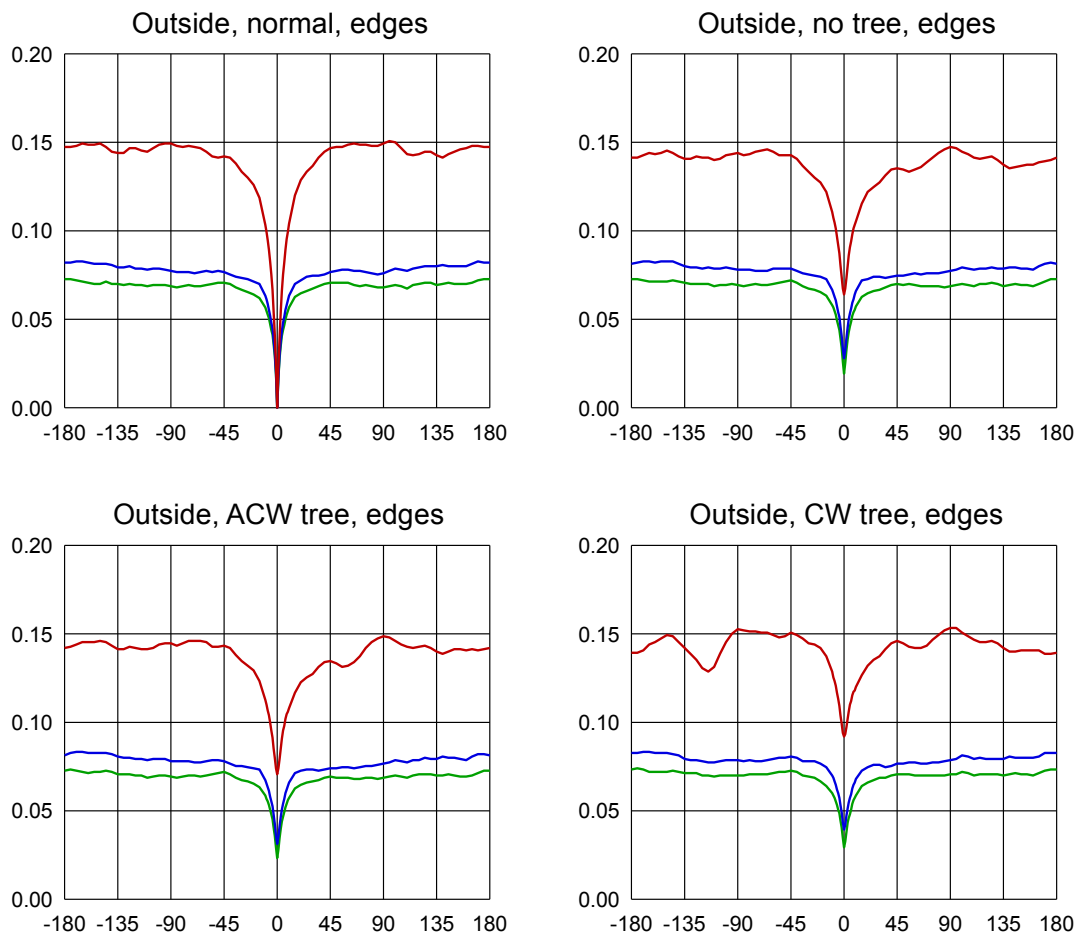


*Fig. 84: rotIDF of edges detected on the outdoor scenes*

*The red curve represents the UV component. The curves are similar to the rotIDF of the images without edge detection, but the difference levels are lower (note that on these panels the vertical axis goes from 0 to 0.2 while on Fig. 78 the scale goes up to 0.5).*

The results of applying the edge detector on the scenes projected by the Antarium are shown in *Fig 85*.

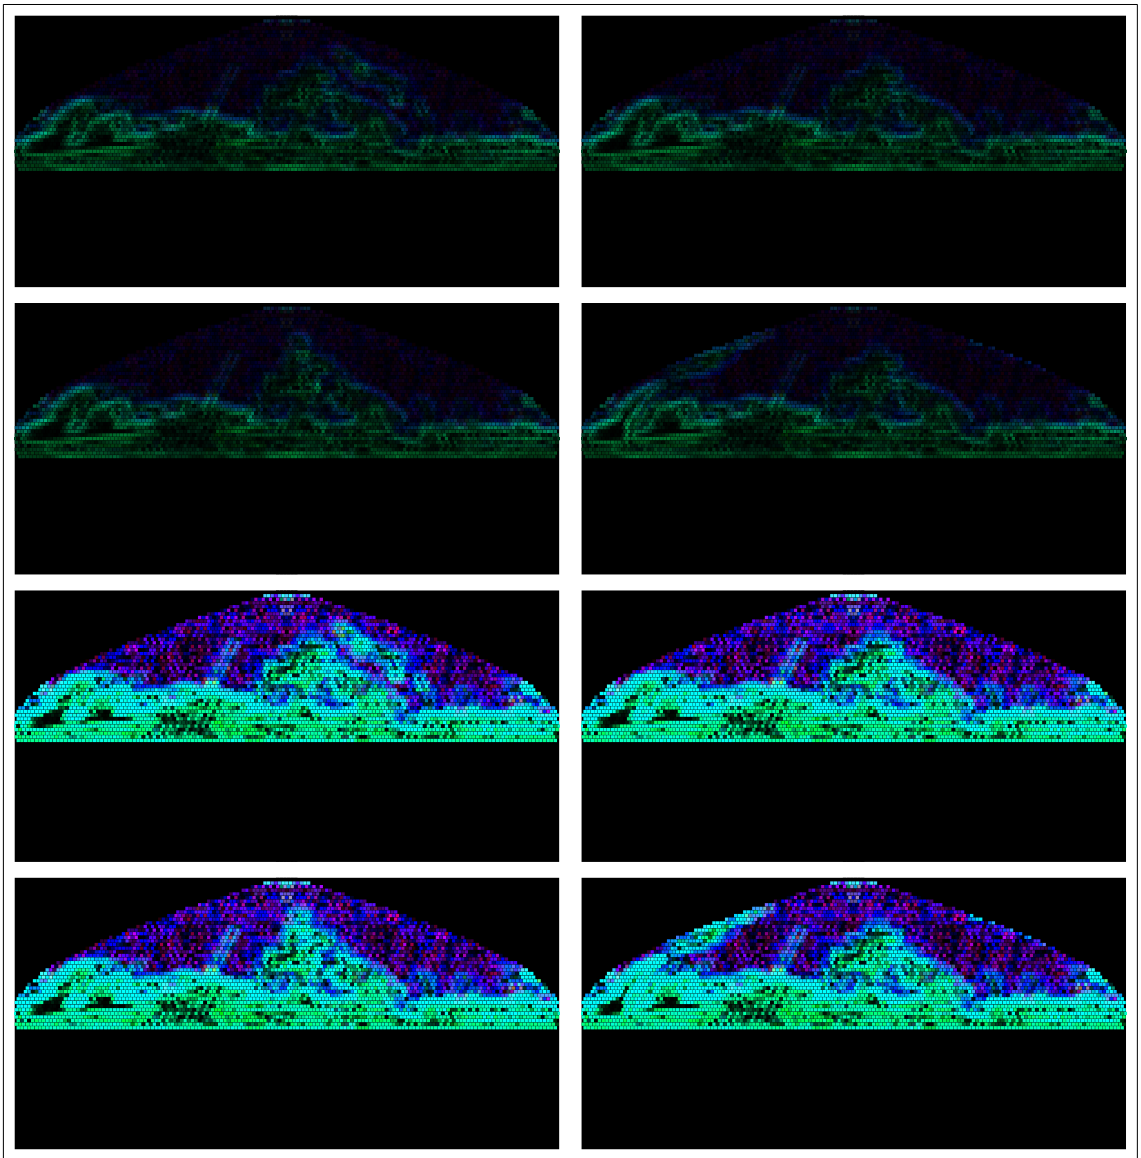

*Fig. 85: Edge detector output on the Antarium's projections*

*The top four panels, clockwise from top-left show the scenes where the tree is at its normal location, removed from the scene, rotated clockwise and rotated anticlockwise. The bottom four panels show the same, but they are contrast equalised for better visibility.*

The rotational image difference functions for the edge detector output on the scenes projected by the Antarium are shown in *Fig. 86*:

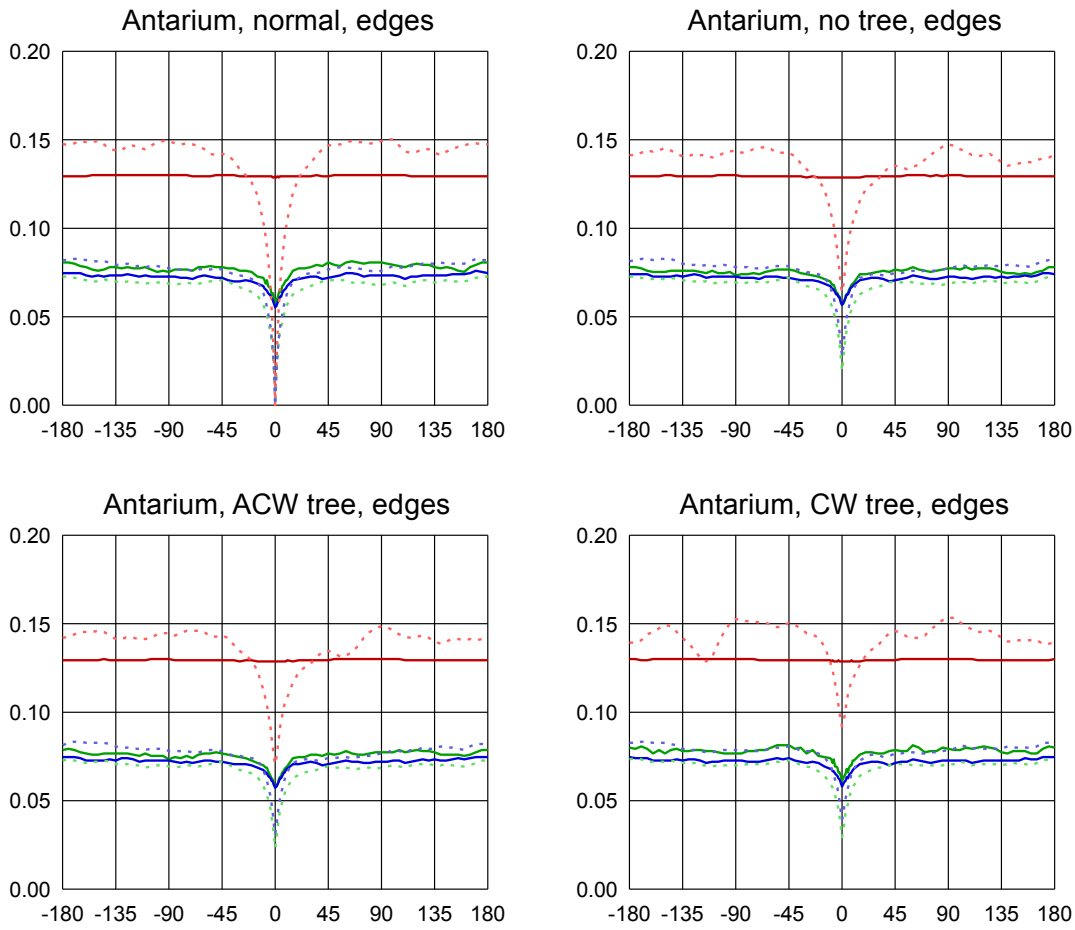


*Fig. 86: rotIDF of edge detector output on the Antarium's projections*

*The red line represents the UV. The solid lines are the rotIDF of the edges detected on the Antarium's projections. The pale dashed lines are the rotIDF on the edge detector output on the outdoor scenes, repeated from Fig. 84 for reference. The Antarium's projections provide no usable UV signal and while the blue and green channels have a global minimum at 0, it is much shallower than for the outdoor images.*

We can conclude that edge detection sharpened the global minima of the green and blue channels, but those minima are still very much reduced compared to the outdoor reference image rotIDF minima.

## Discussion

Ants which were motivated to forage were placed in the Antarium. In the virtual world depicting their natural habitat they were placed at a spot on the foraging corridor, not far from the nest. They were shown four scenes where the foraging tree was at its usual location, removed from the scene or rotated clockwise or anticlockwise around the release point.

The results indicate that approximately 25% of the animals successfully navigated in the Antarium, which is in accordance with the observations from the previous experiments. A method was devised to select the good navigators, based on their performance when confronted with the unmodified scene.

Fine-grained analysis of the paths of the navigating animals shows that when the foraging tree is absent or moved to a different location, the animals still move towards the location where the tree should be. It is in agreement with the conclusions drawn by the literature and personal observations, detailed in the Introduction section on page 223. That suggests that the animals use the entire panorama to determine their heading. However, the experiment also clearly indicated that the animals do respond to the absence or relocation of the tree, which, again, is in agreement with previous observations in a real world setting.

The central concept in panorama based navigation is the image difference function (IDF), as was discussed in the Introduction on page 71. To assess the information content of the images, a simulated ommatidial arrangement which approximates the eye of the animals was constructed and the panoramic images were mapped to it. The rotational image difference function (rotIDF) calculated on the ommatidial map indicated that the natural panorama does provide reliable navigational information even when the foraging tree is removed or displaced. The function still shows a well-defined global minimum point, although the minimum is elevated relative to the scene on which the tree is at its normal position.

In order to investigate why the majority of the animals cannot successfully navigate in the Antarium, the projected image was also mapped to the ommatidial grid, taking into account the spectral mismatch between the ants'

eyes and the Antarium as well as the particular LED positions used in the device. The resulting images showed that the UV information content of the scenes is completely lost and the visible spectrum is shifted towards the green. Furthermore, a moiré pattern between the simulated ommatidia and the device's LEDs manifested itself as noise, possibly exacerbated by the Antarium's LEDs being on the vertices of a hexagonal grid instead of at the centres of the hexagons.

As a consequence, the minima of the rotIDF of the projected images are very shallow. Under natural circumstances the minimum of the rotIDF becomes shallower and shallower as the animal gets further and further from the reference point (i.e. the nest). So it is possible that in the Antarium she becomes very uncertain of her location, triggering a random search behaviour.

Furthermore, the presence of UV in the panorama is very important. Barta and Horváth (2004) pointed out that the UV reflectance of terrestrial objects is generally low, thus UV provides higher sky-ground contrast than visible light. Differt and Möller (2015) analysed the information content of the UV-green contrast versus UV only in panoramic images and concluded that the UV on its own is almost as good as UV+green. Schultheiss e*t al* (2016a) have experimentally shown the diminishing navigational capability of the desert ant *Melophorus bagoti* when the UV content of the panorama was blocked. Stone *et al* (2014, 2016) examined the possibility of using the UV skyline for navigating in a city by sky segmentation and built a self-navigating robot using skyline extraction .

The animals in the Antarium do not receive useful UV information, partially because the UV LEDs excite their UV receptors only marginally and partially because there are very few UV LEDs compared to blue/green ones. If indeed *Myrmecia* too relies heavily on the UV content of the panorama, then it is understandable why 75% of the animals cannot navigate in it.

While it is not known what image features are extracted by the medulla and what transformations are performed by it, it is very unlikely that the animal actually memorises raw images. A very simple processing step that greatly decreases the information volume of natural images is edge detection. A basic

edge detector, based on the well-known Sobel filter was implemented and the natural scenes were processed with it. The resulting images were subjected to rotIDF analysis. The resulting functions were very similar to the functions obtained from the unfiltered images, indicating that the navigational information content of the images was not affected by the edge detection operation. That was also true for the images projected by the Antarium. The rotIDF functions derived from the edges detected in the projected images had very shallow minima, just like the images without edge detection.

## Experiment 5

The aim of this experiment was to examine the animals' reaction to colour changes in the image. The experiment was performed in January, 2019 on four non-consecutive evenings at the ANU Campus Field Station. All together 35 animals were used.

### Method

*M. pyriformis* were captured at the foraging tree on the outbound leg of their foraging trip. The animals were fed and then placed on the trackball using the magnetic mounting method.

The Antarium projected static images showing the panorama from roughly the middle of the foraging corridor. However, the colour of the images was altered in various ways.

*Fig. 87: Panoramic image from the middle of the foraging corridor.*

*The centre of the image is North. The large tree right of centre is the foraging tree, East of the camera. The nest is not visible at this resolution.*

A panoramic image was taken halfway between the nest and the foraging tree, as shown in *Fig. 87*. The image was taken around sunset. Note the lack of shadows and also the fact that the sky is not blue, more like a bluish-tinged grey. That means that it has a significant green content.

As the animals cannot see red, the red channel was removed. *Fig. 88* shows the result.

*Fig. 88: The panoramic image with the red channel removed.*

As the original image contains no UV channel, one was artificially created that covered the sky. The red channel of the image was used to represent the UV content, that's why in *Fig. 89* the sky has a purple hue. The UV intensity was set to be the same as that of the blue channel. Due to the high green content of the sky, adding the UV (on the red channel) turns it into a somewhat pinkish grey.

*Fig. 89: Panorama with UV added to the sky.*

To test whether the blue or the green colour is more important, two images were created. On one of them only the blue channel is enabled, on the other only green channel, as shown in *Fig. 90*. The UV channel was turned off. Note that the sky is not black on the green-only image, it is lit up.

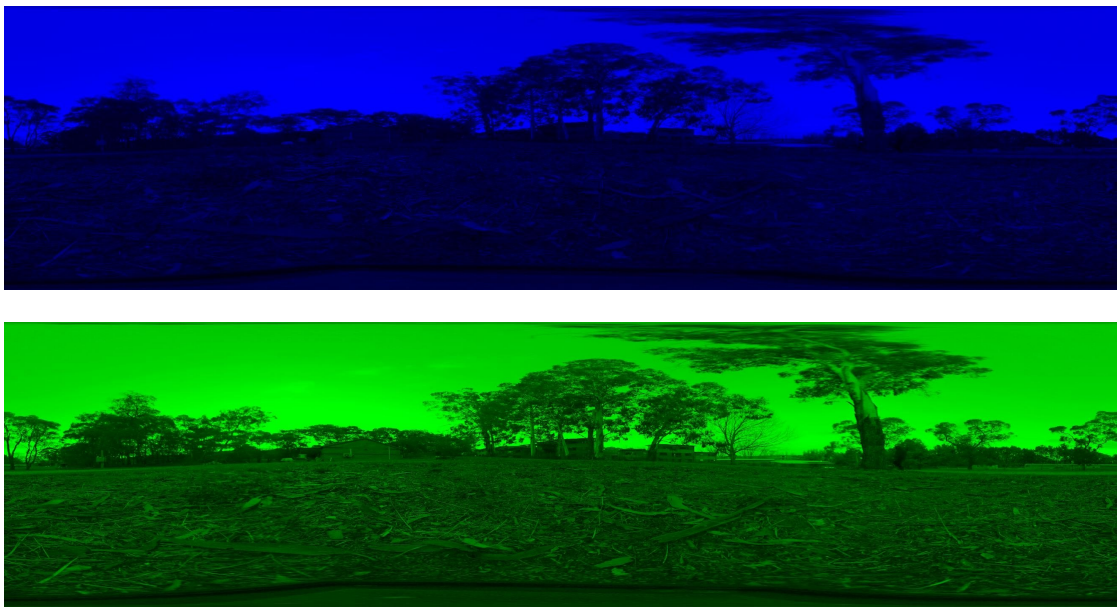

*Fig. 90: Panorama with only the blue (top) or green (bottom) channel enabled*

To see whether artificially increased blue-green contrast would make a difference, an image was generated where first the UV channel was turned off, then for each pixel the blue and green intensities were compared. Whichever was higher was retained and the other channel was set to zero for that pixel, resulting in full colour saturation, as shown in *Fig 91*.

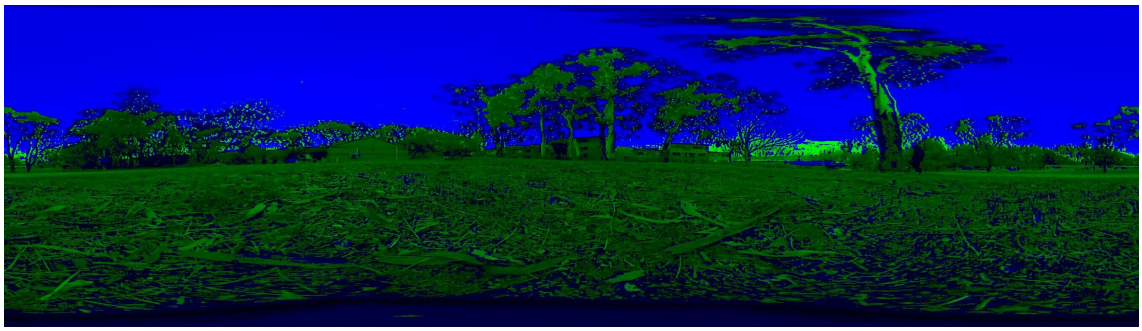

*Fig. 91: Panorama with full green-blue discrimination*

A blue channel only and a green channel only versions of that image were also created, shown in *Fig 92*.

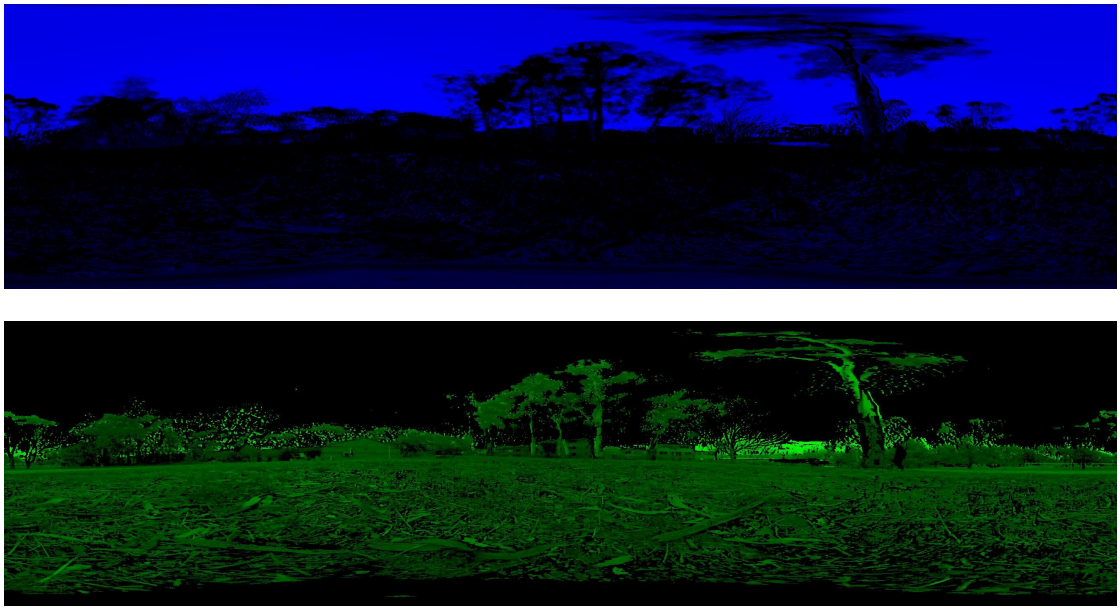

*Fig. 92: Panorama with full blue-green discrimination, blue or green channel only*

Finally, *Fig. 93* shows an image that is supposed to be extremely confusing for the animals. First, the full green-blue discrimination image was created and then the blue and the green channels were swapped:

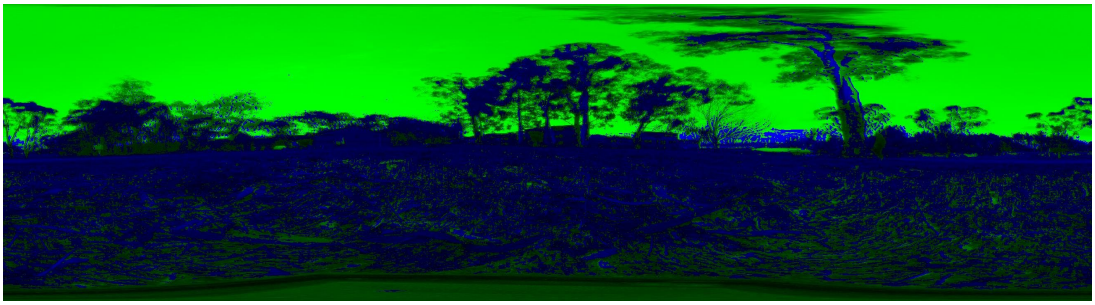

*Fig. 93: Panorama, full colour discrimination, blue-green exchanged.*

The Antarium projected all 8 scenes to the animal in randomised order. Each scene was displayed for about 40 seconds. The scene with the artificially created UV content was shown twice. The orientation of the scenes (that is, the alignment of the ball and scene coordinate systems) were chosen randomly from 0, 90, 180 or 270 degrees. When the scene changed, first the old image

was faded to black and after a 2 second pause the new image was faded in. Thus, an animal was in the Antarium for about 7 minutes.

Due to a software bug, on the first day the scene with the artificial UV content was not shown to the animals but the unaltered scene was shown twice. Consequently, on the first day only 8 scenes were shown to each animal. The bug was discovered and rectified by the second day. Because the analysis uses cumulative statistics of the scenes, it was decided that results of the first day do not need to be discarded, despite the lack of UV-containing scene data.

The animals' paths on the trackball were recorded together with the scene orientation and scene order. In addition, the animals were filmed from above for behavioural analysis to be performed later.

## Results

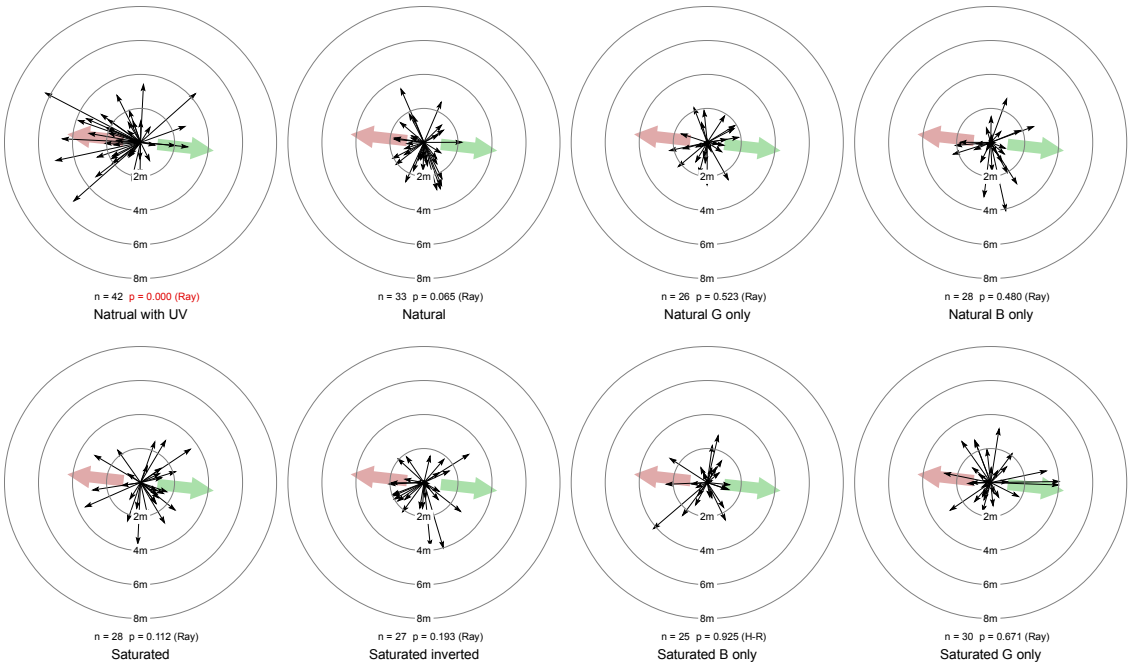


*Fig. 94: Final bearings at the end of the scene.*

*Only the scene with the artificial UV content shows statistically significant deviation from the circularly uniform. The pale red arrow indicates the nest direction, the pale green indicates the foraging tree direction.*

*Fig. 94* shows the bearing of the animals after each scene. Samples where the ant did not move or walked a very short distance are omitted. Statistical analysis against circularly uniform distribution indicates that the animals only showed orientation while facing the scene with the artificially added UV content.

The paths of the animals were also subjected to fine-grain momentary direction analysis, as on page 195 .

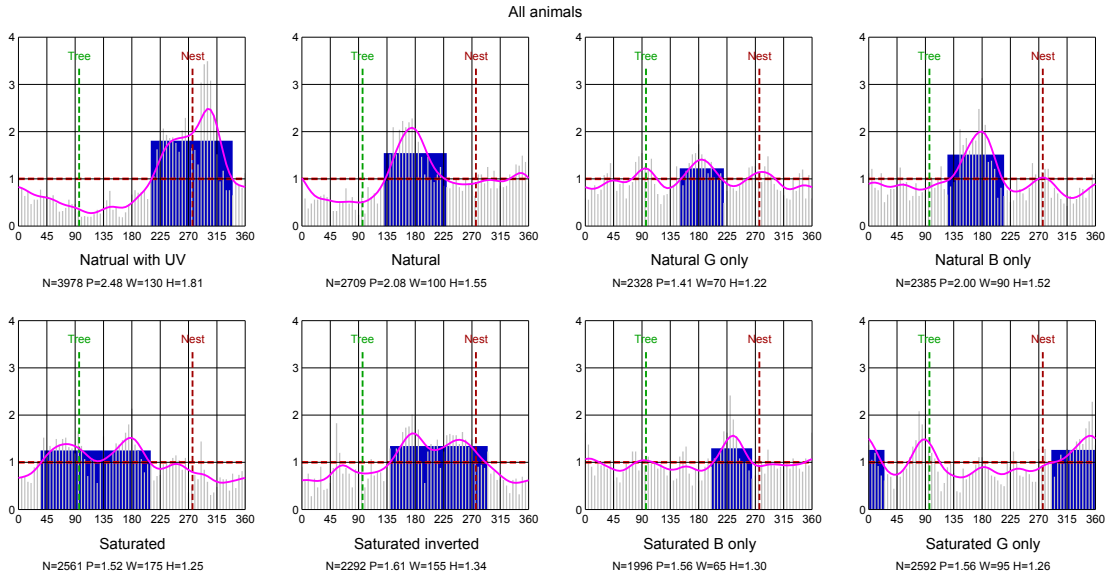


*Fig. 95: Momentary heading probability density functions for the scenes.*

*The grey vertical lines represent the 5° bins into which the vectors were sorted. The vertical axis is the frequency, relative a uniform circular distribution (marked with the red horizontal dashed line). The magenta curve is the empirical density function created by applying a Gaussian smoothing kernel on the bin counts.*

*The vertical dashed red line marked Nest indicates the direction towards the nest, the green Tree line indicates the direction to the foraging tree.*

*N is the number of vectors sorted into the bins, P is the height of the largest peak, W is the angular distance between the points where the curve crosses the uniform line left and right of the peak and H is the mean of the curve over that interval. The blue rectangle is W wide and H high.*

*Fig. 95* shows the empirical probability densities of the momentary headings for the various scenes, with all animals pooled. Like it was done for the bearings, cases where the animal did not move or travelled very little during the scene were not included in the analysis.

The scene with the added UV components shows a peak in the nest-ward direction. The other scenes have smaller peaks. In the previous experiments it was established that peak values below around 2.8 can occur for non-navigating animals while good navigating animals produce peaks above 3.6 (see page 203). Indeed, the density curves for individual animals have large peaks, indicating strong orientation. However, the animals chose very different targets, which after pooling their results flattens the overall curve. An example animal is shown in *Fig. 96*.

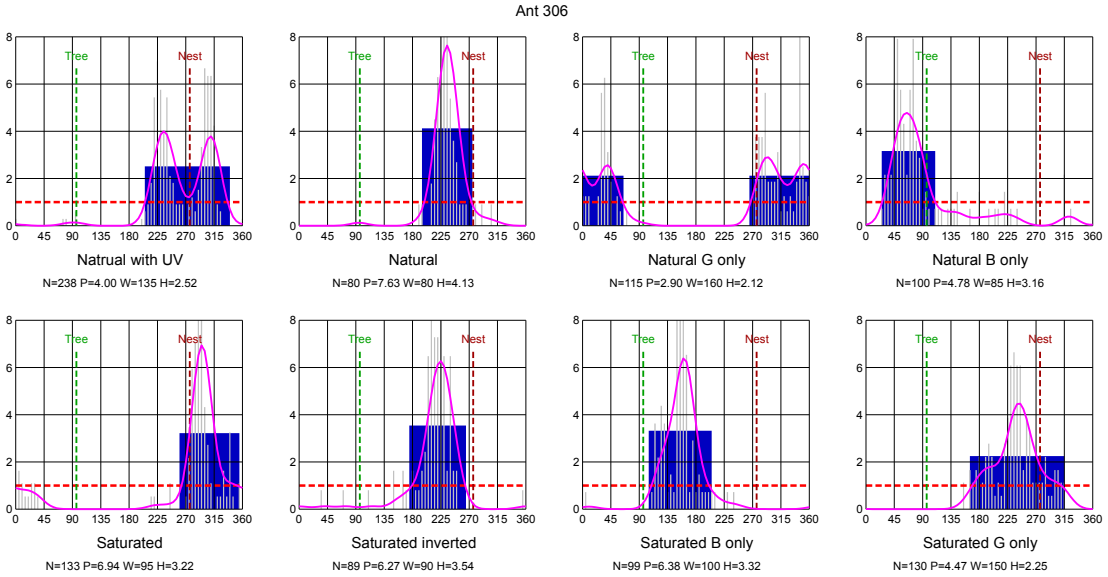


*Fig. 96: Momentary heading probability densities of an animal with strong orientation*

On all scenes but the green-only, the animal shows strong orientation. The double peak on the with-UV panel is the result of her seeing the scene twice and choosing two distinct directions (indicated by the two peaks). If the two runs were plotted individually, each peak would be twice as high.

Not all animals have shown such strong orientation. An example of an animal which was less oriented is shown on *Fig. 97*.

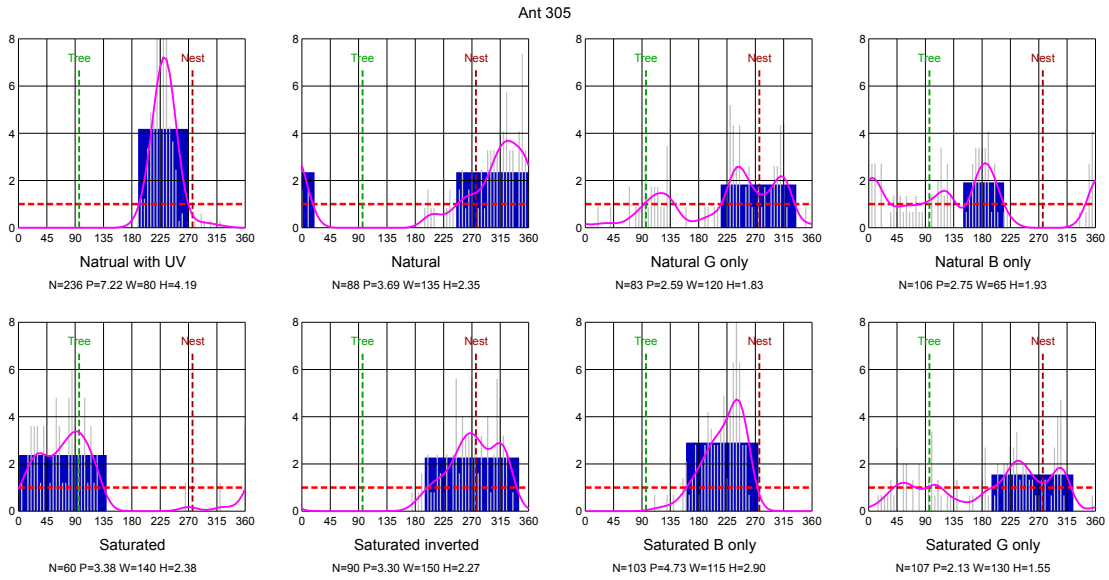


*Fig. 97: Momentary heading densities for an animal showing less orientation*

Nevertheless, looking at the individual density functions suggests that by and large most animals have shown strong orientation on the majority of the scenes. To see why pooling the data for individually oriented animals resulted in curves indicating weak orientation, the peak values and the directions of the peaks were taken as random variables and displayed as violin plots, shown in *Fig. 98*.

The peak values show a large spread, they vary between 2 and 8. Notably, for the saturated image the spread is much less, but still, the standard deviation of the values is not different from that of the other scenes. As it was established in Experiment 2, values over 3.6 indicate oriented movement and values below 2.8 are likely from random searches. Interestingly, the saturated scene evoked oriented movement for almost all animals. For all other scenes the majority of the animals have shown orientation (the median for all scenes is above 3.6), but, unlike the saturated scene, a non-negligible portion of the cohort was disoriented.

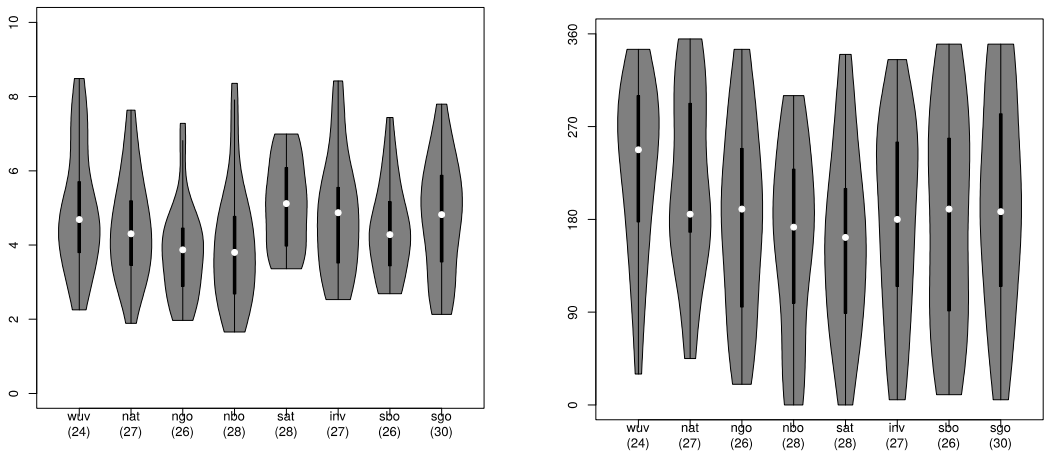


*Fig. 98: Probability density peak value and direction distributions*

*The left panel shows the peak value, the right panel the peak direction distributions for the various scenes. The scene name abbreviations are the following: wuv, with UV; nat, natural scene; ngo, natural scene green only; nbo, natural scene blue only; sat, saturated; inv, inverted saturated; sbo, saturated blue only; sgo, saturated green only. The numbers in parenthesis under the scene code are the number of samples for the scene.*

*The thick black bar is the standard deviation around the mean, the white dot is the median and the width of the grey outline shows the empirical probability density function.*

Considering the direction panel, it is obvious that even though the animals were oriented, the direction they followed was almost random. That observation agrees with the statistical results for the bearings in *Fig. 94* on page 255. Just like on that figure, on the violin plots the scene with the added UV is an outlier. The mean direction for all scenes is around 180°, which is expected if uniform circular data is processed as a linear quantity between 0 and 360. However, the mean for the UV scene is around 250°, indicating that the distribution was not uniform, again agreeing with the bearing analysis result.

Further analysis can be performed on the peak values and directions to see whether there are statistically significant differences between the scenes. The results are shown in *Table 13* and *Table 14*.

| | wuv | nat | ngo | nbo | sat | inv | sbo | sgo |
|---|---|---|---|---|---|---|---|---|
| **wuv** | | 1.000 | 1.000 | 1.000 | 0.050 | 1.000 | 1.000 | 1.000 |
| **nat** | 0.100 | | 1.000 | 1.000 | 1.000 | 1.000 | 1.000 | 1.000 |
| **ngo** | 0.100 | 0.100 | | 1.000 | 1.000 | 1.000 | 1.000 | 1.000 |
| **nbo** | 0.100 | 0.100 | 0.100 | | 1.000 | 1.000 | 1.000 | 1.000 |
| **sat** | 0.010 | 0.100 | 0.100 | 0.100 | | 1.000 | 1.000 | 1.000 |
| **inv** | 0.100 | 0.100 | 0.100 | 0.100 | 0.100 | | 1.000 | 1.000 |
| **sbo** | 0.100 | 0.100 | 0.100 | 0.100 | 0.100 | 0.100 | | 1.000 |
| **sgo** | 0.100 | 0.100 | 0.100 | 0.100 | 0.100 | 0.100 | 0.100 | |

*Table 13: Pairwise statistical comparison of peak direction distributions*

*The peak direction data for the scenes were pairwise compared using the Watson $U^2$ nonparametric test for circular data. The R implementation of that test returns a p-value range; the upper triangle of the matrix indicates the maximum of the range for the row-column pair, the lower triangle is the minimum of the range. The only statistically significant difference is between the scene with UV added and the one with the fully saturated blue and green channels.*

*The scene name abbreviations are the following: wuv, with UV; nat, natural scene; ngo, natural scene green only; nbo, natural scene blue only; sat, saturated; inv, inverted saturated; sbo, saturated blue only; sgo, saturated green only.*

| | wuv | nat | ngo | nbo | sat | inv | sbo | sgo |
|---|---|---|---|---|---|---|---|---|
| **wuv** | | 0.692 | 0.027 | 0.121 | 0.121 | 0.661 | 0.218 | 0.646 |
| **nat** | 0.373 | | 0.142 | 0.151 | 0.106 | 0.754 | 0.994 | 0.642 |
| **ngo** | 0.006 | 0.063 | | 0.749 | 0.000 | 0.005 | 0.172 | 0.004 |
| **nbo** | 0.033 | 0.124 | 0.764 | | 0.026 | 0.131 | 0.181 | 0.091 |
| **sat** | 0.542 | 0.081 | 0.001 | 0.003 | | 0.436 | 0.068 | 0.332 |
| **inv** | 0.963 | 0.571 | 0.017 | 0.040 | 0.319 | | 0.492 | 0.864 |
| **sbo** | 0.242 | 0.909 | 0.060 | 0.122 | 0.068 | 0.474 | | 0.292 |
| **sgo** | 0.870 | 0.541 | 0.013 | 0.040 | 0.472 | 0.968 | 0.382 | |

*Table 14: Pairwise statistical comparison of peak value distributions*

*The upper triangle of the matrix contains the p-values of the non-parametric Mann-Whitney test (a.k.a. Wilcoxon test) and the lower triangle that of the Kolmogorov-Smirnov test. Red marks statistically significant difference identified by both tests (after applying the Bonferroni correction), purple marks values where only the K-S test, which is sensitive to the shape of the distribution, could tell the two distributions apart. The scene name abbreviations are the same as in Table 13.*

While the statistical tests found statistically significant differences of the peak value distributions of certain scene pairs, overall, most pairings did not differ enough to be picked up by the tests. In particular, the peak levels of the natural scene were not statistically significantly different from the levels collected in any

of the modified scenes. That seems to indicate that the animals could not extract usable navigational information from any of the scenes.

To understand why that is the case, the rotational image difference (rotIDF) functions were calculated for the scenes, as seen by the animals in the Antarium. The image processing details and the rotIDF calculations were explained in the Rotational image difference function analysis section on page 233; for the sake of brevity those details are not repeated here.

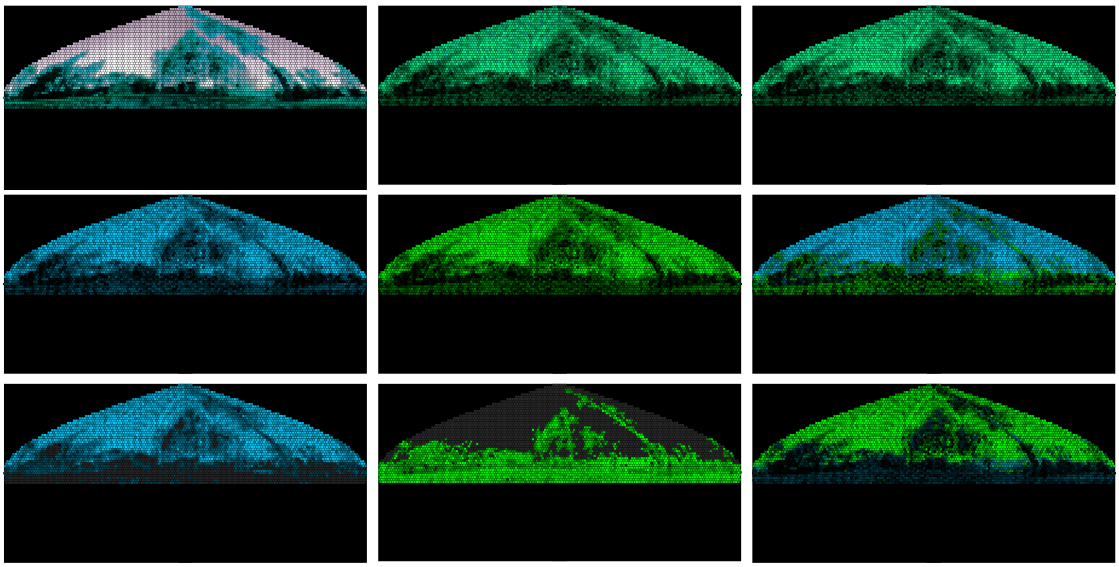

*Fig. 99: Scenes mapped to the ommatidial model*

*The top-left image is the UV-augmented natural scene, as seen outdoors. The rest of the images, going left to right and top to bottom, show what the animals saw in the Antarium: natural scene with added UV, natural scene without UV, blue only, green only, saturated, saturated blue only, saturated green only and saturated inverted. The images have been normalised using histogram equalisation.*

*The bottom halves of the images, below the horizon line, have been blacked out as it is assumed that no useful long-range navigational information can be extracted from nearby ground features.*

*Fig. 99* shows the image that the animal might have been familiar with (the UV-augmented outdoor image mapped to the ommatidial arrangement) and the scenes that she faced in the Antarium, which are the colour-manipulated outdoor images mapped to the Antarium's LEDs then, in turn, mapped to the ommatidial matrix and colour-transformed to cater for the mismatch between the LED emission spectra and the animals' photoreceptor spectral sensitivities.

After that mapping there is no perceivable difference between the with-UV and without-UV images (top row, centre and right) or between the blue-only and saturated blue-only (left, centre and bottom rows). On all Antarium images the moiré noise is prominent on the sky area.

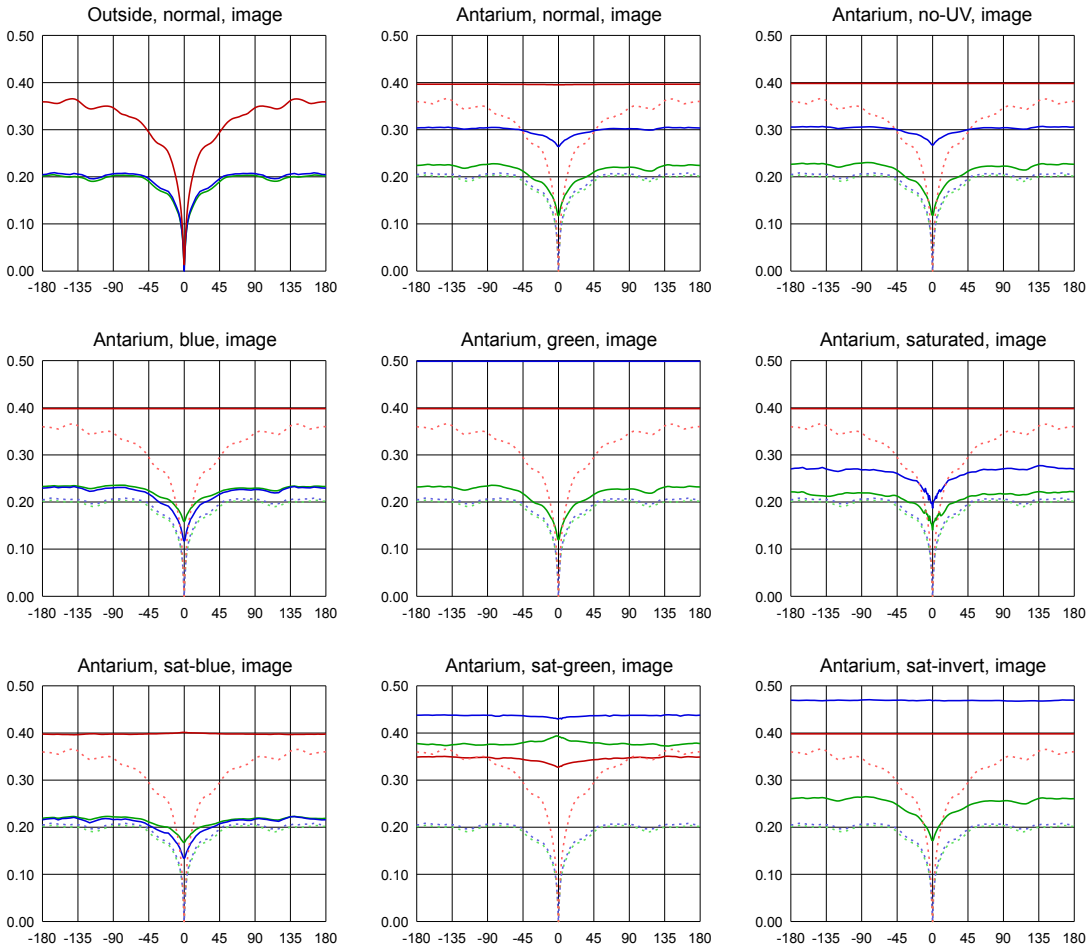


*Fig. 100: Rotational image difference functions of the scenes seen by the animals*

*The outside image (top-left) shows the image difference function of the outdoor scene, augmented with UV. The other panels show the rotIDFs of the scenes shown to the animal in the Antarium. On those panels the dashed lines are that of the rotIDF of the outdoor image, repeated from the top-left panel, for reference purposes.*

*Fig. 100* shows the rotIDF functions if the animal compares the scene projected in the Antarium to the memorised outdoor view. The rotIDF is calculated for the UV, blue and green channels separately. All images were normalised using histogram equalisation. It should be noted that the rotIDF value of 0 indicates a

perfect match between the two images while 0.5 is the expected value of two completely unrelated images.

It is obvious that the UV channel, which supplies the most salient information on the outdoor scene provides no information whatsoever in the Antarium. In most Antarium scenes the blue channel is also devoid of navigational content, the minimum of that channel is not much below its maximum and even at the minimum point it is higher than what the animal can expect as worst-case value when she is rotating outdoors.

The only channel with any useful navigational content is the green channel, which, for most scenes, is roughly at the same level as the outdoor scene green channel, except that its minimum is very shallow compared to that of the outdoor scene. In fact, the saturated green-only scene has a positive peak at the point of no rotation, although at a rotIDF value around 0.4 the actual shape of the curve matters very little, as that range indicates no similarity between the images.

Comparing the edge detector output (see *Edge extraction* on page 243) rotIDF shows that the projected scenes are similarly parsimonious with navigational information content.

The edge detector images are omitted, as they show no particularly interesting features but the resulting rotational image difference functions are shown in *Fig. 101*.

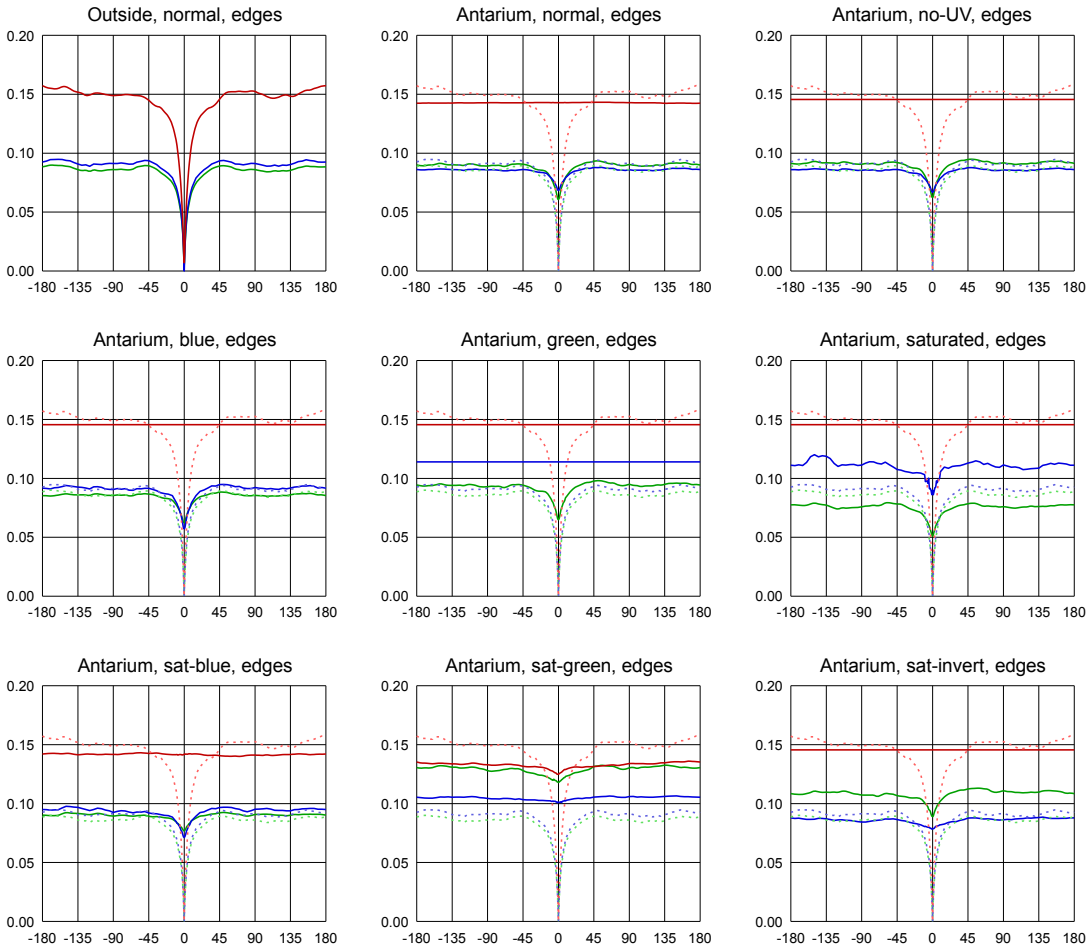


*Fig. 101: Rotational image difference functions of edge-detector processed scenes*

*The outside image (top-left) shows the image difference function of the outdoor scene, augmented with UV. The other panels show the rotIDFs of the scenes shown to the animal in the Antarium. On those panels the dashed lines are that of the rotIDF of the outdoor image, repeated from the top-left panel, for reference purposes.*

The rotIDF functions for the edge detector output are very similar to the rotIDFs obtained by processing the raw images. The UV channel gives the most prominent rotational information to the animal in her natural environment while that channel is completely silent in the Antarium.

The blue channel provides some information when the Antarium projects images that exercise the blue LEDs, but the minimum of the function is much shallower than in the outdoor case. Obviously, for scenes where the blue channel is off or it is driven with incorrect information, the animal receives no rotational information from the channel.

The green channel behaves similarly to the blue one, but because the green photoreceptors get excited by both the blue and the green LEDs, the green channel supplies some rotational information even when the green LEDs are turned off and only the blue ones are active. Nevertheless, the very shallow minimum of the channel still makes the rotational information unreliable.

Thus, the rotational image difference function analysis explains why the animals cannot obtain navigational information from the scenes. However, there remain two issues which cannot be explained.

One, at the majority of the scenes the animals have shown strong orientation, their heading remained within a reasonably small range. If the animal is robbed all navigational information, she is expected to perform a random search, frequently changing direction. It is known from previous experiments that sometimes they keep running at the same direction even when they lack navigational data, but that is not what most of them do. Yet in this experiment that was the norm and the reason for that is not known.

The second discrepancy is the stark contrast between the bearing distribution of the animals when they face the UV-augmented natural scene and the same scene without UV. As it is obvious from *Fig. 99, 100* and *101* there is no discernible difference between the perceived images or the corresponding rotIDF functions. Yet, the animals have shown a statistically significant, nest-directed bias in their bearing when the UV-augmented scene was projected while their bearings had no statistically detectable deviation from circularly uniform distribution when they faced a UV-less image. There currently is no explanation to this discrepancy.

## Discussion

In this experiment the animals were projected a familiar scene which showed the panoramic view from the middle of their foraging corridor. Since the animals were fed, if they recognise the scene then their expected behaviour is to walk back to the nest, or, possibly, continue their trip to the foraging tree.

While the projected image has always shown the same scene, its colour composition was changed in various ways. Each animal faced all colour

variations in the Antarium, in randomised order. Between scene changes the Antarium was darkened so the animals were not subjected to abrupt visual changes.

Bearing analysis of the paths indicated that the animals had a nest-ward tendency when the projected scene was unchanged from its natural colours and UV content was added to the scene. When the animals faced the natural colours without UV or any of the scenes where the colours were manipulated, their bearing distribution was not different from a circularly uniform random distribution.

Fine-grained analysis of the momentary headings of the animals throughout their paths have shown that most animals have shown strong orientation at the majority of the scenes they saw, but there was no correlation between the scene and their preferred direction; in fact the direction seemed to be randomly chosen by the animal every time she saw a new scene.

No solid statistical difference could be established between the behaviour at the various scenes. Rotational image difference function analysis indicated that due to the spectral mismatch between the Antarium's LEDs and the spectral sensitivities of the photoreceptors of the animals the projected visual provided very little, if any, navigational information to the animal. The presence of visual noise due to moiré patterns between the Antarium LED arrangement and the animal's eyes might also be a factor.

While the rotIDF analysis explained why the animals could not navigate in the visual environment they were placed into, it did not explain why the animals showed strong orientation even without visual references or why the animals had a nest-directed bearing distribution at one scene, which had no more navigational information content than any other. These questions remain open.

## Summary

Two experiments were performed using the Antarium that could only be done with great difficulty without a virtual reality arena.

The first experiment tested the animals' reaction to removing the foraging tree from the scene as well as to the tree being displaced in the scene. Although

removal of trees sometimes occur in real life, an experiment involving the re-location of a 20m high Eucalyptus tree is not feasible in practice.

The results were in line with real-life observations when trees were felled near the foraging routes of animals. Furthermore, the experiment supplied supporting evidence for the assertion that only about 25% of the animals can extract navigational information from the images projected by the Antarium in its current form.

The second experiment intended to test the navigational ability of the animals in an environment where the colours of a natural, familiar image was manipulated. This experiment would be very hard to perform in real life, as changing the sky to green while turning green foliage blue is simply not feasible in an outdoor setting. Unfortunately, the results could not answer the question, because the incorrect LED wavelengths rendered the projected images all but unusable from the animals' point of view. However, the experiment was still useful, as it offered two surprising and as yet not explained observations:

The animals' behaviour when shown a natural scene without any colour manipulation was significantly different depending on whether the UV channel of the Antarium was activated or turned off. The Antarium's UV LEDs are very sparse (1 UV pixel for every 40 blue/green pixels) and its 405nm wavelength is barely detectable by the UV sensitive photoreceptors of the animals (Ogawa et al., 2015). Nevertheless, with those LEDs active, the animals' bearings were statistically significantly different from the case when the UV LEDs were turned off.

The other surprising result was that the animals have shown strong orientation during their time in the Antarium. An ant with no visual reference is expected to perform a random search. Sometimes animals at an unfamiliar location choose to run in a straight line, but most of the time they wander around. However, during this experiment most animals chose a random direction and kept their heading during the period of scene projection. There is no explanation for that behaviour.

To summarise, Experiment 4 demonstrated that the Antarium, even in its current form, can be used to perform an experiment that is very hard to do in the natural environment of the animal. Experiment 5, which would be very hard to devise in an outdoor setting, could not achieve its intended goals due to the limitations of the Antarium. However, it presented two intriguing questions, the answers to which have not been found yet and further experiments are needed to examine them in more detail.

# The design of Antarium Mk-II

After the shortcomings of the Antarium were discovered, it was decided that a new device, with all known problems of the Antarium rectified, should be built. The budget for the device was doubled compared to the existing device, so that budgetary constraints affect the design choices less.

In this chapter I will describe the design of the new device and show why the solutions are superior to those used in the existing Antarium. The new Antarium is currently in the design and procurement phase and is expected to be built before the Australian experimental season starts in late 2021.

The most prominent problems with the existing device are the incorrect wavelengths of the LEDs, the low number of UV pixels, the large parallax of the LEDs and the power mismatch of the LED outputs at different wavelengths. In addition to fixing these problems, it would be beneficial if the device were capable of a higher frame rate with lower latency than before. These problems were addressed by using custom designed LEDs, adding UV to each pixel, changing the geometry of the device, using custom optics and re-arranging the electronics of the panels. The rest of the chapter will detail the solutions and validate that they will indeed improve the device.

## Colour matching and LED selection

The new Antarium's LEDs need to match the photoreceptor spectral sensitivities of *Myrmecia* ants. Ogawa *et al* (2015) identified three spectrally different photoreceptors using electrophysiological measurements. They examined both the diurnal (*M. croslandi*) and the nocturnal (*M. vindex*) *Myrmecia* species. The spectral sensitivity curves were measured with both intra-cellular and electro-retinogram recordings for the night active *M. vindex* and with electroretinogram for the day active *M. croslandi*. The results indicated that the sensitivity curves are, although not identical, but very similar for both species.

For the LED selection the sensitivity curves obtained by intracellular recordings and published in that paper were used as a baseline. They are shown in *Fig. 102*.

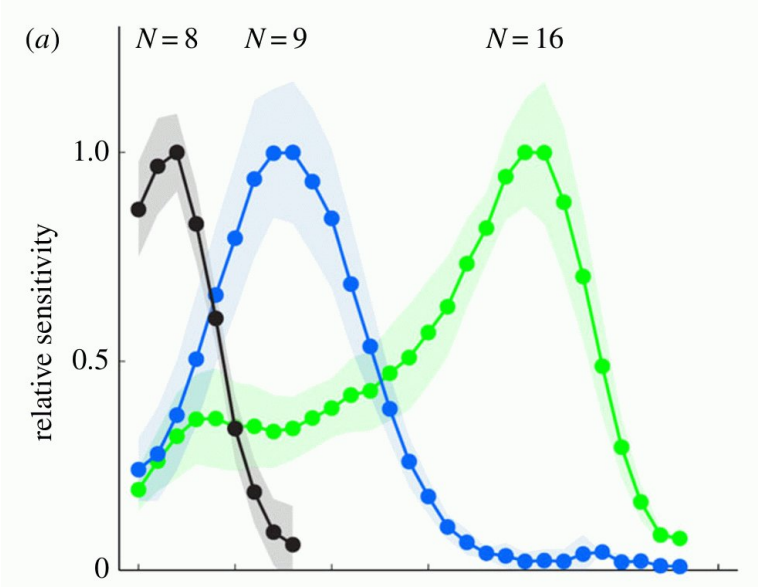


*Fig. 102: Photoreceptor spectral sensitivities*

*The spectral sensitivities of the UV, blue and green photoreceptors of M. vindex, obtained using intracellular recording. The number above each curve is the number of recoded photoreceptor cells. Shaded areas around the curves indicate the standard deviation. The measured wavelength range is 350nm to 630nm. Adapted from (Ogawa et al., 2015).*

The sensitivity was not measured below 340nm. Therefore, for the present purpose of modelling how LED spectral properties are matched to spectral sensitivities the UV sensitive curve was artificially extended towards the shorter wavelengths. It also should be noted that the text in the article claims that the peak sensitivity of the blue receptor is at 450nm, the figure showing the sample points peaks at 425nm. In this work the sample points of the figure were used. Based on the sample points a continuous smooth curve was created, using a cubic interpolation method that does not generate overshoots (Kruger, 2002). The sensitivity curves used for all further calculations are shown in *Fig. 103*.

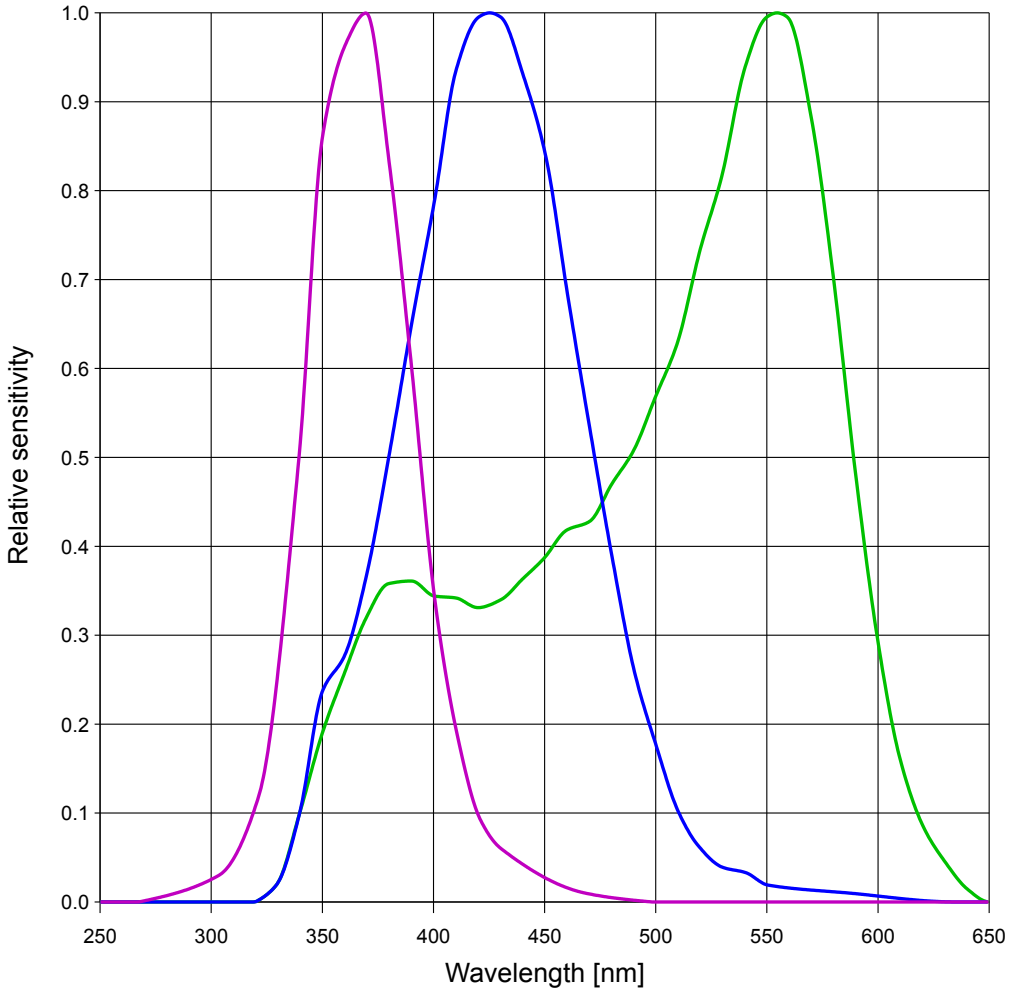


*Fig. 103: Sensitivity curves used for the modelling*

To assess the extent to which the LEDs are matched to the spectral sensitivities of *Myrmecia* ants and whether the LEDs that we are planning to use can reproduce most or all the colours the animal can encounter in real life, I modelled the ants' colour space and how LEDs and natural spectra are represented in it. Before going into details, a couple of concepts should be clarified.

Light is electromagnetic radiation. Visible (i.e. that humans can see) light occupies about one octave of the spectrum, roughly between the wavelengths of 400nm and 800nm (750THz to 375THz). Light can be described using

*radiometric* and *spectrophotometric* terms that describe physical parameters of any electromagnetic radiation, visible or otherwise. However, for practical purposes, such as mixing paint, designing coloured lights and similar, we must take the peculiarities of human vision into account and use *photometry* and *colorimetry* instead. The similarity of the radiometric and photometric terms is a source of much confusion. For example, *luminosity* is a radiometric concept, describing the total electromagnetic radiant power of an object, measured in watts. *Luminous flux*, on the other hand, belongs to photometry, for it is the *perceived* brightness of an object, measured in lumens. As such, it depends on the characteristics of the human eye: a green object will be perceived much brighter than a red one with the same radiant power (luminosity), because the human eye is more sensitive to green than to red. Fortunately, the terms used by spectrophotometry and colorimetry are quite distinct. They both deal with the spectral composition of the electromagnetic radiation. Spectrophotometry describes the spectrum in absolute terms while colorimetry focuses on the perceived colour that a given spectral composition evokes.

Photometry focuses on the brightness of things and is interested in their colours only as much as the perceived brightness depends on the colour.

Colorimetry deals with colours that we see. It would be independent of photometry if colour were not affected by intensity – but it is. For example, the colour brown if basically the same as the colour orange, except lower intensity. In addition, humans have scotopic (dim light) and photopic (bright light) vision. The former is mediated by the rods in the retina and is characterised by lacking colour perception, thus out of the scope of colorimetry[20]. Photopic vision is delivered by the cones in the retina. There are three kinds of cones; one most sensitive to red light, another to green and the third to blue. That is, humans have tri-chromatic vision. *Myrmecia* are also tri-chromatic so colorimetric terms are suitable to describe their colour vision[21].

---

20 *While we lose our colour vision in dim light, many animals do not. For example, the nocturnal moth,* Deilephia elpenor *has been shown to have colour vision even in moonless starlight conditions (Kelber et al., 2002).*

21 *Most of what will be said about chromaticity can also be found in, for example, (Kelber et al., 2003) and (Vasas et al., 2019). The colour vision of ants was reviewed in (Aksoy and Camlitepe, 2018).*

Assuming that the three kinds of receptors can be independently excited by light between nothing and a maximum excitation level (denoted as 0 and 1), all possible lights that the animal can sense can be represented by a 3-dimensional orthonormal coordinate system, the three positive axes being the excitation levels of the 3 receptors as a response to stimulus. That is shown in *Fig. 104*.

The length of a vector in the all-positive octant is the *luminance* of the stimulus light while its direction is the perceived *chromaticity*[22]; these terms are often colloquially shortened to *luma* and *chroma* in colour signal processing circles.

Because the excitation levels, by definition, are between 0 and 1, every possible light stimulus that the animal can theoretically perceive is enclosed in a unit length cube. Since we are interested in the colour reproduction, that is, the chromaticity, we will ignore the luminance for now. The (1,0,0), (0,1,0) and (0,0,1) points (the unit excitations of a single receptor) span a plane and define an equilateral triangle on it. A line between any point in the cube and the origin intersects the plane within the triangle, thus the triangle covers all possible chromaticities.

22 *Luminance and chromaticity are terms normally used in the context of human light perception and are subject to not only the physiology but also the psychology of vision. In the case of ants the psychological component, which can only be established by evaluating the subjective assessments of stimuli by many individuals, will not be taken into account.*

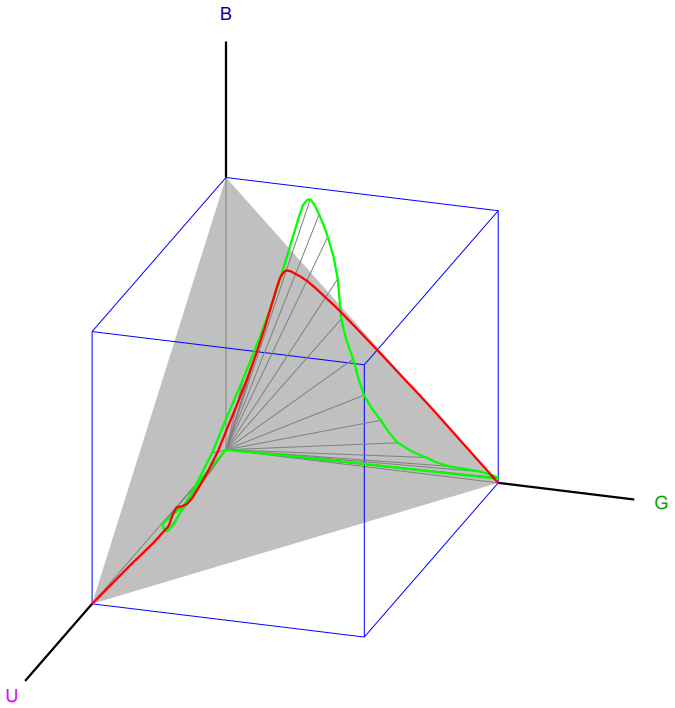


*Fig. 104: Colour cube*

*The three axes represent the individual receptor excitations, between 0 and 1. The blue cube contains all possible excitation combinations. A vector from the origin to any point within the cube represents a receptor excitation combination; the length of a vector is the luminance and its direction is the chromaticity. The grey triangle on the plane spanned by its vertices contains all vector directions within the cube and thus all possible chromaticities. If we trace the excitation by unit intensity monochromatic light as a function of wavelength we get a 3D curve. In the figure that is shown in green. We can then project that curve to the triangle, which is shown in red. Any perceivable spectrum is a linear combination of the monochromatic components, thus the gamut can be determined from that curve.*

Now we need to introduce another term, the colour *gamut*[23]. In colorimetry it means the full range of chromaticities that can be perceived or represented by some system. Since the triangle represents all chromaticities, a colour gamut always lies within the triangle. To determine the colour gamut of *Myrmecia* eyes, we need to trace the excitation combinations for uniform intensity monochromatic light for all perceivable wavelengths. In *Fig. 104* that is

23 *The word gamut comes from the Greek-Latin 'gamma ut' and was used to describe the first note of a (medieval) musical scale. Over time the meaning changed to the scale as a whole. By extension it is used to describe a complete range, as in 'a whole gamut of emotions'.*

represented by the green curve. We then project that curve to the chromaticity triangle, which in the figure is the red curve. From now on we can focus on the chromaticity triangle and work in two dimensions instead of three.

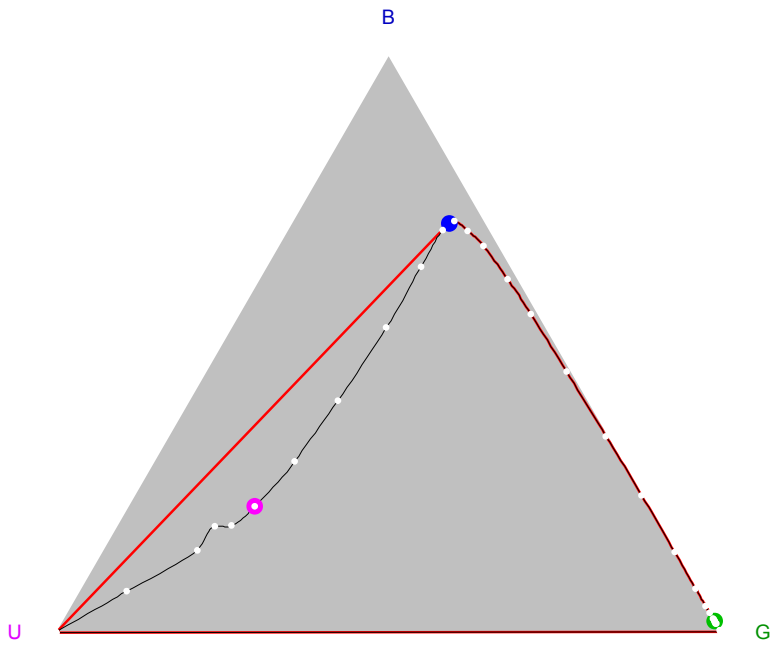


*Fig. 105: Colour gamut of the ant*

*The black curve traces the chromaticity of a monochromatic stimulus from 250nm to 600nm. The white dots are 10nm apart. The red curve encloses the colour gamut. The purple, blue and green dots show the peak sensitivities of the UV, blue and green photoreceptors (370, 425 and 550nm, respectively).*

The monochromatic excitation projected to the colour triangle, which was the red curve in *Fig. 104* is shown as a black curve in *Fig. 105.*

Any perceivable colour is the linear combination of monochromatic components (possibly an infinite number of them), with non-negative coefficients. Any linear combination (with non-negative coefficients) of two chromaticity points on the triangle will lie on the line connecting the two points. Therefore, where the monochromatic curve is convex, the resulting point will be inside the area enclosed by the curve. Had it been convex everywhere, it would have enclosed the colour gamut. But, on the UV-B side of the triangle the curve is concave and linear combinations of points on the monochromatic curve can fall on the outside. That means that we need to extend the gamut area to form a convex

hull so that a linear combination of any two points enclosed by it will now remain inside it. The red curve in *Fig. 105* shows the result of that process. That curve encloses all chromaticities that the animal can possibly see. The question arises: why are there chromaticities that the ant cannot see, even though in theory with 3 independent receptors she should be able to see them? The answer is, the receptors are not independent. There is no wavelength or combination of wavelengths that excites the blue receptor but neither the UV nor the green (see *Fig. 103* on page 271*)*. Therefore, areas of the chromaticity triangle that would require the animal to have a blue-only or mostly-blue excitation patterns are not perceivable for her, due to the physiology of her photoreceptors. *Myrmecia* are not alone with this, humans also have a limited colour gamut, as does any tri-chromatic system where the long and short wavelength sensitivity curves overlap and thus the middle wavelength sensor cannot be excited on its own. In fact, if we could somehow excite the green receptors of a human eye without also exciting the blue and red ones, the subject would experience a new colour that, literally, is “impossibly green”.

We want the arena to be able to reconstruct as large an area of the gamut as possible. If we had LEDs which emitted monochromatic lights at the bottom two vertices of the triangle and one at the apex of the gamut curve, then we could reconstruct almost the entire gamut. Unfortunately, that is not feasible. On the one hand, we want the LEDs to be effective, that is, close to the peak sensitivities of the receptors. Furthermore, the emission wavelengths of LEDs depend on the discrete energy levels of excited states of electrons in a doped semiconductor. Those levels are determined by quantum physics, and not all wavelengths are available. Also for quantum physical reasons the efficacy of an LED depends on its colour and for some colours it is just not practical to make an LED. Furthermore, LEDs are not really monochromatic. Their emission spectrum is spread out due to, again, quantum physical reasons. Last, but not least, the material that the light emitting crystal is enclosed in, has its own absorption spectrum and may not be transparent to all wavelengths. Most common plastics that are transparent to visible light are opaque to UV.

I approached several LED manufacturers to make custom LEDs for the Antarium. The large manufacturers simply ignored me or sent back an astronomical dollar figure for merely considering the issue. However, some smaller Chinese manufacturers responded and even sent me samples. I measured those and after careful consideration I settled with Shenzen Starsealand Opto Electronic Co. Ltd.

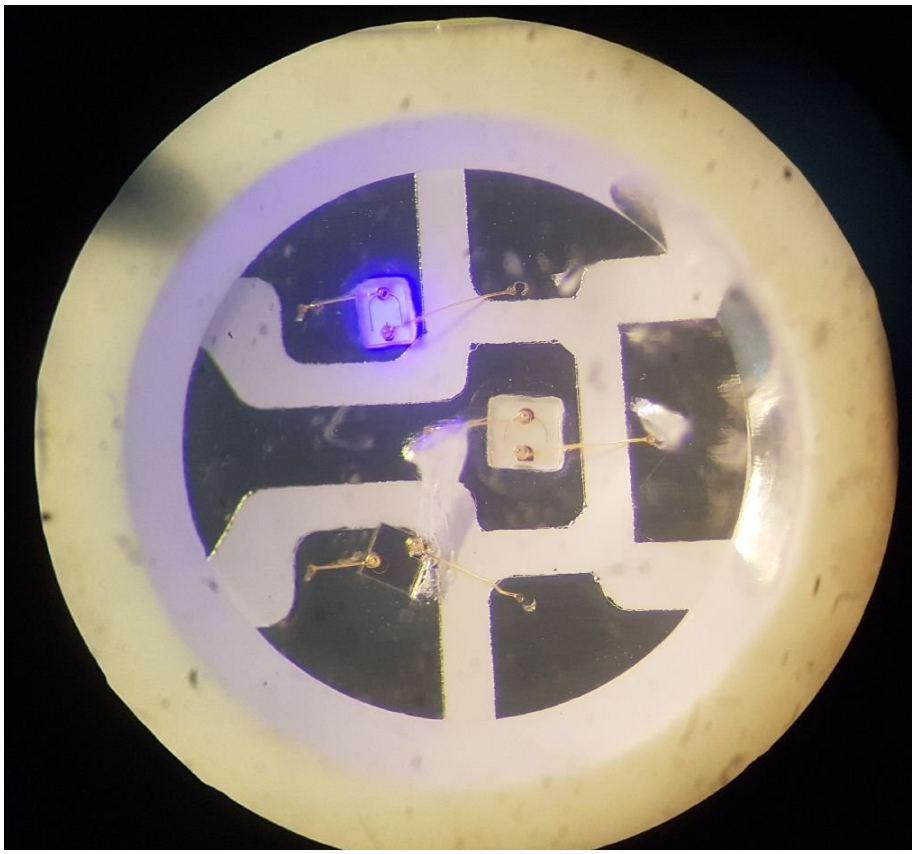

*Fig. 106: Microscopic image of the custom LED*

*The white part is the 2.8mm by 3.5mm plastic substrate, most of which is not visible. In it there is a well with sloping walls. It is 0.8mm deep and 2mm in diameter at the bottom. On the bottom there are metal surfaces (the dark, silvery areas) which are connected to the external metal pads that are then soldered onto the circuit board. The three cubes are the light emitting crystals (the actual LEDs), the blue one at the top is slightly glowing. The electrical connections between the LED chips and the metallisation are gold wires (gold is chemically inert and highly pliable), also clearly visible. The well is then filled with a silicone rubber material, which is transparent to UV. The silicone protects the semiconductor chips, which are sensitive to both atmospheric oxygen and water vapour.*

I told them what we needed, they told me what was possible from a practical and financial point of view. After a few iterations with measuring the emission

spectra of samples and changing the emitting crystals, we ended up with an LED that is adequate for the Antarium and also fits in the budget.

They built a custom tri-colour LED, with UV, blue and green emitters in a single 3.5 by 2.8 mm package. The physical arrangement of the LEDs in the package is shown in *Fig. 106* while their emission spectra are shown in *Fig. 107*.

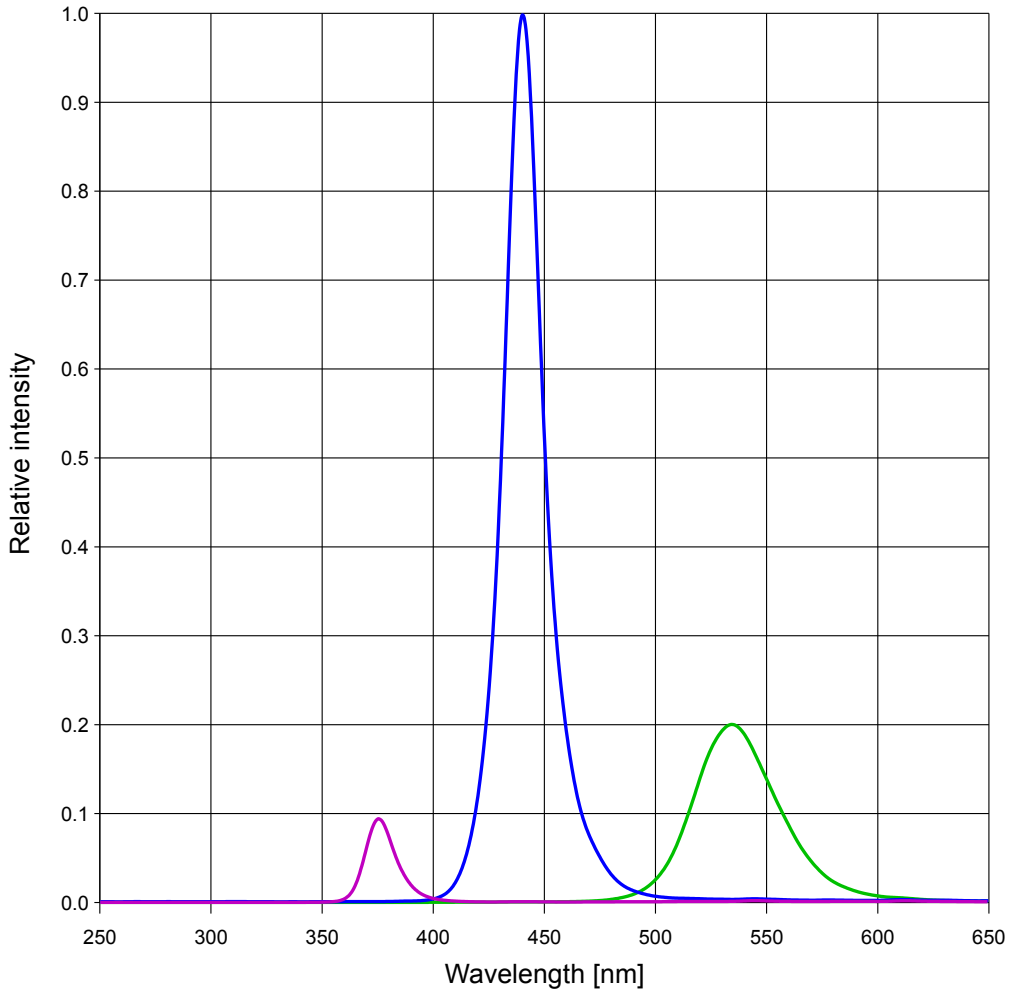


*Fig. 107: Emission spectra of the Antarium LEDs at their nominal 20mA current*

Obviously, the blue LED is much more efficient than either the green or the UV one. But making a LED less bright is easy, one just needs to decrease the drive current. Nominally, the maximum drive current for all three LEDs is 20mA. In the Antarium, however, the green, blue and UV LEDs will be driven by 20mA,

3.5mA and 13mA, respectively (how the figures were derived will be explained later). The resulting emission spectra are shown in *Fig. 108*.

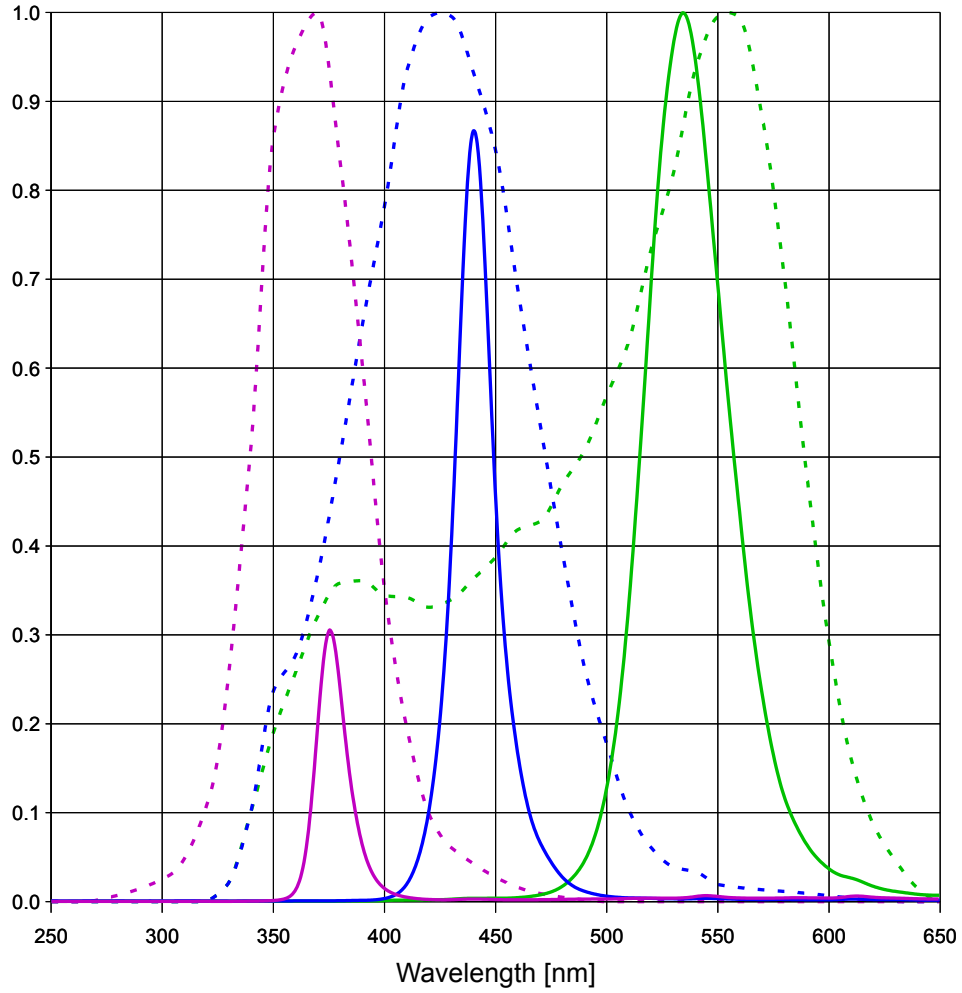


*Fig. 108: Emission spectra of the LEDs with their drive currents adjusted*

*The solid lines are the LED emission spectra. The dashed lines are the normalised photoreceptor sensitivity curves from Fig. 103 on page 271, repeated here for reference. Although due to technological limitations the LEDs do not match perfectly the sensitivity peaks, they are quite close to optimal wavelength.*

The excitation level of the photoreceptors for any polychromatic light can be calculated by evaluating the integral[24]:

$$E_X = \int S_X(\lambda) \cdot I(\lambda) \, d\lambda \quad ; \quad X \in \{U, B, G\}$$

24 *The derivation of that integral formula can be found in any radiation physics or photometry book.*

where $E_X$ is the excitation level of receptor $X$, $S_X$ is its sensitivity curve and $I$ is the spectrum of the polychromatic light.

If we perform that for the three LED spectra and map the excitations on the chromaticity triangle, then we get the area of the gamut that the LEDs can re-create, shown in *Fig. 109*.

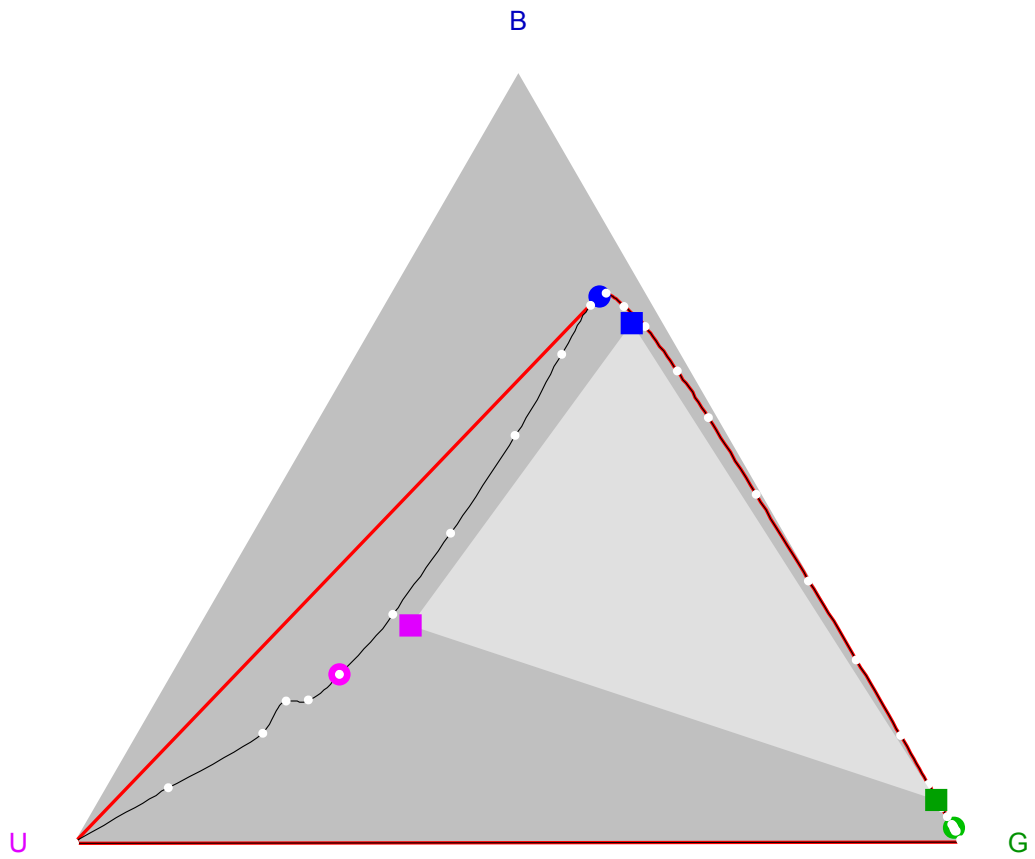


*Fig. 109: Colour gamut of the selected LEDs*

*The black curve traces the chromaticity of a monochromatic stimulus from 250nm to 600nm. The white dots are 10nm apart. The red curve encloses the colour gamut of the animal. The purple, blue and green dots show the peak sensitivities of the UV, blue and green photoreceptors. The light grey triangle is the gamut that the LEDs can represent. The purple, blue and green squares at the vertices of the triangle are the chromaticity positions of the LEDs.*

At first sight, the colour gamut covered by the LEDs is very small. However, as it will be shown, they are adequate to represent what the Antarium is expected to project.

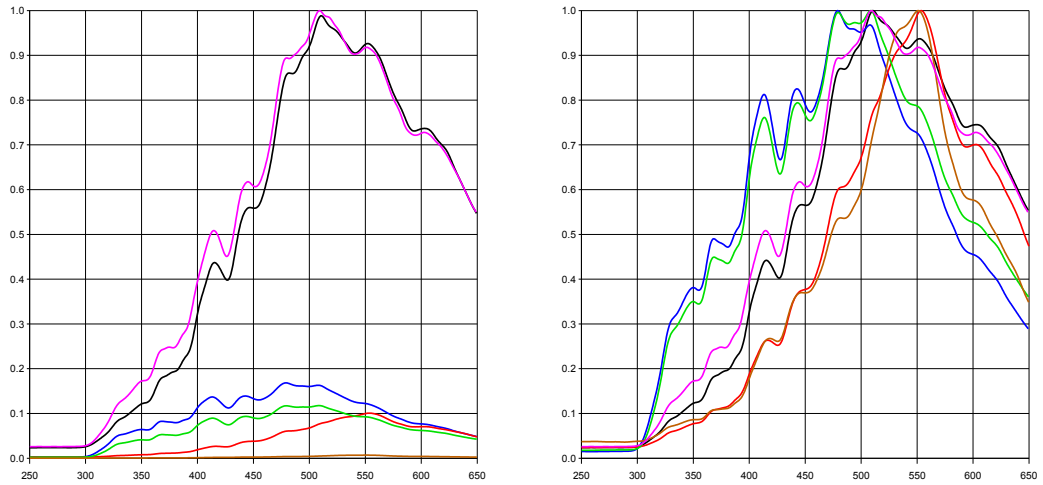


*Fig. 110: Spectra of selected natural objects*

*The left panel shows the spectra as measured, the right panel shows each spectrum normalised to its maximum. Below 315nm the ozone layer absorbs most of the Sun's radiation, that's why there the spectra are flat and only the dark current of the spectrometer is present.*

*Purple: a white reflectance standard (LabSphere) at direct sunlight; black: the Sun on a cloudless day, mid-afternoon; blue: clear morning sky approx 90° away from the Sun; green: the sky in a completely overcast afternoon; red: a bush lit by the morning sun; brown: the same bush at close to sunset.*

The device is supposed to render the natural environment of the animals, with artificial modifications that support the actual experiment. Therefore, we need to examine the spectrum of the light environment in which the animals normally operate. *Fig. 110* shows the raw and normalised spectra of certain natural objects that the animal might encounter, as measured with a USB-4000 spectrometer (Ocean Optics).

On the left panel, the black (direct sunlight) and purple (a white reference) curves are almost identical, the slight difference is being that humans see the Sun a little bit yellow, therefore “white” actually needs a slight elevation of the UV-blue part of the spectrum. (Interestingly, *Myrmecia* see our white a little bit green.) It is obvious that most of the energy is in the 470nm – 570nm (roughly green) region. In the 330nm – 400nm range, where the ants’ UV receptor is most sensitive, the intensity is about 20% that of the green region.

To get a better picture of what the Antarium needs to produce, we can now look at the normalised spectra, the right panel of *Fig. 110*. The peak is contained in

the 450-550nm range, that is, blue to green. In particular, vegetation (red and brown curves) is shifted towards the green from the sunlight or white reference, while the sky away from the sun (blue curve) is shifted to the blue, even in overcast conditions (green).

It has been pointed out by Möller (2002) that insects could exploit the UV-green contrast for navigation. Examining the UV range shows why UV is good for navigational purposes. The intensity of the 370nm component is about 10% of the peak for direct sunlight. For light reflected from vegetation that drops to about 7%, while for the sky away from the sun or under cloud cover the figure is roughly 28%. Thus, if an ommatidium looks at vegetation, it will see a UV to blue-green ratio of 0.07 while an ommatidium looking at the sky (but not directly to the sun) will detect a ratio of 0.28. Therefore the skyline formed by vegetation is marked by a sudden jump of relative UV content, by a factor of 4. At 450nm[25] the intensities are around 38% (vegetation) and 80% (sky), marking the skyline with a blue content with a factor of only 2. For the peek sensitivity of the green receptor, at 550nm, the figures are 100% (vegetation) and 75% (sky), a factor of a mere 1.33. From an ant's point of view, the UV delivers by far the largest discrimination between vegetation and sky, which was also the conclusion in (Barta and Horváth, 2004) and (Differt and Möller, 2015).

On page 278 I set up arbitrary LED current ratios. Those ratios came from the spectra demonstrated in *Fig. 110*. *Table 15* shows the relative excitation levels of the G, B and UV photoreceptors for the spectra measured at the scenes discussed above. It also shows the excitation levels for the LEDs when they are at full brightness (i.e. all three are driven by 20mA) and when their drive current is reduced to the 20mA for G, 3.5mA for B and 13mA for UV. As it can be seen from the table, with the chosen currents the excitation of the photoreceptors almost perfectly matches that of the white reference.

25 *At 425nm, where the photoreceptor sensitivity peak is, there is a dip in the spectrum of natural light, so the smoother 450nm point was used for the comparison.*

| Scene | G | B | UV |
|---|---|---|---|
| White reference | 1.000 | 0.466 | 0.112 |
| Direct sunlight | 1.000 | 0.508 | 0.142 |
| Morning sky | 1.000 | 0.727 | 0.280 |
| Overcast sky | 1.000 | 0.677 | 0.251 |
| Well-lit bush | 1.000 | 0.354 | 0.078 |
| Sunset bush | 1.000 | 0.357 | 0.083 |
| LEDs @ 20mA each | 0.783 | 1.000 | 0.111 |
| LEDs @ 20/3.5/13 mA | 1.000 | 0.478 | 0.113 |

*Table 15: Relative receptor excitation levels at various scenes*

This now leads us back to the issue of the seemingly limited colour gamut covered by the selected LEDs. If we also plot the colour of the natural scenes on the chromaticity triangle, we will see that it is not as limited as it seems.

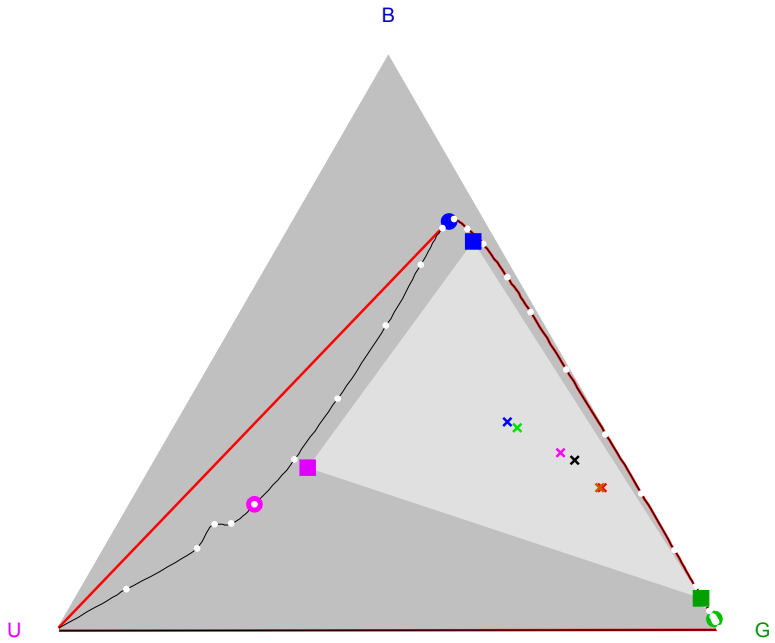


*Fig. 111: Colour gamut with natural scenes marked*

*The black curve traces the chromaticity of a monochromatic stimulus from 250nm to 600nm. The white dots are 10nm apart. The red curve encloses the colour gamut of the animal. The purple, blue and green dots show the peak sensitivities of the UV, blue and green photoreceptors. The light grey triangle is the gamut that the LEDs can represent. The purple, blue and green squares at the vertices of the triangle are the chromaticity positions of the LEDs.*

*The coloured crosses represent the chromaticities of objects found in natural scenes. All of them are within, and far from the sides of, the triangle spanned by the LED chromaticities. See the text for a more in-depth analysis.*

*The colour of the crosses represent: black: a white reflectance standard (LabSphere) at direct sunlight; purple: the Sun on a cloudless day, mid-afternoon; blue: morning sky approx 90 ° away from the Sun; green: the sky in a completely overcast afternoon; red: a bush lit by the morning sun; brown: the same bush at close to sunset.*

It is clear from both the spectral curves of *Fig. 110* and the chromaticity triangle in *Fig. 111*, that natural scenes have a relatively small chromatic spread (Chiao et al., 2000). It is understandable, as they are dominated by the spectrum of sunlight (filtered by the atmosphere). I have measured the reflected spectra of other natural objects (ground, tree trunks, etc), not shown on the figures, and they, too, show very similar characteristics. To move to the bottom-left area of the colour gamut would require a surface that is highly reflective of UV but absorbs most of the visible spectrum. Such surfaces are not occurring readily in nature (Barta and Horváth, 2004).

Another issue is the spectral representation of night scenes. No night measurements were taken (the spectrometer is not sensitive enough), but it is known that the spectral composition of moonlight is very similar to that of sunlight, except that shorter wavelengths are suppressed somewhat (Ciocca and Wang, 2013). Ants cannot see red, thus from their point of view moonlight has a somewhat elevated green component, but otherwise it is the same as sunlight. Since the reflective properties of objects remain the same regardless of what shines on them, the new Antarium will still be able to represent every night colour the animal can see.

In conclusion, the chosen LEDs and their adjusted drive currents are capable of rendering natural scenes and allow a wide range of colour manipulations, except moving into the very UV-heavy unnatural part of the gamut.

## The geometry of the new Antarium

The existing Antarium is composed of 55 identical isosceles triangles which form a biscribed pentakis dodecahedron with an opening at the bottom. That polyhedron was chosen because that is the largest number of facets achievable using identical polygons. Manufacturing cost increases with the number of different circuit boards and the very tight budged limited our choices. The budget of the Mk-II device is more flexible, allowing us to explore alternatives.

First, let us consider the trackball apparatus on which the experimental animal is placed. It is a self-contained unit, a cube with approximately 170mm edges. The ball protrudes ~5mm from the surface of the top face of the cube and the animal's head is not more than 10mm above the ball. That means that her downward visual field is limited by the trackball enclosure:

$$\alpha = \tan^{-1} \frac{-15\,mm}{170\,mm/2} = -10^{o}$$

The Antarium's elevation limit was -50° which is much more than what was necessary. We do not need to go down that much, simply because the animal's downward view is obscured by the trackball.

Even if we shrink the trackball apparatus to a 100mm cylinder (with a 100mm ball that is as small as we can go) and take into account that the animal's head

is offset by half of its body-length, i.e. by 12mm, the limit angle is still -22° (see *Fig. 112*).

Thus, if the elevation limit of the new projector is set to -22°, the device will still function as expected.

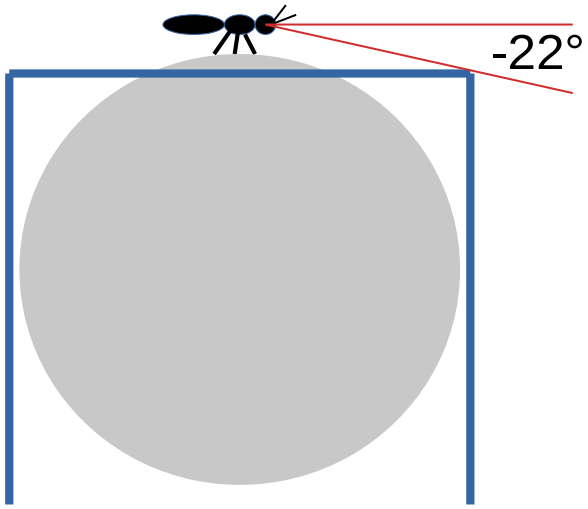


*Fig. 112: The downward elevation limit of the visual field of the animal on the trackball*

With that in mind, could we find another polyhedron that is a better approximation than the pentakis dodecahedron and still not too expensive? For the Mk-II device the chosen polyhedron is the pentakis-hexakis truncated icosahedron. The truncated icosahedron is the shape of the soccer ball, 20 identical, regular hexagons and 12 identical, regular pentagons approximate a sphere. We then replace the polygons with hexagonal or pentagonal pyramids, to get an even better approximation of the sphere. Depending on how tall we make the pyramids, we can have different shapes. We set the pyramids so that the resulting polyhedron has neither an inscribed nor a circumscribed sphere, that is, the surfaces of the triangles do not all touch the same sphere and the vertices do not all lie on the same sphere, but it has an edge-scribed sphere, that is, a sphere that all edges are a tangent of. That choice gives us certain geometric advantages that will become clear later. In short, the pyramids themselves are built from identical isosceles triangles and the heights of the pyramids are set to values that offer optimal properties for projection.

The triangles for the hexagonal and pentagonal pyramids are different, thus we need two kinds of triangles. The full polyhedron is built from 120 triangles of the hexagonal (H) kind and 60 triangles of the pentagonal (P) kind, as show in *Fig. 113*.

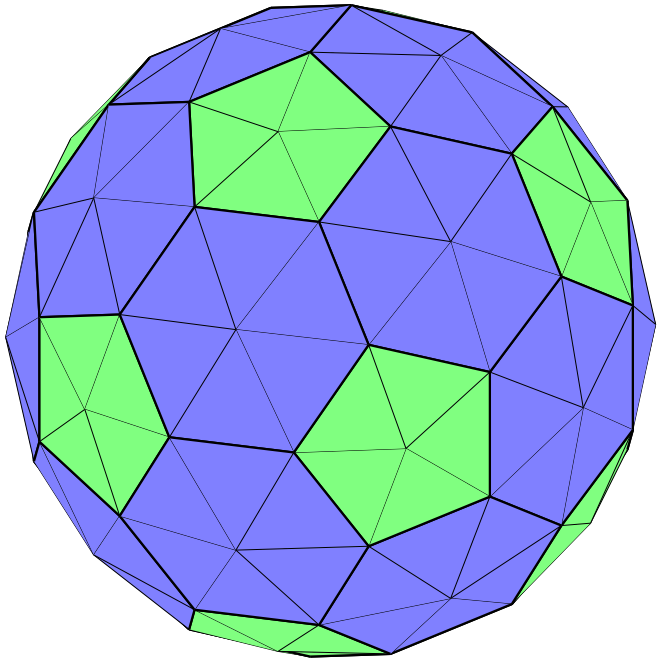

*Fig. 113: Edge-scribed hexakis-pentakis truncated icosahedron*

*The polyhedron contains 120 identical isosceles 'H' triangles (blue) that form the 20 hexagonal pyramids and 60 identical isosceles 'P' triangles (green) that form the 12 pentagonal pyramids. Note that the edges shared by P and H triangles are always their bases. Also note that there are three kinds of vertices: a) where the apices of 5 P triangles meet, b) where the apices of 6 H triangles meet, and c) where the bases of 4 H and 2 P triangles meet, with 2H+1P with their right hand side bottom vertices and 2H+1P with their left hand side bottom vertices.*

We rotate the solid so that a pentagon is at the top and another at the bottom. We then remove from the bottom 15 type P and 30 type H triangles to create an opening. Unfortunately, the circumference of the opening is not flat, that is, the vertices that form it are not all on the same plane. That somewhat complicates the mechanical construction of the device, but the gain from changing the polyhedron outweigh that cost. The vertices of the opening are at -26.57° and -30.99° elevation, exceeding the -22° requirement that was calculated previously.

Thus, the Mk-II device will contain 90 type H and 45 type P triangles, more than doubling the number of triangular facets compared to the existing Antarium. The final shape is shown in *Fig. 114*.

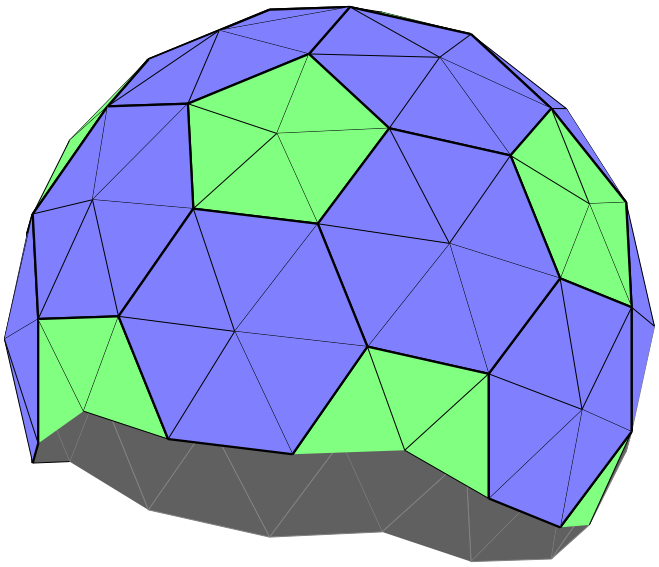

*Fig. 114: The shape of the Mk-II Antarium*

*15 P and 30 H triangles were removed from the bottom, leaving 45 P and 90 H triangles to form the hull of the device. The image is slightly rotated for better visibility of the opening.*

Let us turn our attention to the triangles. *Fig. 115* shows the P and the H variants, with an equilateral triangle in the middle for comparison.

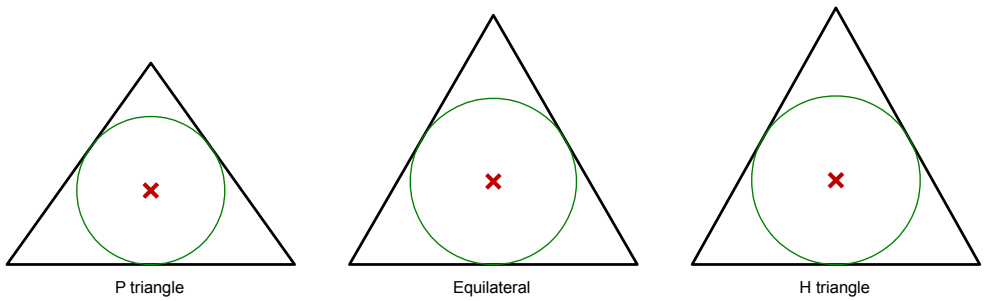


*Fig. 115: The geometry of the P and H triangles*

*The red cross is the point where the surface normal points to the centre of the sphere. The green circle is the largest inscribed circle in the triangle, its centre is the red cross. The equilateral triangle is shown for comparison.*

The cross in the middle of the triangle is where the radius of the edge-scribed sphere intersecting the triangle is parallel with the triangle's normal (i.e. where the triangle is at right angles with the radius). The circles in the triangles are centred around that point, they are the largest circle that can be drawn inside the triangle, the importance of which will be discussed later.

As it is obvious from the figure, the H triangle is almost equilateral, it is only 3.05% taller. The P triangle is more different, it is 19.2% flatter.

An important aspect of the geometry of the device is the optical parallax, that is, the angle at which the corners of the triangles are seen from the centre of the sphere, relative to the normal of the triangle going through the centre of the sphere. It is important because the optical axes of the LEDs are always parallel to the normal, therefore that angle is the parallax at which the corner LEDs are seen by the animal. For the old Antarium that angle was over 20°. For the new device from the centre of the sphere the base corners of the H triangles are seen at 13.52° off from the centre of the triangle, the apex is seen at 13.97°, the base corners of the P triangle are at 13.11° off centre while its apex is at 10.39°. That is a very significant improvement compared to the old Antarium.

If we select the baseline for the triangles to be 210mm, then the centre points (where the normal is parallel to the radius) of the H triangles are 506.0mm and that of the P triangles are 506.8mm from the spherical centre, so the Mk-II Antarium retains the ~1m diameter sphere approximation, which has been proven to be a convenient size from the point of view of mechanical construction, usability and heat dissipation.

With a 210mm base the height of the H triangle is 187.4mm and that of the P triangle is 147.0mm while the radii of the inscribed circles are 61.5mm for the H and 54.0mm for the P variants.

## LED arrangement

Ideally, we would like to have an LED arrangement where the projections of the LEDs form a uniform regular hexagonal lattice on the surface of the edge-scribed sphere. Euler proved over 300 years ago that that is not possible, the best achievable is to cover the sphere with hexagons plus 12 pentagons. Even

worse, in general the hexagons will be neither regular nor uniform (the truncated icosahedron is the exception, not the rule).

On top of the constraints posed by mathematics, we have further restrictions: we want to manufacture only one kind of H and one kind of P triangles, therefore the same kind of triangles must have the same LED arrangement. In addition, due to manufacturing technology limits, the LEDs cannot be too close to the edges of the triangles.

First, let us try to lay a regular hexagonal pattern on the triangles and we will deal with the spherical projection issues later. Let us start with the bases of the triangles, which will form the edges of the polyhedron where pyramids meet. The bottom row of the LEDs will there meet with a 180° rotated variant of itself. We need to place the LEDs in the bottom such that across the edge they form a hexagonal lattice with their own rotated copy, as shown in *Fig. 116*.

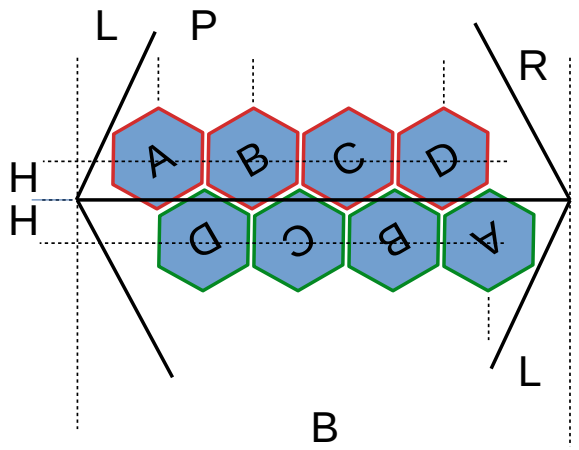


*Fig. 116: Hexagonal arrangement at the base of the triangles*

*For details, see text.*

In the figure B is the length of the base of the triangle, H is how far the bottom row of LEDs is from the edge, L and R are the horizontal distance between the LED and the left and right corner of the triangle, respectively, and P is the period of the hexagons, i.e. the distance between adjacent LEDs (which are at the centre of each hexagon).

There are some restriction on the figures, determined by geometry. Let N be the number of LEDs in the bottom row. Then the following equations must all be true:

$$B = L + R + P \cdot (N - 1)$$
$$R = L + \frac{1}{2} \cdot P$$
$$H = \frac{\sqrt{3}}{4} \cdot P \approx 0.433 \cdot P$$

Since B is given (210mm), R, L and H are dependent on P, so we need to choose P and the largest N which still confines the bottom row inside the triangle. The chosen value for P is 11.3mm. With that figure the H and P triangles will contain 171 and 133 LEDs, respectively. The pattern of LEDs on the triangles using an initially strict hexagonal arrangement is shown in *Fig. 117*.

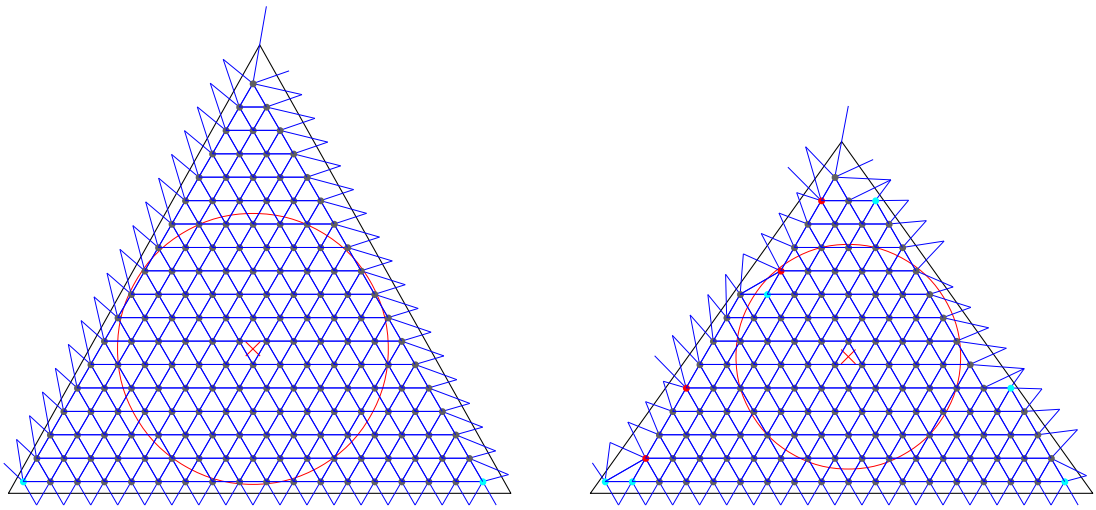

*Fig. 117: Original LED arrangement*

*The H-triangle (left) is tiled with 171 LEDs while the P-triangle (right) has 133. The LEDs are marked by the dots. The lines between the dots represent the adjacency relations of the LEDs. Leds marked with dark blue dots have 6 neighbours, cyan ones have only 5 while red ones have 7. Links that go outside the triangle's boundaries connect a LED to one on an adjacent triangle.*

*The red cross marks a spot which is at the centre of a triangle formed by 3 LEDs; the LED triangle is chosen by being the closest to the optical centre of the panel. The LEDs within the red circle are marked as fixed, see the text for further details.*

The LED positions are then optimised. What we want is the LEDs to be at identical angular distance from each other, not on the planar surface of the triangular panels, but their apparent position projected to a sphere. A software was written to perform the optimisation. For reasons that will be explained later, some LEDs were fixed and the optimiser could not move them; these LEDs are the ones that are enclosed by the red circle on the figure.

The LED arrangement after the optimisation is shown in *Fig. 118*.

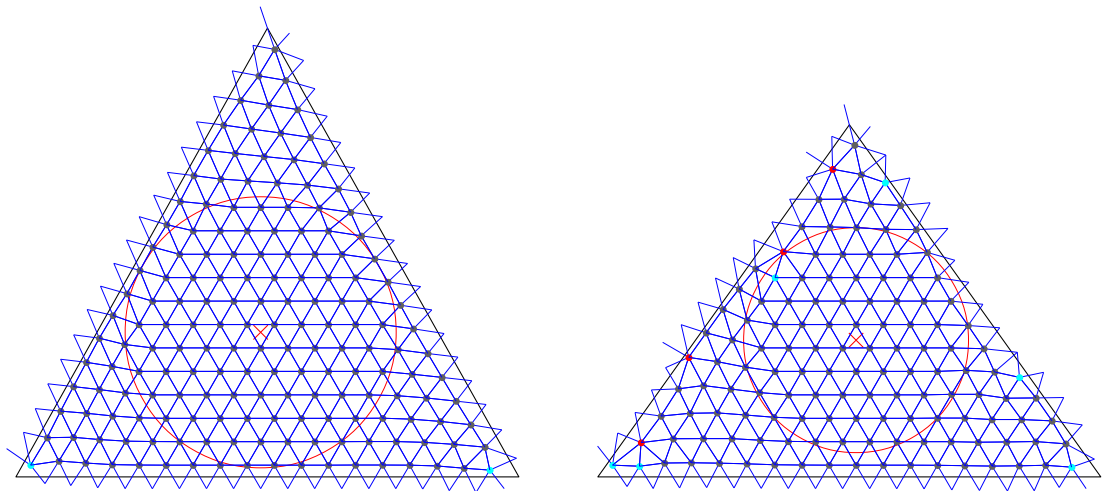

*Fig. 118: Opimised LED positions*

*The positions of the LEDs on the H (left) and P (right) triangles after the optimiser made the angular spread between adjacent LEDs projected to a sphere as small as possible. The LEDs within the red circles were not moved from their original position.*

To verify that the optimisation process worked, all LEDs on all panels were projected to a sphere. For each LED its nearest neighbour was found and the angular distance between the LED and the neighbour was calculated. Then a histogram was created where the horizontal axis is the angle (sorted into 0.0125° wide bins) and the vertical axis is the number of LEDs where the angular distance falls into the bin. That process was repeated for the second, third, …, sixth closest neighbour[26].

26 *For the LEDs at the very bottom, where the Antarium ends, only 4 neighbours were sought.*

If all the LEDs are on an ideal hexagonal grid, then all neighbours of all LEDs are at equal angular distance and each histogram would contain a single non-zero column, which is as high as many LEDs there are (21,375). We know that such an arrangement is not achievable even in theory, there will be some angular spread. However, if the optimiser worked, we still expect the angles to be contained within a narrow range. *Fig. 119* shows the histograms for the 6 closest neighbours.

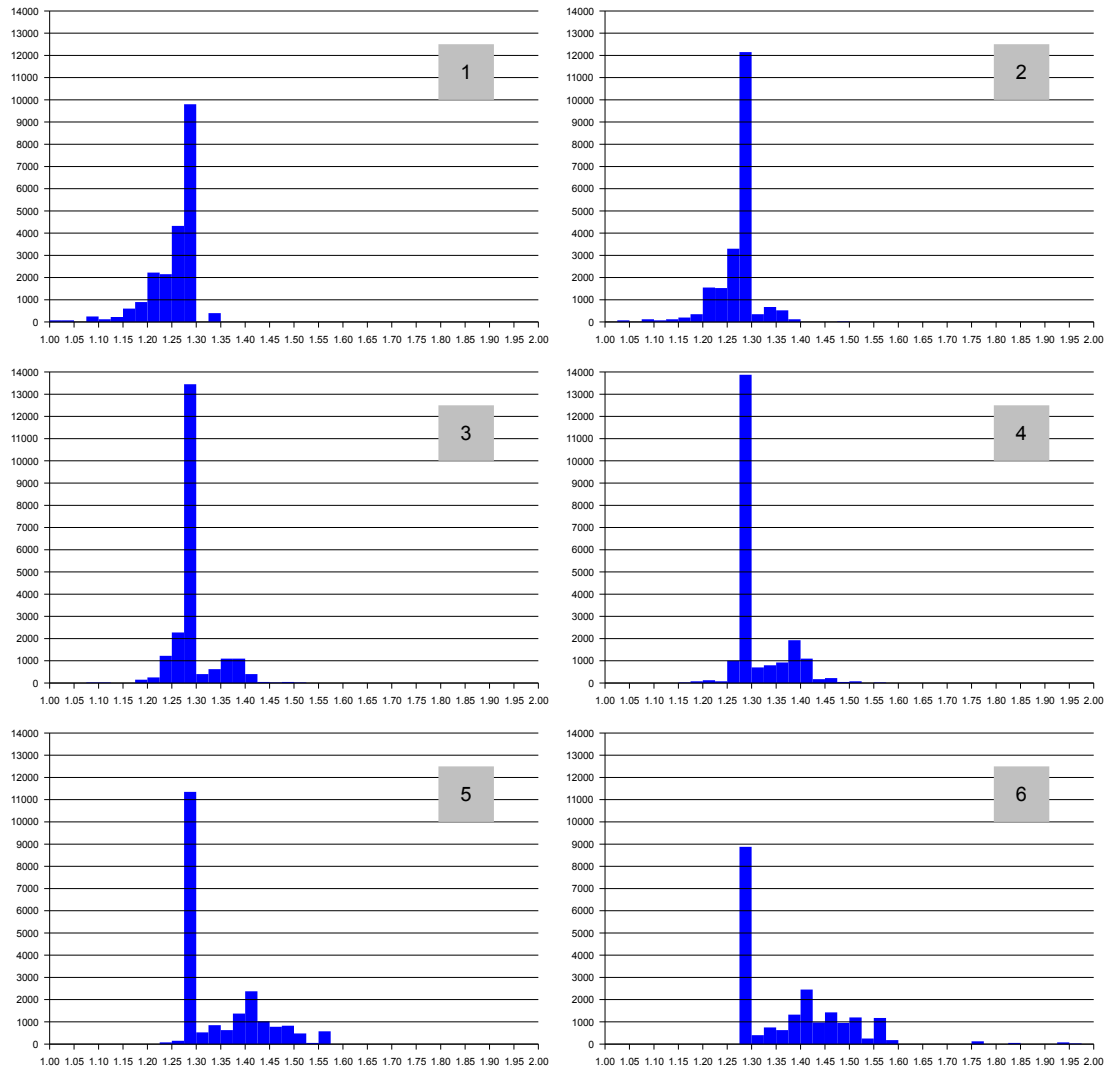


*Fig. 119: Nearest neighbours angular distribution*

*The number of LEDs of which the 1st, 2nd, … 6th closest neighbouring LED is at the given angle bin. The bins are 0.0125° wide and there are 21,375 LEDs all together.*

As the figure shows the optimiser did a reasonable job, considering the constraints it had to work under. The vast majority of angular distances fall

between 1.2° and 1.3° for the 2nd to 4th closest neighbours, while for the closest neighbour the distribution is skewed towards the lower angles (down to 1°), for the 5th and the 6th closest the histograms spread to higher angles (up to 1.6° and 2°, respectively). It is a significant improvement over the original Antarium, for which the angle was centred around 1.6° for the first 3 closest neighbours and above 2° for all others. The smaller neighbour angles result from an increased number of LEDs (21,375 vs. 19,855) being distributed over a smaller surface area (more of the sphere is missing at the bottom). The 6 close neighbours instead of 3 is due to the way the initial LED arrangement was designed.

*Fig. 120* shows the centre of the panoramic image that was used in the proof of concept experiments.

*Fig. 120: Reference image for Fig. 121.*

The new LED arrangement is demonstrated in *Fig 121*, where the scene is shown to be projected to the sphere.

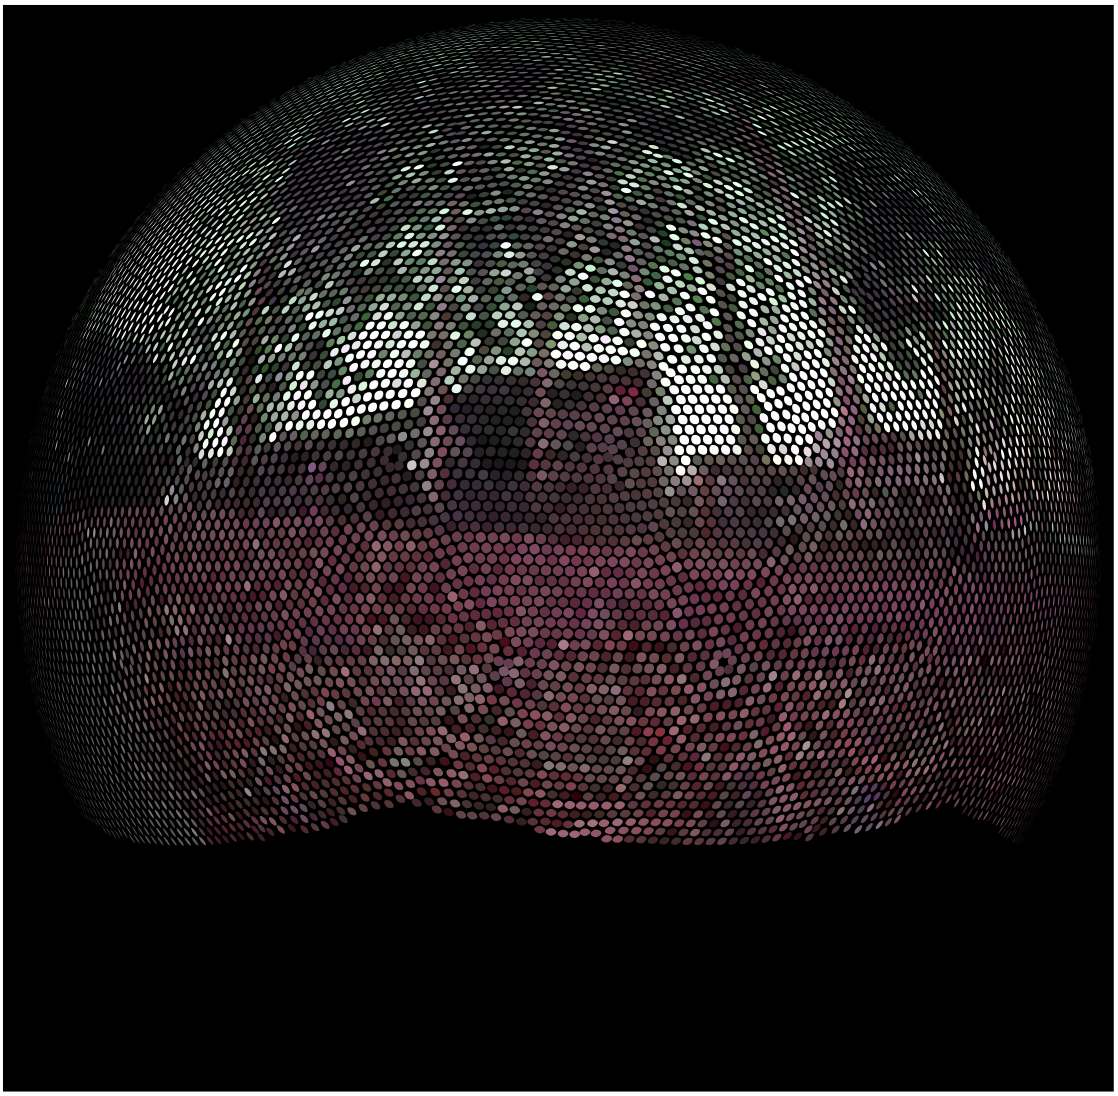

*Fig. 121: Natural scene projected to a sphere with the new LED arrangement*

*For comparison with the old Antarium, see Fig. 44 on page 180.*

To see whether this arrangement is indeed superior, the same mapping and image difference calculations were carried out as for the altered reality experiment, described on page 233.

In *Fig. 122* the outside image is shown to be mapped to an ommatidial arrangement with 3,500 ommatidia per eye (which is realistic for *M. pyriformis*). The figure also shows the same image first mapped to the Antarium's LEDs and those LEDs subsequently mapped to the ommatidia. Both the old and new Antarium are evaluated so that their visual fidelity can be compared.

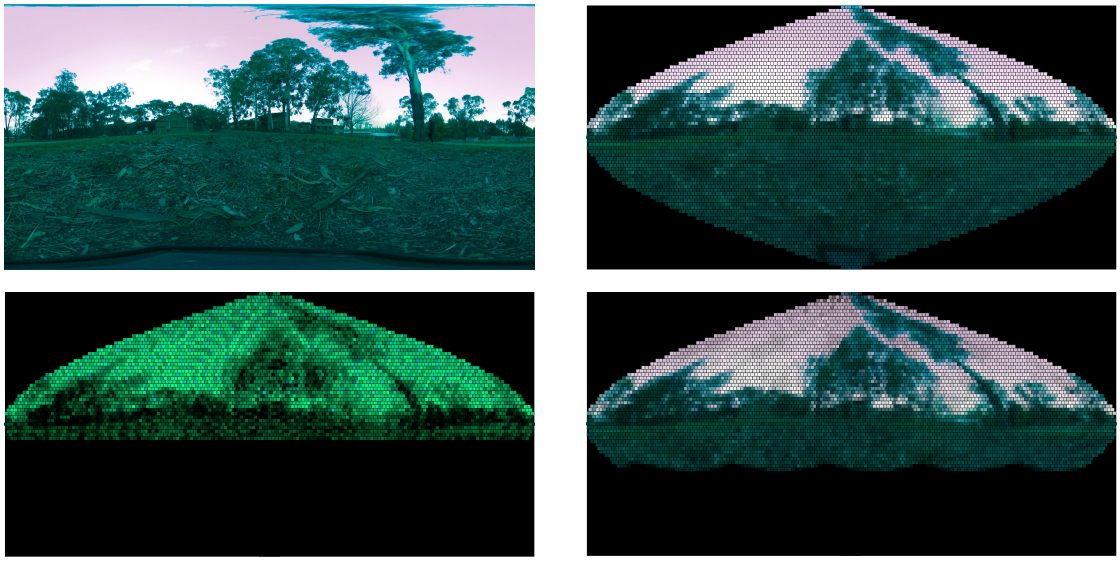

*Fig. 122: What the ant sees in the old and the new Antarium*

*Top left: A panoramic image. The UV is represented by the red channel, hence the pink hue of the sky. Top right: The image mapped to 7,000 ommatida, that is, the panorama as the animal would see it. Bottom left: The image mapped to the old Antarium's LEDs then the LEDs mapped to the 7,000 ommatidia, that is, what the animal sees when she is in the Antarium projecting the panoramic image. Note the green sky due to the spectral mismatch of the LEDs. Also note the moiré noise due to the sub-optimal LED arrangement. Bottom right: What the animal will see in the new Antarium. The resolution is better and there is no moiré noise. There is a slight intensity variation, most prominent on the sky, which will be electronically compensated after a calibration process when the device is built. Note that the spectral match of the new LEDs made the UV component visible for the animals inside the device, there is no colour distortion.*

In the altered reality experiment it was demonstrated that the rotational image difference function minima are very shallow in the old Antarium. *Fig. 123* compares the rotIDFs calculated outdoor, in the old Antarium and in the new one.

As it was discussed in the altered reality experiment, the Antarium's colour mismatch and LED arrangement issues render the rotIDF all but useless for navigational purposes. On the other hand, the Mk-II device provides a robust rotIDF. Not clearly visible on the figure, but the global minima (which are 0 for the outdoor setting, all colours) are 0.05 for UV, 0.06 for blue and 0.06 for green while for the existing Antarium they are 0.40, 0.26 and 0.12, respectively. More important than the absolute figures are the ratio of the minimum and the maximum.

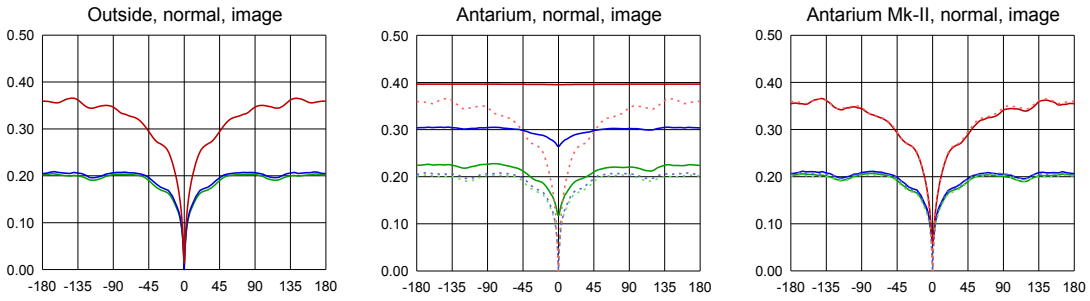


*Fig. 123: rotIDF comparison*

*The red, blue and green lines are the rotIDF calculated for the UV, blue and green content of the image. The dashed curves on the centre and right panels are the curves of the left panel, plotted as a reference.*

If we define the metric of discrimination ability as

$$D = \frac{max - min}{max}$$

then 1.0 indicates excellent discrimination and 0 means a completely flat rotIDF.

| | Outside | Mk-I | Mk-II |
|---|---|---|---|
| **UV** | 1.00 | 0.00 | 0.86 |
| **Blue** | 1.00 | 0.13 | 0.71 |
| **Green** | 1.00 | 0.48 | 0.71 |

*Table 16: rotIDF discrimination comparison*

As shown in *Table 16*, the old device offered no discrimination in UV, which is known to be the most important navigation aid for the animal. The blue channel offers minimal discrimination and the green channel is the best, although not very good. The new device offers convincing values for all three channels.

## Custom optics

The custom LEDs in the new Antarium are very wide-angle, due to technological limitations. That means that only a small fraction of the light output is directed towards the animal at the centre. A custom designed lens array can be used to focus the LED output on the one hand and to direct the beams towards the centre on the other.

A problem is that most visually transparent plastics are not transparent for UV. Fortunately, poly-(methyl methacrylate) is a widely used thermoplastic which is

naturally UV resistant and at the same time UV transparent. It is sold under the name of Perspex and Plexiglass among others. Pure PMMA is transparent above about 300nm and its refraction index, which is around 1.49, remains fairly constant in the 300 to 700nm range.

So, we can use PMMA to build a lens in front of each LED, with the required optical properties. *Fig. 124* shows a lens that collimates a single LED, without bending the beam.

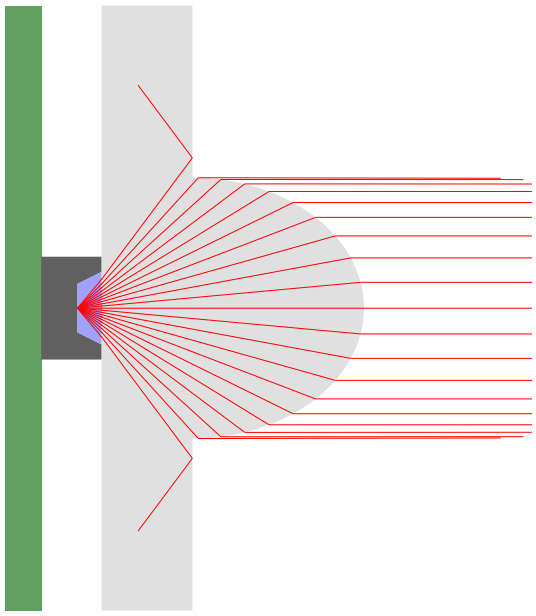

*Fig. 124: Lens example*

*The green bar of the left represents the printed circuit board. On it is shown the LED, in which the actual light emitting crystal is at the bottom of small well, filled with silicone rubber with a refraction index of $n_s$=1.6. In front of the LED, without any air gap, is a 3mm thick Perspex ($n_p$=1.49) plate on which the lens is formed. In this example the lens was calculated to collimate the light. The red lines show the ray-traced paths of the light emitted by the LED. Also shown is that rays outside the lens area suffer a total internal reflection and will be bouncing back and forth between the planar surfaces until exhausted.*

Unlike the example shown, the new Antarium lenses need not fully collimate the light as it cannot be guaranteed that the experimental animal's head is exactly

at the centre of the device. Rather, a roughly 6° divergence is sought so that the entire area where the animal can possibly be is illuminated.

Furthermore, the lenses also need to bend the light towards the centre. The lens shape must be calculated for each LED separately, as each LED has a slightly different angle by which the beam needs to be bent.

The mathematical details of the calculations are omitted for brevity, but the basic idea is that we can calculate the beams as they need to go towards the centre of the Antarium and we can also calculate the beams as they leave the LED and travel through the Perspex. We then pair each ray leaving the LED with a ray going to the centre, calculate their intercept and at that point we can calculate the required surface normal of the perspex to achieve the necessary refraction. If we do that for very many pairs of rays, we can then stitch these elementary small surfaces together to form the geometrical shape of the lens.

As we are not designing precision optics, only “classical” refraction is taken into account which simplifies the calculations. Still, it is not an easy task as the silicone in the LED has a refraction index different from that of the perspex, therefore the rays in the perspex do not appear to be coming from a single point source. Nevertheless, calculating the lens is doable on computer by an iterative algorithm.

Once each lens is calculated, a metal master can be machined and polished to a mirror-smooth surface. It can then be sent to a plastic manufacturer for replication using Perspex.

## Changes to the electronics

Instead of having a central distributor and an FPGA on each panel, a small daughter board (a board that plugs into connectors on the “motherboard”, in this case the triangular panels with the LEDs) was designed. It contains the power supply for the LEDs, a microcontroller to drive the triangle and two Ethernet connectors, both capable of 100Mb/s. The daughter board can drive either kind of triangle. The two Ethernet ports allow the boards to be daisy-chained.

The separate daughter board is a cost increase, but it allows the triangular LED panels to be simpler. That will offset (but not completely negate) the increase. In addition, as each triangle now has a 100Mb/s connection, the frame rate can be increased. If the rendering computer has a 1Gb/s Ethernet interface (commonplace today) then employing an off-the-shelf intelligent Ethernet switch to distribute the packets to the panels will allow us to send the frames faster than before.

Although the theoretical network speed is now 10 times that of the Mk-I device, the speed increase is less, as the amount of data is significantly higher (each pixel now receives three 16-bit intensity values instead of the Mk-I device's 16-bit intensity plus 8-bit hue; there are more pixels; and an increased protocol overhead) but 500 frames per second is achievable. That also means that in a closed loop scenario the total lag between the ant's moving on the trackball and the consequent scene change can be kept below 5ms if the rendering computer is fast enough.

## Polarisation

This is one issue of the new Antarium that has not been solved yet and it is not because of the lack of trying.

In the original Antarium the polarisation of the UV light was provided by polarising film discs mounted on rotating actuators. For each UV pixel 2 UV LEDs were employed, one affected by the polariser and another that wasn't allowing to polarisation contrast to be set by the intensity ratio of the two LEDs. There were significantly fewer UV pixels than blue-green ones, so the high cost of the actuators could be born by the project.

In Antarium Mk-II all pixels contain a UV LED but only one. The above method could not be used. There is an additional problem: in the old Antarium the UV LEDs emitted 405nm while in the new one they are at 370nm. That makes a huge difference, as it turns out.

In Antarium Mk-II the polariser in front of each pixel must be wide bandwidth: it must be transparent from ~350nm to ~600nm. That in itself is a major problem. Cheap visual light polariser sheets are usually transparent only above 400nm,

very sharply falling off below that so they would block the UV. Commercially available UV polarisers are usually not made from plastic, are very expensive and not transparent in the visual spectrum, not to mention that most of them are only available in small sizes. This is a long standing issue. All sorts of tricks have been tried to make polariser sheets that worked (to some degree) in both the UV and the visible range for at least 60 years, attested, for example, by (Makas, 1962).

Nevertheless, let us assume that we can get broadband polarisation materials and see how could we provide both variable angle and variable contrast.

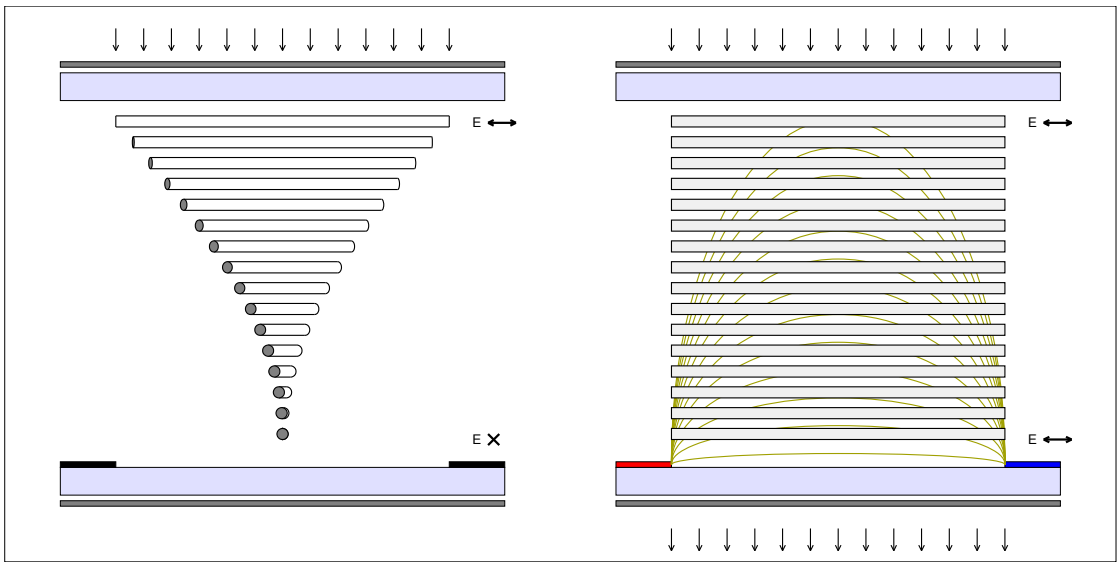


*Fig. 125: In-plane switching liquid crystal display*

*The panel, from top to bottom, is composed of a polariser (dark grey), a glass panel (light blue), the liquid crystal (rods), another glass panel (light blue) with electrodes (black) and another polariser (dark gray), aligned the same way as the top one. Further explanation is in the text.*

Liquid crystals contain long, rod-like molecules that are polar (have an electric charge difference between their two ends). They have the remarkable property that in a thin layer these elongated molecules organise themselves to a helix across the layer. If polarised light passes through the helix, its polarisation angle changes. If the thickness of the liquid crystal layer is such that the helix twists by a quarter turn, the polarisation angle of the light passing through will be rotated by 90°. Liquid crystal displays are based on this phenomenon, shown in *Fig. 125*.

The light enters from the top. As it passes through the top polariser its E-vector assumes the direction that is parallel with the horizontal axis of the page, indicated by the E vector direction just below the top glass on both panels.

On the left panel, the electrodes are not energised so the liquid crystal molecules form a quarter-turn helix and rotate the E-vector by 90° so that it is now perpendicular to the page, as indicated just above the bottom glass. Because that direction is orthogonal to the bottom polariser's alignment, no light is passed through.

On the right panel, the electrodes are energised and the resulting electrical field (indicated by the greenish arches) rotate the polar liquid crystal molecules so that they are all aligned parallel. That results in the light's E-vector not being rotated while passing through the liquid crystal layer, therefore when it exits through the bottom glass, its polarisation is aligned with the bottom polariser and can pass through it. Applying smaller voltages on the electrodes result in a partial alignment of the molecules, and a rotation of the E-vector by more than 0° but less than 90°. Consequently light will emerge at the bottom, but with diminished intensity. By varying the voltage, we can control the amount of light the panel lets through.

If we now remove the bottom polariser, the light will always emerge unimpeded, but its E-vector direction will be determined by the electrode potentials. If we put such an arrangement in front of the lens array of the Antarium, we can control the intensity of the light (by means of controlling the LEDs) and the polarisation direction (by controlling the liquid crystal panel).

We still need to control the polarisation contrast, though. The pixels of a liquid crystal panels are small, thus we can easily fit about 100 pixels within each light beam as it leaves the lens array. If we set all 100 pixels into the same direction, the light will be 100% polarised. If we set all 100 pixels into random directions, the beam will contain all sorts of polarisation angles, thus at the end, appear completely un-polarised. By varying the alignment of the pixels, we can achieve a varying degree of polarisation contrast, with good resolution.

This scheme requires that there exist glass and liquid crystal materials which work in the 350nm-600nm range. Whether that is true or not is being investigated. I am aware of liquid crystal panels operating at 405nm, but I do not yet know for certain whether going down as low as 350nm is achievable or not.

There is also fallback solution for the case where a suitable film is available but no liquid crystal display can be manufactured. It is very far from perfect, but offers at least a rudimental sky polarisation.

The polarisation of the sky varies slowly. If we polarise the light on a triangular panel by triangular panel basis instead of a LED by LED basis, we would segment the sky (taken as everything above the horizon) to 90 regions with individual polarisation characteristics. There is no experimental data available on the spatial resolution of polarisation sensitivity of the ants, so it is not yet certain that that is sufficient. However, we do know that both the dorsal rim area and the ocelli are out of focus, so we can assume that they look at the polarisation of large areas, not small points. If that is true, we have a good chance that a per-triangle segmentation is sufficient.

In that case, we can put a single, large polariser disc in front of the triangle. The centre of the disc is the centre of the triangle, so the disc is as large as it can be. For both H and P triangles the disc covers close to 60% of the LEDs. That, of course, means that we cannot produce 100% polarised light. However, the sky itself is never 100% polarised. Its polarisation contrast peaks at about 60% for clear sky and atmosphere, and around 55% for hazy conditions (Gambling and Billard, 1967). That is in line with what the circular disc can produce.

Still, we need to solve the problem of varying not only the E-vector direction but also the polarisation contrast. It was mentioned earlier that when the software optimised the LED positions on the panels, it was instructed to leave certain LEDs alone, at fixed positions. The LEDs that were at fixed positions happen to be on concentric circles around the middle of the polariser disc. We can then laser-cut holes into the polariser disc at the LED positions. When the disc is aligned so that the holes are above the LEDs, there is no polarisation. As we rotate the disc, more and more of the beam emanating from the lens in front of the LED is covered by the polariser, so the beam becomes more and more

polarised. If we then cut holes at regular intervals between the LEDs we can select both the angle and the contrast with limited, but probably sufficient resolution.

At the time of writing this thesis, I am trying to find and contact companies that specialise on wide bandwidth polarisers and also on speciality liquid crystal panels. It is not yet clear which solution will be chosen and who will supply the necessary components.

# Other considerations

## Manufacturing cost

The original Antarium was built on a budget of approximately AU $50,000, at the time corresponding to US $35,000, as shown in *Table 17*.

| Cost item | AUD |
|---|---|
| Initial prototype panel | 3,000 |
| Second prototype panel | 1,400 |
| 60 final panels @ AUD 533 ea | 32,000 |
| Power supplies | 1,600 |
| Mechanical construction | 9,000 |
| Sundry items (shipping, import duty, cables, etc) | 3,000 |
| **Total** | **50,000** |

*Table 17: Cost breakdown of the Antarium*

Antarium Mk-II is a more sophisticated device, where the custom LEDs as well as the custom optics increase the cost significantly. Also, it is a work in progress, therefore only cost estimation can be provided. *Table 18* shows the predicted budget of the new device, with the cost of polarisers being unknown.

| Item | AUD |
|---|---|
| 23,000 custom LEDs @ AU $2 ea | 46,000 |
| Prototype panels (2 iterations) | 6,000 |
| 95H+50P LED panels (not including the LEDs) | 15,000 |
| 140 control panels | 11,000 |
| Custom optics prototypes (2 iterations) | 10,000 |
| Custom optics 95H+50P | 22,000 |
| Power supplies | 3,000 |
| Mechanical construction | 15,000 |
| Polarisers | 0 |
| Sundry items | 5,000 |
| **Total** | **133,000** |

*Table 18: Cost estimate for the Antarium Mk-II device*
*Note that the cost of the polarisation solution is not yet known.*

## Other species

The Antarium Mk-II device was designed with *Myrmecia* in mind; the LEDs were selected to match the sensitivity of the eyes of those ants.

However, the device could be useful for other insects with similar spectral sensitivities as well. A recent review paper by van der Kooi *et al* (2021) lists the spectral sensitivities of over 200 insects. Several of the species listed have tri-colour vision with spectral sensitivity peaks near to the LED emission peaks of the device. That includes other Hymenoptera: several bee, wasp and ant species have a matching sensitivity – unfortunately, *Apis mellifera* has a mismatch with their UV peak being at 346nm. Lepidoptera and Odonata tend to employ more than 3 different opsins, so it is unlikely that the Antarium would be suitable for them. The bi-colour cockroaches are also out, because while their UV sensitivity matches, the ~500nm peak of their other receptor is about midway between the Antarium's blue and green LEDs. *Locusta migratoria* is a reasonably good match. The Antarium Mk-II device has a very high frame rate, which would be suitable for fast flying and rapidly turning insects like flies. The horse fly *Tabanus bromius* is a good match, with similar spatial resolution as *Myrmecia* (Meglič et al., 2019). *D. melanogaster* is a complete mismatch, though.

So the device can be used for some other species as well. In addition, an Antarium can be built for a specific experimental animal in mind. The electronics, mechanical and optical design can be left unchanged and only the custom LEDs need to be made with matching wavelength. But then, that Antarium would not be suitable for the ants we plan to experiment with.

In conclusion, the Antarium is not a universal device. Its spectral composition and spatial resolution cannot be adjusted at will. Such an Omnitarium would require a completely different (and likely very expensive) design. The Antarium was designed with specific species of ants in mind and whether it is suitable for any particular other species is a question of luck.

## Electrophysiology in the Antarium

The Antarium was envisioned for behavioural experiments. That, however, does not exclude the possibility of performing electrophysiology experiments in it. There are two issues that should be considered.

Electrophysiology experiments usually require the animal to be fixed relative to the measurement jig, as the microelectrodes must be kept in place. For flying insects the animal's intended path is determined from measuring torque along the spatial axes, during which the animal does not actually moves, but the scenery is adjusted according to the measured torque. A fixed walking insect on a trackball seems to be similar: she rotates the ball along its two horizontal axes as she walks and rotates the ball around its vertical axis as she turns. The problem, however, is that when she walks, the arm length is the diameter of the ball but when she turns, the arm length is only the distance between her legs. A ball designed to match the momentum of the animal for walking will feel extremely sluggish for her when she wants to turn. This can possibly distort the experimental result.

The other issue is the electrical noise generated by the Antarium. The machine switches the current on and off for over 60,000 LED chips at a very high frequency, which then will result in emitting a strong, wideband electromagnetic noise. Electrophysiology equipments operate below-MHz frequencies, thus a large portion of the radiated noise would not affect them. The lower frequency noise would still pose a problem though. Thus, to minimise the Antarium's interference with the measurement equipment, a Faraday cage must be employed.

A frame made of relatively thick copper wire along the inner edges of the polyhedron and triangular fine copper mesh panels soldered to it would form a self-sustaining structure acting as a Faraday cage, that could shield the inner space from radiated noise with a relatively small loss of light.

# Outlook

In this thesis I presented a “virtual reality” machine for ants that can be used to perform visual navigation experiments by presenting the animal its natural visual surroundings. The machine also allows the scenery to be arbitrarily manipulated thereby presenting the animal with challenges to which her response can be measured and analysed.

The Antarium was built and tested and it was proven to be a viable device to perform controlled behaviour experiments. However, it has some serious shortcoming, which are addressed by a new device currently being built.

The Antarium Mk-II device will allow us to seek answers to questions which would be hard to interrogate without a reconstructed visual reality. Some of these questions are listed below:

- What do the optic lobes do? The optic lobes heavily process the raw image but it is not know what information they extract. In the Antarium we can modify the image, by increasing or decreasing the edge contrast, change colour contrast, re-colour the image or apply non-linear or fractal transformations and see where the animals’ navigation ability breaks down. That will certainly help us to understand the role of the optic lobes in the processing of visual information and possibly identify the local and global features extracted from the image.
- How do animals compensate head attitude around the roll and pitch axes? Ardin *et al* (2015) found that rotIDF discrimination ability diminishes very rapidly as the pitch changes even by small amounts, yet they found no experimental evidence of ants attempting to stabilise their head against body pitch changes. Raderschall *et al* (2016) reported that the ability of night active *M. pyriformis* to keep their head horizontal decreases with decreasing light levels. They also showed that the minimum of the rotIDF between a horizontal reference image and its roll-rotated self becomes undetectable at more than 25° roll misalignment. In the Antarium we can change the orientation of the image at will and can therefore perform experiments to systematically map the animals’

navigation ability against roll and pitch angles, providing the experimental data needed to devise theoretical models that explain the observed behaviour.

- How ants monitor the changing view as they walk? To this date I have not systematically performed closed loop experiments where the animals actually do experience the visual consequences of their translational movements. We can then displace them in that world without actually physically disturbing them and observe them truly navigating and correcting their paths depending on scene changes, something that does not happen with the static scenes I have been using so far. Closing the loop will also allow us to investigate in detail when and how often the ants' behaviour is guided by a comparison between memorised and currently experienced views.
- How do ants respond to impossible visual experiences? We will be able to test reactions to unusual or impossible situations, such as scene changes or the movement of objects that are not the result of the animals' own movements. We can rotate celestial cues, such as the sun and the polarised skylight relative to the landmark panorama and test current models of visual navigation that are using metrics of view familiarities, such as the IDF, by smoothly changing scenes via homomorphic fractal transformations that change the spatial, but not other statistical scene properties.

Clearly, as we start experimenting with the new device, more questions and experimental opportunities will arise, with the eventual goal being of recording from the ants' brain while they are navigating inside the Antarium.

At this stage, the Antarium is specifically designed to address the visual system properties of Australian *Myrmecia* ants. However, provided with sufficient funding, it is possible to design a projection device along the same principles that has higher resolution and above all, light sources that can be tuned to match the spectral sensitivities of different animals. For this the LED based approach would need to be abandoned and replaced with a projector using lasers as light sources and an elaborate mechano-optical arrangement directing

the laser beams to a projection surface surrounding the animal. The unique aspect of the Antarium, however, namely the ability to present ants with their familiar visual navigational environment, will be an interesting challenge to replicate for other animals.

# Bibliography

Aepli, F., Labhart, T., and Meyer, E.P. (1985). Structural specializations of the cornea and retina at the dorsal rim of the compound eye in Hymenopteran insects. Cell and Tissue Research *239*, 19–24. .

Aksoy, V. (2014). Experience Based Use of Landmark and Vector Based Orientation During Homing by the Ant *Formica cunicularia* (Hymenoptera: Formicidae). Journal of Insect Behavior *27*, 357–369. https://doi.org/10.1007/s10905-013-9432-2.

Aksoy, V., and Camlitepe, Y. (2018). Spectral sensitivities of ants – a review. Animal Biology *68*, 55–73. https://doi.org/10.1163/15707563-17000119.

Al Toufailia, H., Couvillon, M.J., Ratnieks, F.L.W., and Grüter, C. (2013). Honey bee waggle dance communication: signal meaning and signal noise affect dance follower behaviour. Behavioral Ecology and Sociobiology *67*, 549–556. https://doi.org/10.1007/s00265-012-1474-5.

Ardin, P., Mangan, M., Wystrach, A., and Webb, B. (2015). How variation in head pitch could affect image matching algorithms for ant navigation. Journal of Comparative Physiology A *201*, 585–597. https://doi.org/10.1007/s00359-015-1005-8.

Ardin, P., Peng, F., Mangan, M., Lagogiannis, K., and Webb, B. (2016). Using an Insect Mushroom Body Circuit to Encode Route Memory in Complex Natural Environments. PLOS Computational Biology *12*, e1004683. https://doi.org/10.1371/journal.pcbi.1004683.

Arena, P., Patane, L., and Strauss, R. (2013). The Insect Mushroom Bodies: a Paradigm of Neural Reuse. In Advances in Artificial Life, ECAL 2013, (MIT Press), pp. 765–772.

Armett-Kibel, C., Meinertzhagen, I.A., and Dowling, J.E. (1977). Cellular and Synaptic Organization in the Lamina of the Dragon-Fly *Sympetrum rubicundulum*. Proceedings of the Royal Society B: Biological Sciences *196*, 385–413. https://doi.org/10.1098/rspb.1977.0047.

Arshavsky, V.Y. (2010). Vision: The Retinoid Cycle in *Drosophila*. Current Biology *20*, R96–R98. https://doi.org/10.1016/j.cub.2009.12.039.

Aso, Y., Sitaraman, D., Ichinose, T., Kaun, K.R., Vogt, K., Belliart-Guérin, G., Plaçais, P.-Y., Robie, A.A., Yamagata, N., Schnaitmann, C., et al. (2014a). Mushroom body output neurons encode valence and guide memory-based action selection in *Drosophila*. ELife *3:e04580*. https://doi.org/10.7554/eLife.04580.

Aso, Y., Hattori, D., Yu, Y., Johnston, R.M., Iyer, N.A., Ngo, T.-T., Dionne, H., Abbott, L., Axel, R., Tanimoto, H., et al. (2014b). The neuronal architecture of the mushroom body provides a logic for associative learning. ELife *3:e04577*. https://doi.org/10.7554/eLife.04577.

Assisi, C., Stopfer, M., and Bazhenov, M. (2020). Optimality of sparse olfactory representations is not affected by network plasticity. PLoS Comput Biol *16*, e1007461. https://doi.org/10.1371/journal.pcbi.1007461.

Baddeley, B., Graham, P., Philippides, A., and Husbands, P. (2011). Holistic visual encoding of ant-like routes: Navigation without waypoints. Adaptive Behavior *19*, 3–15. https://doi.org/10.1177/1059712310395410.

Baddeley, B., Graham, P., Husbands, P., and Philippides, A. (2012). A Model of Ant Route Navigation Driven by Scene Familiarity. PLoS Computational Biology *8*, e1002336. https://doi.org/10.1371/journal.pcbi.1002336.

Barta, A., and Horváth, G. (2004). Why is it advantageous for animals to detect celestial polarization in the ultraviolet? Skylight polarization under clouds and canopies is strongest in the UV. Journal of Theoretical Biology *226*, 429–437. https://doi.org/10.1016/j.jtbi.2003.09.017.

Barta, A., Farkas, A., Száz, D., Egri, Á., Barta, P., Kovács, J., Csák, B., Jankovics, I., Szabó, G., and Horváth, G. (2014). Polarization transition between sunlit and moonlit skies with possible implications for animal orientation and Viking navigation: anomalous celestial twilight polarization at partial moon. Applied Optics *53*, 5193–5204. https://doi.org/10.1364/AO.53.005193.

Batschelet, E. (1970). Recent Statistical Methods for Orientation Data. In Proceedings of Conference on Animal Orientation and Navigation Capabilities, (Wallop Station, VA, USA: NASA), p. 31.

Beckers, R., Goss, S., Deneubourg, J.-L., and Pasteels, J.-M. (1989). Colony size, communication and ant foraging strategy. Psyche: A Journal of Entomology *96*, 239–256. .

Boeddeker, N., Lindemann, J.P., Egelhaaf, M., and Zeil, J. (2005). Responses of blowfly motion-sensitive neurons to reconstructed optic flow along outdoor flight paths. Journal of Comparative Physiology A *191*, 1143–1155. https://doi.org/10.1007/s00359-005-0038-9.

Bolek, S., and Wolf, H. (2015). Food searches and guiding structures in North African desert ants, *Cataglyphis*. Journal of Comparative Physiology A *201*, 631–644. https://doi.org/10.1007/s00359-015-0985-8.

Borst, A. (2009). *Drosophila*'s View on Insect Vision. Current Biology *19*, R36–R47. https://doi.org/10.1016/j.cub.2008.11.001.

Borst, A., and Egelhaaf, M. (1989). Principles of visual motion detection. Trends in Neurosciences *12*, 297–306. https://doi.org/10.1016/0166-2236(89)90010-6.

Borst, A., and Haag, J. (2002). Neural networks in the cockpit of the fly. Journal of Comparative Physiology A: Sensory, Neural, and Behavioral Physiology *188*, 419–437. https://doi.org/10.1007/s00359-002-0316-8.

Buehlmann, C., Cheng, K., and Wehner, R. (2011). Vector-based and landmark-guided navigation in desert ants inhabiting landmark-free and landmark-rich environments. Journal of Experimental Biology *214*, 2845–2853. https://doi.org/10.1242/jeb.054601.

Buehlmann, C., Hansson, B.S., and Knaden, M. (2012a). Path Integration Controls Nest-Plume Following in Desert Ants. Current Biology *22*, 645–649. https://doi.org/10.1016/j.cub.2012.02.029.

Buehlmann, C., Hansson, B.S., and Knaden, M. (2012b). Desert Ants Learn Vibration and Magnetic Landmarks. PLoS ONE *7*, e33117. https://doi.org/10.1371/journal.pone.0033117.

Buehlmann, C., Fernandes, A.S.D., and Graham, P. (2018). The interaction of path integration and terrestrial visual cues in navigating desert ants: what can we learn from path characteristics? The Journal of Experimental Biology *221*, jeb167304. https://doi.org/10.1242/jeb.167304.

Buehlmann, C., Wozniak, B., Goulard, R., Webb, B., Graham, P., and Niven, J.E. (2020). Mushroom Bodies Are Required for Learned Visual Navigation, but Not for Innate Visual Behavior, in Ants. Current Biology *30*, 3438-3443.e2. https://doi.org/10.1016/j.cub.2020.07.013.

Caron, S.J. (2013). Brains Don't Play Dice—or Do They? Science *342*, 574–574. .

Caron, S.J.C., Ruta, V., Abbott, L.F., and Axel, R. (2013). Random convergence of olfactory inputs in the *Drosophila* mushroom body. Nature *497*, 113–117. https://doi.org/10.1038/nature12063.

Cartwright, B.A., and Collett, T.S. (1982). How honey bees use landmarks to guide their return to a food source. Nature *295*, 560–564. https://doi.org/10.1038/295560a0.

Cartwright, B.A., and Collett, T.S. (1983). Landmark learning in bees: Experiments and models. Journal of Comparative Physiology ? A *151*, 521–543. https://doi.org/10.1007/BF00605469.

Cartwright, B.A., and Collett, T.S. (1987). Landmark maps for honeybees. Biological Cybernetics *57*, 85–93. https://doi.org/10.1007/BF00318718.

Cheeseman, J.F., Millar, C.D., Greggers, U., Lehmann, K., Pawley, M.D.M., Gallistel, C.R., Warman, G.R., and Menzel, R. (2014a). Way-finding in displaced clock-shifted bees proves bees use a cognitive map. Proceedings of the National Academy of Sciences *111*, 8949–8954. https://doi.org/10.1073/pnas.1408039111.

Cheeseman, J.F., Millar, C.D., Greggers, U., Lehmann, K., Pawley, M.D.M., Gallistel, C.R., Warman, G.R., and Menzel, R. (2014b). Reply to Cheung et al.: The cognitive map hypothesis remains the best interpretation of the data in

honeybee navigation. Proceedings of the National Academy of Sciences *111*, E4398–E4398. https://doi.org/10.1073/pnas.1415738111.

Chen, P.-J., Belušič, G., and Arikawa, K. (2020). Chromatic information processing in the first optic ganglion of the butterfly *Papilio xuthus*. Journal of Comparative Physiology A *206*, 199–216. https://doi.org/10.1007/s00359-019-01390-w.

Cheung, A. (2014). Animal path integration: A model of positional uncertainty along tortuous paths. Journal of Theoretical Biology *341*, 17–33. https://doi.org/10.1016/j.jtbi.2013.09.031.

Cheung, A., Zhang, S., Stricker, C., and Srinivasan, M.V. (2007). Animal navigation: the difficulty of moving in a straight line. Biological Cybernetics *97*, 47–61. https://doi.org/10.1007/s00422-007-0158-0.

Cheung, A., Hiby, L., and Narendra, A. (2012). Ant Navigation: Fractional Use of the Home Vector. PLoS ONE *7*, e50451. https://doi.org/10.1371/journal.pone.0050451.

Cheung, A., Collett, M., Collett, T., Dewar, A., Dyer, F., Graham, P., Mangan, M., Narendra, A., Philippides, A., Stürzl, W., et al. (2014). Still no convincing evidence for cognitive map use by honeybees. Proceedings of the National Academy of Sciences *111*, 4396–4397. .

Chiao, C.-C., Cronin, T.W., and Osorio, D. (2000). Color signals in natural scenes: characteristics of reflectance spectra and effects of natural illuminants. Journal of the Optical Society of America A *17*, 218. https://doi.org/10.1364/JOSAA.17.000218.

Ciocca, M., and Wang, J. (2013). By the light of the silvery Moon: fact and fiction. Physics Education *48*, 360–367. https://doi.org/10.1088/0031-9120/48/3/360.

Collett, M. (2012). How Navigational Guidance Systems Are Combined in a Desert Ant. Current Biology *22*, 927–932. https://doi.org/10.1016/j.cub.2012.03.049.

Collett, M. (2014). A desert ant's memory of recent visual experience and the control of route guidance. Proceedings of the Royal Society B: Biological Sciences *281*, 20140634. https://doi.org/10.1098/rspb.2014.0634.

Collett, T.S. (2019). Path integration: how details of the honeybee waggle dance and the foraging strategies of desert ants might help in understanding its mechanisms. The Journal of Experimental Biology *222*, jeb205187. https://doi.org/10.1242/jeb.205187.

Collett, M., and Collett, T.S. (2018). How does the insect central complex use mushroom body output for steering? Current Biology *28*, R733–R734. https://doi.org/10.1016/j.cub.2018.05.060.

Collett, T.S., and Zeil, J. (2018). Insect learning flights and walks. Current Biology *28*, R984–R988. https://doi.org/10.1016/j.cub.2018.04.050.

Collett, M., Harland, D., and Collett, T.S. (2002). The use of landmarks and panoramic context in the performance of local vectors by navigating honeybees. Journal of Experimental Biology *205*, 807–814. .

Collett, M., Chittka, L., and Collett, T.S. (2013). Spatial Memory in Insect Navigation. Current Biology *23*, R789–R800. https://doi.org/10.1016/j.cub.2013.07.020.

Collett, T.S., Collett, M., and Wehner, R. (2001). Guidance of desert ants by extended landmarks. Journal of Experimental Biology *204*, 1635–1639. .

Collett, T.S., Graham, P., and Durier, V. (2003). Route learning by insects. Current Opinion in Neurobiology *13*, 718–725. https://doi.org/10.1016/j.conb.2003.10.004.

Collett, T.S., Graham, P., and Harris, R.A. (2007). Novel landmark-guided routes in ants. Journal of Experimental Biology *210*, 2025–2032. https://doi.org/10.1242/jeb.000315.

Cronin, T.W., Warrant, E.J., and Greiner, B. (2006). Celestial polarization patterns during twilight. Applied Optics *45*, 5582–5589. https://doi.org/10.1364/AO.45.005582.

Crozier, R.H., Dobric, N., Imai, H.T., Graur, D., Cornuet, J.-M., and Taylor, R.W. (1995). Mitochondrial-DNA Sequence Evidence on the Phylogeny of Australian Jack-Jumper Ants of the *Myrmecia pilosula* Complex. Molecular Phylogenetics and Evolution *4*, 20–30. https://doi.org/10.1006/mpev.1995.1003.

Cruse, H., and Wehner, R. (2011). No need for a cognitive map: decentralized memory for insect navigation. PLoS Computational Biology *7*, e1002009. .

Dacke, M., Byrne, M.J., Baird, E., Scholtz, C.H., and Warrant, E.J. (2011). How dim is dim? Precision of the celestial compass in moonlight and sunlight. Philosophical Transactions of the Royal Society B: Biological Sciences *366*, 697–702. https://doi.org/10.1098/rstb.2010.0191.

Dacke, M., Baird, E., Byrne, M., Scholtz, C.H., and Warrant, E.J. (2013a). Dung Beetles Use the Milky Way for Orientation. Current Biology *23*, 298–300. https://doi.org/10.1016/j.cub.2012.12.034.

Dacke, M., Byrne, M., Smolka, J., Warrant, E., and Baird, E. (2013b). Dung beetles ignore landmarks for straight-line orientation. Journal of Comparative Physiology A *199*, 17–23. https://doi.org/10.1007/s00359-012-0764-8.

Dahmen, H., Wahl, V.L., Pfeffer, S.E., Mallot, H.A., and Wittlinger, M. (2017). Naturalistic path integration of *Cataglyphis* desert ants on an air-cushioned lightweight spherical treadmill. The Journal of Experimental Biology *220*, 634–644. https://doi.org/10.1242/jeb.148213.

Dan, O., Hopp, E., Borst, A., and Segev, I. (2018). Non-uniform weighting of local motion inputs underlies dendritic computation in the fly visual system. Scientific Reports *8:5787*. https://doi.org/10.1038/s41598-018-23998-9.

Darwin, C. (1873). Origin of Certain Instincts. Nature 417–418. https://doi.org/10.1038/007417a0.

DeVoe, R.D., and Ockleford, E.M. (1976). Intracellular responses from cells of the medulla of the fly, *Calliphora erythrocephala*. Biological Cybernetics *23*, 13–24. https://doi.org/10.1007/BF00344147.

Dewar, A.D.M., Wystrach, A., Graham, P., and Philippides, A. (2015). Navigation-specific neural coding in the visual system of Drosophila. Biosystems *136*, 120–127. https://doi.org/10.1016/j.biosystems.2015.07.008.

Dickinson, M.H., and Lighton, J.R.B. (1995). Muscle efficiency and elastic storage in the flight motor of *Drosophila*. Science *268*, 87–90. https://doi.org/10.1126/science.7701346.

Differt, D., and Möller, R. (2015). Insect models of illumination-invariant skyline extraction from UV and green channels. Journal of Theoretical Biology *380*, 444–462. https://doi.org/10.1016/j.jtbi.2015.06.020.

Differt, D., and Stürzl, W. (2020). A generalized multi-snapshot model for 3D homing and route following. Adaptive Behavior 105971232091121. https://doi.org/10.1177/1059712320911217.

Dombeck, D.A., and Reiser, M.B. (2012). Real neuroscience in virtual worlds. Current Opinion in Neurobiology *22*, 3–10. https://doi.org/10.1016/j.conb.2011.10.015.

Dornhaus, A., and Chittka, L. (2004). Why do honey bees dance? Behavioral Ecology and Sociobiology *55*, 395–401. https://doi.org/10.1007/s00265-003-0726-9.

Dreyer, D., Frost, B., Mouritsen, H., Günther, A., Green, K., Whitehouse, M., Johnsen, S., Heinze, S., and Warrant, E. (2018). The Earth's Magnetic Field and Visual Landmarks Steer Migratory Flight Behavior in the Nocturnal Australian Bogong Moth. Current Biology *28*, 2160-2166.e5. https://doi.org/10.1016/j.cub.2018.05.030.

Dreyer, D., Frost, B., Xu, J., Mouritsen, H., Green, K., Whitehouse, M., Chahl, J.S., and Warrant, E. (2019). The starry sky provides compass information during long-distance navigation in the Australian Bogong moth. In Fourth Internationa Conference on Invertebrate Vision, (Bäckaskog Castle, Sweden), p.

Dujardin, F. (1850). Mémoire sur le système nerveux des insectes. Ann. Sci. Nat. Zool *14*, 195–206. .

Dupeyroux, J., Serres, J., and Viollet, S. (2018). A Hexapod Walking Robot Mimicking Navigation Strategies of Desert Ants *Cataglyphis*. In Biomimetic and Biohybrid Systems, V. Vouloutsi, J. Halloy, A. Mura, M. Mangan, N. Lepora, T.J. Prescott, and P.F.M.J. Verschure, eds. (Cham: Springer International Publishing), pp. 145–156.

Durier, V., Graham, P., and Collett, T.S. (2003). Snapshot Memories and Landmark Guidance in Wood Ants. Current Biology *13*, 1614–1618. https://doi.org/10.1016/j.cub.2003.08.024.

Ehmer, B., and Gronenberg, W. (2002). Segregation of visual input to the mushroom bodies in the honeybee (*Apis mellifera*). The Journal of Comparative Neurology *451*, 362–373. https://doi.org/10.1002/cne.10355.

Ehmer, B., and Gronenberg, W. (2004). Mushroom body volumes and visual interneurons in ants: Comparison between sexes and castes. The Journal of Comparative Neurology *469*, 198–213. https://doi.org/10.1002/cne.11014.

Eichler, K., Li, F., Litwin-Kumar, A., Park, Y., Andrade, I., Schneider-Mizell, C.M., Saumweber, T., Huser, A., Eschbach, C., Gerber, B., et al. (2017). The complete connectome of a learning and memory centre in an insect brain. Nature *548*, 175–182. https://doi.org/10.1038/nature23455.

Esch, H.E., and Burns, J.E. (1995). Honeybees use optic flow to measure the distance of a food source. Naturwissenschaften *82*, 38–40. .

Fain, G.L., Hardie, R., and Laughlin, S.B. (2010). Phototransduction and the Evolution of Photoreceptors. Current Biology *20*, R114–R124. https://doi.org/10.1016/j.cub.2009.12.006.

Falibene, A., Roces, F., and Rössler, W. (2015). Long-term avoidance memory formation is associated with a transient increase in mushroom body synaptic complexes in leaf-cutting ants. Frontiers in Behavioral Neuroscience *9*, 13. https://doi.org/10.3389/fnbeh.2015.00084.

Farris, S.M. (2011). Are mushroom bodies cerebellum-like structures? Arthropod Structure & Development *40*, 368–379. https://doi.org/10.1016/j.asd.2011.02.004.

Felsenberg, J., Jacob, P.F., Walker, T., Barnstedt, O., Edmondson-Stait, A.J., Pleijzier, M.W., Otto, N., Schlegel, P., Sharifi, N., Perisse, E., et al. (2018). Integration of Parallel Opposing Memories Underlies Memory Extinction. Cell *175*, 709-722.e15. https://doi.org/10.1016/j.cell.2018.08.021.

Fent, K., and Wehner, R. (1985). Ocelli: A Celestial Compass in the Desert Ant *Cataglyphis*. Science *228*, 192–194. .

Fiore, V.G., Kottler, B., Gu, X., and Hirth, F. (2017). *In silico* Interrogation of Insect Central Complex Suggests Computational Roles for the Ellipsoid Body in Spatial Navigation. Frontiers in Behavioral Neuroscience *11:142*. https://doi.org/10.3389/fnbeh.2017.00142.

Fitak, R.R., and Johnsen, S. (2017). Bringing the analysis of animal orientation data full circle: model-based approaches with maximum likelihood. The Journal of Experimental Biology *220*, 3878–3882. https://doi.org/10.1242/jeb.167056.

Fleischmann, P.N., Christian, M., Müller, V.L., Rössler, W., and Wehner, R. (2016). Ontogeny of learning walks and the acquisition of landmark information in desert ants, *Cataglyphis fortis*. The Journal of Experimental Biology *219*, 3137–3145. https://doi.org/10.1242/jeb.140459.

Fleischmann, P.N., Grob, R., Wehner, R., and Rössler, W. (2017). Species-specific differences in the fine structure of learning walk elements in *Cataglyphis* ants. The Journal of Experimental Biology *220*, 2426–2435. https://doi.org/10.1242/jeb.158147.

Fleischmann, P.N., Grob, R., Müller, V.L., Wehner, R., and Rössler, W. (2018a). The Geomagnetic Field Is a Compass Cue in *Cataglyphis* Ant Navigation. Current Biology *28*, 1440-1444.e2. https://doi.org/10.1016/j.cub.2018.03.043.

Fleischmann, P.N., Rössler, W., and Wehner, R. (2018b). Early foraging life: spatial and temporal aspects of landmark learning in the ant *Cataglyphis noda*. Journal of Comparative Physiology A *204*, 579–592. https://doi.org/10.1007/s00359-018-1260-6.

Fleischmann, P.N., Grob, R., and Rössler, W. (2020). Magnetoreception in Hymenoptera: importance for navigation. Animal Cognition *23*, 1051–1061. https://doi.org/10.1007/s10071-020-01431-x.

Foster, J.J., and Dacke, M. (2019). An orientation strategy that is robust to light pollution. In The Fourth International Conference on Invertebrate Vision, (Bäckaskog Castle, Sweden), p. 172.

Foster, J.J., el Jundi, B., Smolka, J., Khaldy, L., Nilsson, D.-E., Byrne, M.J., and Dacke, M. (2017). Stellar performance: mechanisms underlying Milky Way orientation in dung beetles. Philosophical Transactions of the Royal Society B: Biological Sciences *372*, 20160079. https://doi.org/10.1098/rstb.2016.0079.

Foster, J.J., Smolka, J., Nilsson, D.-E., and Dacke, M. (2018). How animals follow the stars. Proceedings of the Royal Society B: Biological Sciences *285*, 20172322. https://doi.org/10.1098/rspb.2017.2322.

Foster, J.J., Tocco, C., Smolka, J., Khaldy, L., Baird, E., Byrne, M.J., Nilsson, D.-E., and Dacke, M. (2021). Light pollution forces a change in dung beetle orientation behavior. Current Biology *31*, 3935-3942.e3. https://doi.org/10.1016/j.cub.2021.06.038.

Freas, C.A., and Cheng, K. (2018). Landmark learning, cue conflict, and outbound view sequence in navigating desert ants. Journal of Experimental Psychology: Animal Learning and Cognition *44*, 409–421. https://doi.org/10.1037/xan0000178.

Freas, C.A., and Cheng, K. (2019). Panorama similarity and navigational knowledge in the nocturnal bull ant *Myrmecia midas*. The Journal of Experimental Biology *222*, jeb193201. https://doi.org/10.1242/jeb.193201.

Freas, C.A., and Spetch, M.L. (2019). Terrestrial cue learning and retention during the outbound and inbound foraging trip in the desert ant, *Cataglyphis velox*. Journal of Comparative Physiology A *205*, 177–189. https://doi.org/10.1007/s00359-019-01316-6.

Freas, C.A., Narendra, A., and Cheng, K. (2017). Compass cues used by a nocturnal bull ant, *Myrmecia midas*. The Journal of Experimental Biology *220*, 1578–1585. https://doi.org/10.1242/jeb.152967.

Freas, C.A., Wystrach, A., Narendra, A., and Cheng, K. (2018). The View from the Trees: Nocturnal Bull Ants, *Myrmecia midas*, Use the Surrounding Panorama While Descending from Trees. Frontiers in Psychology *9:16*. https://doi.org/10.3389/fpsyg.2018.00016.

Freas, C.A., Plowes, N.J.R., and Spetch, M.L. (2019). Not just going with the flow: foraging ants attend to polarised light even while on the pheromone trail. Journal of Comparative Physiology A *205*, 755–767. https://doi.org/10.1007/s00359-019-01363-z.

von Frisch, K. (1946). Die Tänze der Bienen. Österreichische Zoologische Zeitschrift *1*, 1–48. .

von Frisch, K. (1967). The Dance Language and Orientation of Bees (Cambridge, Mass. USA: Harvard University Press).

Fry, S.N., Müller, P., Baumann, H.-J., Straw, A.D., Bichsel, M., and Robert, D. (2004). Context-dependent stimulus presentation to freely moving animals in 3D. Journal of Neuroscience Methods *135*, 149–157. https://doi.org/10.1016/j.jneumeth.2003.12.012.

Fry, S.N., Rohrseitz, N., Straw, A.D., and Dickinson, M.H. (2008). TrackFly: Virtual reality for a behavioral system analysis in free-flying fruit flies. Journal of Neuroscience Methods *171*, 110–117. https://doi.org/10.1016/j.jneumeth.2008.02.016.

Fukushi, T. (2001). Homing in wood ants. Journal of Experimental Biology *204*, 2063–2072. .

Fukushi, T. (2004). Navigation in wood ants *Formica japonica*: context dependent use of landmarks. Journal of Experimental Biology *207*, 3431–3439. https://doi.org/10.1242/jeb.01159.

Galizia, C.G. (2014). Olfactory coding in the insect brain: data and conjectures. European Journal of Neuroscience *39*, 1784–1795. https://doi.org/10.1111/ejn.12558.

Gambling, D.J., and Billard, B. (1967). A Study of the Polarization of Skylight. Australian Journal of Physics *20*, 675–681. .

Gkanias, E., Risse, B., Mangan, M., and Webb, B. (2019). From skylight input to behavioural output: a computational model of the insect polarised light compass. PLOS Computational Biology *15:7*. https://doi.org/10.1371/journal.pcbi.1007123.

Goldschmidt, D., Manoonpong, P., and Dasgupta, S. (2017). A Neurocomputational Model of Goal-Directed Navigation in Insect-Inspired Artificial Agents. Frontiers in Neurorobotics *11*. https://doi.org/10.3389/fnbot.2017.00020.

Gould, J.L. (1986). The Locale Map of Honey Bees: Do Insects Have Cognitive Maps? Science *232*, 861–863. https://doi.org/10.1126/science.232.4752.861.

Graham, P., and Cheng, K. (2009a). Ants use the panoramic skyline as a visual cue during navigation. Current Biology *19*, R935–R937. https://doi.org/10.1016/j.cub.2009.08.015.

Graham, P., and Cheng, K. (2009b). Which portion of the natural panorama is used for view-based navigation in the Australian desert ant? Journal of Comparative Physiology A *195*, 681–689. https://doi.org/10.1007/s00359-009-0443-6.

Graham, P., and Philippides, A. (2017). Vision for navigation: What can we learn from ants? Arthropod Structure & Development *46*, 718–722. https://doi.org/10.1016/j.asd.2017.07.001.

Graham, P., Fauria, K., and Collett, T.S. (2003). The influence of beacon-aiming on the routes of wood ants. Journal of Experimental Biology *206*, 535–541. https://doi.org/10.1242/jeb.00115.

Graham, P., Durier, V., and Collett, T.S. (2004). The binding and recall of snapshot memories in wood ants (*Formica rufa* L.). Journal of Experimental Biology *207*, 393–398. https://doi.org/10.1242/jeb.00771.

Graham, P., Durier, V., and Collett, T. (2007). The co-activation of snapshot memories in wood ants. Journal of Experimental Biology *210*, 2128–2136. https://doi.org/10.1242/jeb.002634.

Graham, P., Philippides, A., and Baddeley, B. (2010). Animal Cognition: Multi-modal Interactions in Ant Learning. Current Biology *20*, R639–R640. https://doi.org/10.1016/j.cub.2010.06.018.

Gray, J.R., Pawlowski, V., and Willis, M.A. (2002). A method for recording behavior and multineuronal CNS activity from tethered insects flying in virtual space. Journal of Neuroscience Methods *120*, 211–223. https://doi.org/10.1016/S0165-0270(02)00223-6.

Greiner, B., Ribi, W.A., Wcislo, W.T., and Warrant, E.J. (2004). Neural organisation in the first optic ganglion of the nocturnal bee *Megalopta genalis*. Cell and Tissue Research *318*, 429–437. https://doi.org/10.1007/s00441-004-0945-z.

Greiner, B., Narendra, A., Reid, S.F., Dacke, M., Ribi, W.A., and Zeil, J. (2007). Eye structure correlates with distinct foraging-bout timing in primitive ants. Current Biology *17*, R879–R880. https://doi.org/10.1016/j.cub.2007.08.015.

Groh, C., and Rössler, W. (2011). Comparison of microglomerular structures in the mushroom body calyx of neopteran insects. Arthropod Structure & Development *40*, 358–367. https://doi.org/10.1016/j.asd.2010.12.002.

Gronenberg, W. (1999). Modality-Specific Segregation of Input to Ant Mushroom Bodies. Brain, Behavior and Evolution *54*, 85–95. https://doi.org/10.1159/000006615.

Gronenberg, W. (2001). Subdivisions of hymenopteran mushroom body calyces by their afferent supply. The Journal of Comparative Neurology *435*, 474–489. https://doi.org/10.1002/cne.1045.

Gronenberg, W. (2008). Structure and function of ant (Hymenoptera: Formicidae) brains: Strength in numbers. Myrmecological News *11*, 25–36. .

Hadeln, J., Hensgen, R., Bockhorst, T., Rosner, R., Heidasch, R., Pegel, U., Quintero Pérez, M., and Homberg, U. (2020). Neuroarchitecture of the central complex of the desert locust: Tangential neurons. Journal of Comparative Neurology *528*, 906–934. https://doi.org/10.1002/cne.24796.

Hamanaka, Y., Shibasaki, H., Kinoshita, M., and Arikawa, K. (2013). Neurons innervating the lamina in the butterfly, *Papilio xuthus*. Journal of Comparative Physiology A *199*, 341–351. https://doi.org/10.1007/s00359-013-0798-6.

Hardcastle, B.J., Omoto, J.J., Kandimalla, P., Nguyen, B.-C.M., Keleş, M.F., Boyd, N.K., Hartenstein, V., and Frye, M.A. (2021). A visual pathway for skylight polarization processing in Drosophila. ELife *10*. https://doi.org/10.7554/eLife.63225.

Harris, R.A., Graham, P., and Collett, T.S. (2007). Visual Cues for the Retrieval of Landmark Memories by Navigating Wood Ants. Current Biology *17*, 93–102. https://doi.org/10.1016/j.cub.2006.10.068.

Hartmann, G., and Wehner, R. (1995). The ant's path integration system: a neural architecture. Biological Cybernetics *73*, 483–497. .

Heinze, S., and Homberg, U. (2008). Neuroarchitecture of the central complex of the desert locust: Intrinsic and columnar neurons. The Journal of Comparative Neurology *511*, 454–478. https://doi.org/10.1002/cne.21842.

Heinze, S., and Pfeiffer, K. (2018). Editorial: The Insect Central Complex—From Sensory Coding to Directing Movement. Frontiers in Behavioral Neuroscience *12:156*. https://doi.org/10.3389/fnbeh.2018.00156.

Heinze, S., Narendra, A., and Cheung, A. (2018). Principles of Insect Path Integration. Current Biology *28*, R1043–R1058. https://doi.org/10.1016/j.cub.2018.04.058.

Held, M., Berz, A., Hensgen, R., Muenz, T.S., Scholl, C., Rössler, W., Homberg, U., and Pfeiffer, K. (2016). Microglomerular Synaptic Complexes in the Sky-Compass Network of the Honeybee Connect Parallel Pathways from the Anterior Optic Tubercle to the Central Complex. Frontiers in Behavioral Neuroscience *10*, 1–15. https://doi.org/10.3389/fnbeh.2016.00186.

Hess, D., Koch, J., and Ronacher, B. (2009). Desert ants do not rely on sky compass information for the perception of inclined path segments. Journal of Experimental Biology *212*, 1528–1534. https://doi.org/10.1242/jeb.027961.

Hoinville, T., and Wehner, R. (2018). Optimal multiguidance integration in insect navigation. Proceedings of the National Academy of Sciences *115*, 2824–2829. https://doi.org/10.1073/pnas.1721668115.

Homberg, U. (2008). Evolution of the central complex in the arthropod brain with respect to the visual system. Arthropod Structure & Development *37*, 347–362. https://doi.org/10.1016/j.asd.2008.01.008.

Honegger, H.-W. (1978). Sustained and transient responding units in the medulla of the cricket *Gryllus campestris*. J. Comp. Physiol. *125*, 259–266. https://doi.org/10.1007/BF00656604.

Honkanen, A., Immonen, E.-V., Salmela, I., Heimonen, K., and Weckström, M. (2017). Insect photoreceptor adaptations to night vision. Philosophical Transactions of the Royal Society B: Biological Sciences *372*, 20160077. https://doi.org/10.1098/rstb.2016.0077.

Honkanen, A., Adden, A., da Silva Freitas, J., and Heinze, S. (2019). The insect central complex and the neural basis of navigational strategies. The Journal of Experimental Biology *222*, jeb188854. https://doi.org/10.1242/jeb.188854.

Huber, R., and Knaden, M. (2015). Egocentric and geocentric navigation during extremely long foraging paths of desert ants. Journal of Comparative Physiology A *201*, 609–616. https://doi.org/10.1007/s00359-015-0998-3.

Huerta, R. (2013). Learning pattern recognition and decision making in the insect brain. In Physics, Computation and the Mind - Advances and Challenges at Interfaces: AIP Conf. Proc. 1510, (La Herradura, Spain), pp. 101–119.

Hulse, B.K., Haberkern, H., Franconville, R., Turner-Evans, D.B., Takemura, S., Wolff, T., Noorman, M., Dreher, M., Dan, C., Parekh, R., et al. (2021). A connectome of the *Drosophila* central complex reveals motif suitable for flexible

navigation and context-dependent action selection. ELife *eLife 2021;10:e66039*, 360. https://doi.org/10.7554/eLife.66039.

Ichinose, T., Kanno, M., Wu, H., Yamagata, N., Sun, H., Abe, A., and Tanimoto, H. (2021). Mushroom body output differentiates memory processes and distinct memory-guided behaviors. Current Biology *31*, 1–9. https://doi.org/10.1016/j.cub.2020.12.032.

Ito, K., Shinomiya, K., Ito, M., Armstrong, J.D., Boyan, G., Hartenstein, V., Harzsch, S., Heisenberg, M., Homberg, U., Jenett, A., et al. (2014). A Systematic Nomenclature for the Insect Brain. Neuron *81*, 755–765. https://doi.org/10.1016/j.neuron.2013.12.017.

Jayatilaka, P. (2014). Individual foraging careers of the jack jumper ant, *Myrmecia croslandi*. PhD Thesis. The Australian National University.

Jayatilaka, P., Narendra, A., Reid, S.F., Cooper, P., and Zeil, J. (2011). Different effects of temperature on foraging activity schedules in sympatric *Myrmecia* ants. Journal of Experimental Biology *214*, 2730–2738. https://doi.org/10.1242/jeb.053710.

Jayatilaka, P., Raderschall, C.A., Narendra, A., and Zeil, J. (2014). Individual foraging patterns of the jack jumper ant *Myrmecia croslandi* (Hymenoptera: Formicidae). Myrmecological News 75–83. .

Jayatilaka, P., Murray, T., Narendra, A., and Zeil, J. (2018). The choreography of learning walks in the Australian jack jumper ant *Myrmecia croslandi*. The Journal of Experimental Biology *221*, jeb185306. https://doi.org/10.1242/jeb.185306.

Johnson, B.R., Borowiec, M.L., Chiu, J.C., Lee, E.K., Atallah, J., and Ward, P.S. (2013). Phylogenomics Resolves Evolutionary Relationships among Ants, Bees, and Wasps. Current Biology *23*, 2058–2062. https://doi.org/10.1016/j.cub.2013.08.050.

el Jundi, B., Pfeiffer, K., Heinze, S., and Homberg, U. (2014). Integration of polarization and chromatic cues in the insect sky compass. Journal of Comparative Physiology A *200*, 575–589. https://doi.org/10.1007/s00359-014-0890-6.

Kakaria, K.S., and de Bivort, B.L. (2017). Ring Attractor Dynamics Emerge from a Spiking Model of the Entire Protocerebral Bridge. Frontiers in Behavioral Neuroscience *11:8*. https://doi.org/10.3389/fnbeh.2017.00008.

Kamhi, J.F., Barron, A.B., and Narendra, A. (2020). Vertical Lobes of the Mushroom Bodies Are Essential for View-Based Navigation in Australian *Myrmecia* Ants. Current Biology *30*, 3432-3437.e3. https://doi.org/10.1016/j.cub.2020.06.030.

Kaushik, P.K., Renz, M., and Olsson, S.B. (2020). Characterizing long-range search behavior in Diptera using complex 3D virtual environments. Proceedings

of the National Academy of Sciences *117*, 12201–12207. https://doi.org/10.1073/pnas.1912124117.

Kelber, A., Balkenius, A., and Warrant, E.J. (2002). Scotopic colour vision in nocturnal hawkmoths. Nature *419*, 922–925. https://doi.org/10.1038/nature01065.

Kelber, A., Vorobyev, M., and Osorio, D. (2003). Animal colour vision – behavioural tests and physiological concepts. Biological Reviews of the Cambridge Philosophical Society *78*, 81–118. https://doi.org/10.1017/S1464793102005985.

Kenyon, F.C. (1896). The brain of the bee. A preliminary contribution to the morphology of the nervous system of the arthropoda. Journal of Comparative Neurology *6*, 133–210. https://doi.org/10.1002/cne.910060302.

Kim, S.S., Rouault, H., Druckmann, S., and Jayaraman, V. (2017). Ring attractor dynamics in the *Drosophila* central brain. Science *356*, 849–853. https://doi.org/10.1126/science.aal4835.

Kinoshita, M., Pfeiffer, K., and Homberg, U. (2007). Spectral properties of identified polarized-light sensitive interneurons in the brain of the desert locust *Schistocerca gregaria*. Journal of Experimental Biology *210*, 1350–1361. https://doi.org/10.1242/jeb.02744.

Kirschfeld, K. (1976). The Resolution of Lens and Compound Eyes. In Neural Principles in Vision, F. Zettler, and R. Weiler, eds. (Berlin, Heidelberg: Springer Berlin Heidelberg), pp. 354–370.

Knaden, M., and Graham, P. (2016). The Sensory Ecology of Ant Navigation: From Natural Environments to Neural Mechanisms. Annual Review of Entomology *61*, 63–76. https://doi.org/10.1146/annurev-ento-010715-023703.

Kodzhabashev, A., and Mangan, M. (2015). Route following without scanning. In Biomimetic and Biohybrid Systems: 4th International Conference, (Barcelona, Spain: Springer Berlin Heidelberg), pp. 199–210.

Koenderink, J. (1986). Optic flow. Vision Research *26*, 161–179. https://doi.org/10.1016/0042-6989(86)90078-7.

Koenig, S., Wolf, R., and Heisenberg, M. (2016). Visual Attention in Flies—Dopamine in the Mushroom Bodies Mediates the After-Effect of Cueing. PLOS ONE *11*, e0161412. https://doi.org/10.1371/journal.pone.0161412.

Kohler, M., and Wehner, R. (2005). Idiosyncratic route-based memories in desert ants, *Melophorus bagoti*: How do they interact with path-integration vectors? Neurobiology of Learning and Memory *83*, 1–12. https://doi.org/10.1016/j.nlm.2004.05.011.

Kollmeier, T., Röben, F., Schenck, W., and Möller, R. (2007). Spectral contrasts for landmark navigation. Journal of the Optical Society of America A *24*, 1. https://doi.org/10.1364/JOSAA.24.000001.

van der Kooi, C.J., Stavenga, D.G., Arikawa, K., Belušič, G., and Kelber, A. (2021). Evolution of Insect Color Vision: From Spectral Sensitivity to Visual Ecology. Annual Review of Entomology *66*, 435–461. https://doi.org/10.1146/annurev-ento-061720-071644.

Krapp, H.G., Hengstenberg, B., and Hengstenberg, R. (1998). Dendritic Structure and Receptive-Field Organization of Optic Flow Processing Interneurons in the Fly. Journal of Neurophysiology *79*, 1902–1917. https://doi.org/10.1152/jn.1998.79.4.1902.

Krapp, H.G., Hengstenberg, R., and Egelhaaf, M. (2001). Binocular Contributions to Optic Flow Processing in the Fly Visual System. Journal of Neurophysiology *85*, 724–734. https://doi.org/10.1152/jn.2001.85.2.724.

Kruger, C.J.C. (2002). Constrained Cubic Spline Interpolation for Chemical Engineering Applications (Korf Industrues Ltd.).

Labhart, T. (1980). Specialized photoreceptors at the dorsal rim of the honeybee's compound eye: Polarizational and angular sensitivity. Journal of Comparative Physiology ? A *141*, 19–30. https://doi.org/10.1007/BF00611874.

Labhart, T. (1999). How polarisation-sensitive interneurones of crickets see the polarisation pattern of the sky: a field study with an opeto-electronic model neurone. Journal of Experimental Biology *202*, 757–770. .

Labhart, T. (2016). Can invertebrates see the e-vector of polarization as a separate modality of light? The Journal of Experimental Biology *219*, 3844–3856. https://doi.org/10.1242/jeb.139899.

Labhart, T., and Meyer, E.P. (1999). Detectors for polarized skylight in insects: a survey of ommatidial specializations in the dorsal rim area of the compound eye. Microscopy Research and Technique *47*, 368–379. https://doi.org/10.1002/(SICI)1097-0029(19991215)47:6<368::AID-JEMT2>3.0.CO;2-Q.

Lambrinos, D., Möller, R., Labhart, T., Pfeifer, R., and Wehner, R. (2000). A mobile robot employing insect strategies for navigation. Robotics and Autonomous Systems *30*, 39–64. https://doi.org/10.1016/S0921-8890(99)00064-0.

Land, M.F. (1997). Visual Acuity in Insects. Annual Review of Entomology *42*, 147–177. https://doi.org/10.1146/annurev.ento.42.1.147.

Land, M.F., and Nilsson, D.-E. (2002). Animal eyes (New York: Oxford University Press).

Landler, L., Ruxton, G.D., and Malkemper, E.P. (2018). Circular data in biology: advice for effectively implementing statistical procedures. Behavioral Ecology and Sociobiology *72*. https://doi.org/10.1007/s00265-018-2538-y.

Landler, L., Ruxton, G.D., and Malkemper, E.P. (2019). The Hermans–Rasson test as a powerful alternative to the Rayleigh test for circular statistics in biology. BMC Ecology *19*. https://doi.org/10.1186/s12898-019-0246-8.

Laughlin, S. (1981). A Simple Coding Procedure Enhances a Neuron's Information Capacity. Z. Naturforsch C *36*, 910–912. https://doi.org/10.1515/znc-1981-9-1040.

Le Möel, F., and Wystrach, A. (2020). Opponent processes in visual memories: A model of attraction and repulsion in navigating insects' mushroom bodies. PLOS Computational Biology *16*, e1007631. https://doi.org/10.1371/journal.pcbi.1007631.

Le Moël, F., Stone, T., Lihoreau, M., Wystrach, A., and Webb, B. (2019). The Central Complex as a Potential Substrate for Vector Based Navigation. Frontiers in Psychology *10:690*. https://doi.org/10.3389/fpsyg.2019.00690.

Lee, D.N., and Kalmus, H. (1980). The Optic Flow Field: The Foundation of Vision [and Discussion]. Philosophical Transactions of the Royal Society B: Biological Sciences *290*, 169–179. .

Legge, E.L.G., Wystrach, A., Spetch, M.L., and Cheng, K. (2014). Combining sky and earth: desert ants (*Melophorus bagoti*) show weighted integration of celestial and terrestrial cues. Journal of Experimental Biology *217*, 4159–4166. https://doi.org/10.1242/jeb.107862.

Lei, Z., Chen, K., Li, H., Liu, H., and Guo, A. (2013). The GABA system regulates the sparse coding of odors in the mushroom bodies of Drosophila. Biochemical and Biophysical Research Communications *436*, 35–40. https://doi.org/10.1016/j.bbrc.2013.05.036.

Lent, D.D., Graham, P., and Collett, T.S. (2013). Phase-Dependent Visual Control of the Zigzag Paths of Navigating Wood Ants. Current Biology *23*, 2393–2399. https://doi.org/10.1016/j.cub.2013.10.014.

Li, F., Lindsey, J.W., Marin, E.C., Otto, N., Dreher, M., Dempsey, G., Stark, I., Bates, A.S., Pleijzier, M.W., Schlegel, P., et al. (2020a). The connectome of the adult Drosophila mushroom body provides insights into function. ELife *9:e62576*, 86. https://doi.org/10.7554/eLife.62576.

Li, J., Mahoney, B.D., Jacob, M.S., and Caron, S.J.C. (2020b). Visual Input into the *Drosophila melanogaster* Mushroom Body. Cell Reports *32*, 108138. https://doi.org/10.1016/j.celrep.2020.108138.

Lindemann, J.P., Kern, R., Michaelis, C., Meyer, P., van Hateren, J.H., and Egelhaaf, M. (2003). FliMax, a novel stimulus device for panoramic and

highspeed presentation of behaviourally generated optic flow. Vision Research *43*, 779–791. https://doi.org/10.1016/S0042-6989(03)00039-7.

MacNeil, M.A., and Masland, R.H. (1998). Extreme Diversity among Amacrine Cells: Implications for Function. Neuron *20*, 971–982. https://doi.org/10.1016/S0896-6273(00)80478-X.

Makas, A.S. (1962). Film polarizer for visible and ultraviolet radiation. Journal of the Optical Society of America A *52*, 43–44. .

Mangan, M., and Webb, B. (2012). Spontaneous formation of multiple routes in individual desert ants (*Cataglyphis velox*). Behavioral Ecology *23*, 944–954. https://doi.org/10.1093/beheco/ars051.

Masland, R.H. (2001). The fundamental plan of the retina. Nature Neuroscience *4*, 877–886. https://doi.org/10.1038/nn0901-877.

Masland, R.H. (2012). The Neuronal Organization of the Retina. Neuron *76*, 266–280. https://doi.org/10.1016/j.neuron.2012.10.002.

Masland, R.H., and Martin, P.R. (2007). The unsolved mystery of vision. Current Biology *17*, R577–R582. https://doi.org/10.1016/j.cub.2007.05.040.

Masland, R.H., and Raviola, E. (2000). Confronting Complexity: Strategies for Understanding the Microcircuitry of the Retina. Annual Review of Neuroscience *23*, 249–284. https://doi.org/10.1146/annurev.neuro.23.1.249.

Meglič, A., Ilić, M., Pirih, P., Škorjanc, A., Wehling, M.F., Kreft, M., and Belušič, G. (2019). Horsefly object-directed polarotaxis is mediated by a stochastically distributed ommatidial subtype in the ventral retina. Proceedings of the National Academy of Sciences *116*, 21843–21853. https://doi.org/10.1073/pnas.1910807116.

Menzel, R., and Blakers, M. (1975). Functional organisation of an insect ommatidium with fused rhabdom. Cytobiologie *11*, 279–298. .

Menzel, R., and Giurfa, M. (2001). Cognitive architecture of a mini-brain: the honeybee. Trends in Cognitive Sciences *5*, 62–71. https://doi.org/10.1016/S1364-6613(00)01601-6.

Menzel, R., Brandt, R., Gumbert, A., Komischke, B., and Kunze, J. (2000). Two spatial memories for honeybee navigation. Proceedings of the Royal Society of London. Series B: Biological Sciences *267*, 961–968. https://doi.org/10.1098/rspb.2000.1097.

Menzel, R., Greggers, U., Smith, A., Berger, S., Brandt, R., Brunke, S., Bundrock, G., Hulse, S., Plumpe, T., Schaupp, F., et al. (2005). Honey bees navigate according to a map-like spatial memory. Proceedings of the National Academy of Sciences *102*, 3040–3045. https://doi.org/10.1073/pnas.0408550102.

Meyer, E.P., and Domanico, V. (1999). Microvillar orientation in the photoreceptors of the ant *Cataglyphis bicolor*. Cell and Tissue Research *295*, 355–361. https://doi.org/10.1007/s004410051242.

Michener, C.D. (1969). Comparative Social Behavior of Bees. Annual Review of Entomology *14*, 299–342. https://doi.org/10.1146/annurev.en.14.010169.001503.

Mittelstaedt, H. (1962). Control Systems of Orientation in Insects. Annual Review of Entomology *7*, 177–198. https://doi.org/10.1146/annurev.en.07.010162.001141.

Mittelstaedt, M.-L., and Mittelstaedt, H. (1980). Homing by path integration in a mammal. Naturwissenschaften *67*, 566–567. https://doi.org/10.1007/BF00450672.

Mizunami, M. (1995). Functional diversity of neural organization in insect ocellar systems. Vision Research *35*, 443–452. https://doi.org/10.1016/0042-6989(94)00192-O.

Mobbs, P.G. (1984). Neural networks in the mushroom bodies of the honeybee. Journal of Insect Physiology *30*, 43–58. https://doi.org/10.1016/0022-1910(84)90107-0.

Möller, R. (2000). Insect visual homing strategies in a robot with analog processing. Biological Cybernetics *83*, 231–243. .

Möller, R. (2002). Insects Could Exploit UV–Green Contrast for Landmark Navigation. Journal of Theoretical Biology *214*, 619–631. https://doi.org/10.1006/jtbi.2001.2484.

Möller, R. (2012). A model of ant navigation based on visual prediction. Journal of Theoretical Biology *305*, 118–130. https://doi.org/10.1016/j.jtbi.2012.04.022.

Morante, J., and Desplan, C. (2008). The Color-Vision Circuit in the Medulla of *Drosophila*. Current Biology *18*, 553–565. https://doi.org/10.1016/j.cub.2008.02.075.

Müller, M., and Wehner, R. (1988). Path integration in desert ants, *Cataglyphis fortis*. Proceedings of the National Academy of Sciences *85*, 5287–5290. https://doi.org/10.1073/pnas.85.14.5287.

Müller, M., and Wehner, R. (2010). Path Integration Provides a Scaffold for Landmark Learning in Desert Ants. Current Biology *20*, 1368–1371. https://doi.org/10.1016/j.cub.2010.06.035.

Munz, T. (2005). The Bee Battles: Karl von Frisch, Adrian Wenner and the Honey Bee Dance Language Controversy. Journal of the History of Biology *38*, 535–570. https://doi.org/10.1007/s10739-005-0552-1.

Murray, T., and Zeil, J. (2017). Quantifying navigational information: The catchment volumes of panoramic snapshots in outdoor scenes. PLOS ONE *12*, e0187226. https://doi.org/10.1371/journal.pone.0187226.

Murray, T., Kócsi, Z., Dahmen, H., Narendra, A., Le Möel, F., Wystrach, A., and Zeil, J. (2020). The role of attractive and repellent scene memories in ant homing ( *Myrmecia croslandi* ). The Journal of Experimental Biology *223*, jeb210021. https://doi.org/10.1242/jeb.210021.

Muser, B., Sommer, S., Wolf, H., and Wehner, R. (2005). Foraging ecology of the thermophilic Australian desert ant, *Melophorus bagoti*. Australian Journal of Zoology *53*, 301. https://doi.org/10.1071/ZO05023.

Namiki, S., and Kanzaki, R. (2016). The neurobiological basis of orientation in insects: insights from the silkmoth mating dance. Current Opinion in Insect Science *15*, 16–26. https://doi.org/10.1016/j.cois.2016.02.009.

Narendra, A. (2007a). Homing strategies of the Australian desert ant *Melophorus bagoti* I. Proportional path-integration takes the ant half-way home. Journal of Experimental Biology *210*, 1798–1803. https://doi.org/10.1242/jeb.02768.

Narendra, A. (2007b). Homing strategies of the Australian desert ant *Melophorus bagoti*. II. Interaction of the path integrator with visual cue information. Journal of Experimental Biology *210*, 2212–2212. https://doi.org/10.1242/jeb.02791.

Narendra, A., and Ramirez-Esquivel, F. (2017). Subtle changes in the landmark panorama disrupt visual navigation in a nocturnal bull ant. Philosophical Transactions of the Royal Society B: Biological Sciences *372*, 20160068. https://doi.org/10.1098/rstb.2016.0068.

Narendra, A., and Ribi, W.A. (2017). Ocellar structure is driven by the mode of locomotion and activity time in *Myrmecia* ants. The Journal of Experimental Biology *220*, 4383–4390. https://doi.org/10.1242/jeb.159392.

Narendra, A., Reid, S.F., and Hemmi, J.M. (2010). The twilight zone: ambient light levels trigger activity in primitive ants. Proceedings of the Royal Society B: Biological Sciences *277*, 1531–1538. https://doi.org/10.1098/rspb.2009.2324.

Narendra, A., Reid, S.F., Greiner, B., Peters, R.A., Hemmi, J.M., Ribi, W.A., and Zeil, J. (2011). Caste-specific visual adaptations to distinct daily activity schedules in Australian *Myrmecia* ants. Proceedings of the Royal Society B: Biological Sciences *278*, 1141–1149. https://doi.org/10.1098/rspb.2010.1378.

Narendra, A., Alkaladi, A., Raderschall, C.A., Robson, S.K.A., and Ribi, W.A. (2013a). Compound Eye Adaptations for Diurnal and Nocturnal Lifestyle in the Intertidal Ant, *Polyrhachis sokolova*. PLoS ONE *8*, e76015. https://doi.org/10.1371/journal.pone.0076015.

Narendra, A., Gourmaud, S., and Zeil, J. (2013b). Mapping the navigational knowledge of individually foraging ants, *Myrmecia croslandi*. Proceedings of the Royal Society B: Biological Sciences *280*, 20130683. https://doi.org/10.1098/rspb.2013.0683.

Narendra, A., Reid, S.F., and Raderschall, C.A. (2013c). Navigational Efficiency of Nocturnal *Myrmecia* Ants Suffers at Low Light Levels. PLoS ONE *8*, e58801. https://doi.org/10.1371/journal.pone.0058801.

Narendra, A., Greiner, B., Ribi, W.A., and Zeil, J. (2016). Light and dark adaptation mechanisms in the compound eyes of *Myrmecia* ants that occupy discrete temporal niches. The Journal of Experimental Biology *219*, 2435–2442. https://doi.org/10.1242/jeb.142018.

Narendra, A., Kamhi, J.F., and Ogawa, Y. (2017). Moving in Dim Light: Behavioral and Visual Adaptations in Nocturnal Ants. Integrative and Comparative Biology *57*, 1104–1116. https://doi.org/10.1093/icb/icx096.

O'Carroll, D.C. (1993). Feature detecting neurons in dragonflies. Nature *362*, 541–543. .

Ogata, K., and Taylor, R.W. (1991). Ants of the genus *Myrmecia* Fabricius: a preliminary review and key to the named species (Hymenoptera: Formicidae: Myrmeciinae). Journal of Natural History *25*, 1623–1673. https://doi.org/10.1080/00222939100771021.

Ogawa, Y., Falkowski, M., Narendra, A., Zeil, J., and Hemmi, J.M. (2015). Three spectrally distinct photoreceptors in diurnal and nocturnal Australian ants. Proceedings of the Royal Society B: Biological Sciences *282*, 20150673. https://doi.org/10.1098/rspb.2015.0673.

Ogawa, Y., Ribi, W., Zeil, J., and Hemmi, J.M. (2017). Regional differences in the preferred e-vector orientation of honeybee ocellar photoreceptors. The Journal of Experimental Biology *220*, 1701–1708. https://doi.org/10.1242/jeb.156109.

Ogawa, Y., Ryan, L.A., Palavalli-Nettimi, R., Seeger, O., Hart, N.S., and Narendra, A. (2019). Spatial Resolving Power and Contrast Sensitivity Are Adapted for Ambient Light Conditions in Australian *Myrmecia* Ants. Frontiers in Ecology and Evolution *7:18*. https://doi.org/10.3389/fevo.2019.00018.

O'Keefe, J. (1991). An allocentric spatial model for the hippocampal cognitive map. Hippocampus *1*, 230–235. .

Osorio, D. (1986). Directionally selective cells in the locust medulla. J. Comp. Physiol. *159*, 841–847. https://doi.org/10.1007/BF00603737.

Owald, D., Felsenberg, J., Talbot, C.B., Das, G., Perisse, E., Huetteroth, W., and Waddell, S. (2015). Activity of Defined Mushroom Body Output Neurons Underlies Learned Olfactory Behavior in *Drosophila*. Neuron *86*, 417–427. https://doi.org/10.1016/j.neuron.2015.03.025.

Owens, A.C.S., and Lewis, S.M. (2018). The impact of artificial light at night on nocturnal insects: A review and synthesis. Ecology and Evolution *8*, 11337–11358. https://doi.org/10.1002/ece3.4557.

Papadopoulou, M., Cassenaer, S., Nowotny, T., and Laurent, G. (2011). Normalization for Sparse Encoding of Odors by a Wide-Field Interneuron. Science *332*, 721–725. https://doi.org/10.1126/science.1201835.

Paulk, A.C., Dacks, A.M., Phillips-Portillo, J., Fellous, J.-M., and Gronenberg, W. (2009). Visual Processing in the Central Bee Brain. Journal of Neuroscience *29*, 9987–9999. https://doi.org/10.1523/JNEUROSCI.1325-09.2009.

Paulk, A.C., Stacey, J.A., Pearson, T.W.J., Taylor, G.J., Moore, R.J.D., Srinivasan, M.V., and van Swinderen, B. (2014). Selective attention in the honeybee optic lobes precedes behavioral choices. Proc. Natl. Acad. Sci. U.S.A. *111*, 5006–5011. https://doi.org/10.1073/pnas.1323297111.

Peckmezian, T., and Taylor, P.W. (2015). A virtual reality paradigm for the study of visually mediated behaviour and cognition in spiders. Animal Behaviour *107*, 87–95. https://doi.org/10.1016/j.anbehav.2015.06.018.

Pegel, U., Pfeiffer, K., and Homberg, U. (2018). Integration of celestial compass cues in the central complex of the locust brain. The Journal of Experimental Biology *221*, jeb171207. https://doi.org/10.1242/jeb.171207.

Pegel, U., Pfeiffer, K., Zittrell, F., Scholtyssek, C., and Homberg, U. (2019). Two Compasses in the Central Complex of the Locust Brain. The Journal of Neuroscience *39*, 3070–3080. https://doi.org/10.1523/JNEUROSCI.0940-18.2019.

Peng, F., and Chittka, L. (2017). A Simple Computational Model of the Bee Mushroom Body Can Explain Seemingly Complex Forms of Olfactory Learning and Memory. Current Biology *27*, 224–230. https://doi.org/10.1016/j.cub.2016.10.054.

Pfeffer, S.E., and Wittlinger, M. (2016). Optic flow odometry operates independently of stride integration in carried ants. Science *353*, 1155–1157. https://doi.org/10.1126/science.aaf9754.

Philippides, A., Baddeley, B., Cheng, K., and Graham, P. (2011). How might ants use panoramic views for route navigation? Journal of Experimental Biology *214*, 445–451. https://doi.org/10.1242/jeb.046755.

Pomozi, I., Horváth, G., and Wehner, R. (2001). How the clear-sky angle of polarization pattern continues underneath clouds: full-sky measurements and implications for animal orientation. Journal of Experimental Biology *204*, 2933–2944. .

Portelli, G., Ruffier, F., and Franceschini, N. (2010). Honeybees change their height to restore their optic flow. Journal of Comparative Physiology A *196*, 307–313. https://doi.org/10.1007/s00359-010-0510-z.

Raderschall, C.A., Narendra, A., and Zeil, J. (2016). Head roll stabilisation in the nocturnal bull ant *Myrmecia pyriformis*: implications for visual navigation. The Journal of Experimental Biology *219*, 1449–1457. https://doi.org/10.1242/jeb.134049.

Ramirez-Esquivel, F., Zeil, J., and Narendra, A. (2014). The antennal sensory array of the nocturnal bull ant *Myrmecia pyriformis*. Arthropod Structure & Development *43*, 543–558. https://doi.org/10.1016/j.asd.2014.07.004.

Reid, S.F. (2010). Life in the dark: Vision and navigation in a nocturnal bull ant. PhD Thesis. The Australian National University.

Reid, S.F., Narendra, A., Hemmi, J.M., and Zeil, J. (2011). Polarised skylight and the landmark panorama provide night-active bull ants with compass information during route following. Journal of Experimental Biology *214*, 363–370. https://doi.org/10.1242/jeb.049338.

Reid, S.F., Narendra, A., Taylor, R.W., and Zeil, J. (2013). Foraging ecology of the night-active bull ant *Myrmecia pyriformis*. Australian Journal of Zoology *61*, 170. https://doi.org/10.1071/ZO13027.

Reiser, M.B., and Dickinson, M.H. (2008). A modular display system for insect behavioral neuroscience. Journal of Neuroscience Methods *167*, 127–139. https://doi.org/10.1016/j.jneumeth.2007.07.019.

Ribi, W.A. (1975). The first optic ganglion of the bee. Cell And Tissue Research *165*, 103–111. .

Ribi, W., and Scheel, M. (1981). The second and third optic ganglia of the worker bee: Golgi studies of the neuronal elements in the medulla and lobula. Cell And Tissue Research *221*, 17–43. https://doi.org/10.1007/BF00216567.

Ribi, W., and Zeil, J. (2017). Three-dimensional visualization of ocellar interneurons of the orchid bee *Euglossa imperialis* using micro X-ray computed tomography. Journal of Comparative Neurology *525*, 3581–3595. https://doi.org/10.1002/cne.24260.

Ribi, W., and Zeil, J. (2018). Diversity and common themes in the organization of ocelli in Hymenoptera, Odonata and Diptera. Journal of Comparative Physiology A *204*, 505–517. https://doi.org/10.1007/s00359-018-1258-0.

Ribi, W., Warrant, E., and Zeil, J. (2011). The organization of honeybee ocelli: Regional specializations and rhabdom arrangements. Arthropod Structure & Development *40*, 509–520. https://doi.org/10.1016/j.asd.2011.06.004.

Richter, S., Loesel, R., Purschke, G., Schmidt-Rhaesa, A., Scholtz, G., Stach, T., Vogt, L., Wanninger, A., Brenneis, G., Döring, C., et al. (2010). Invertebrate neurophylogeny: suggested terms and definitions for a neuroanatomical glossary. Frontiers in Zoology *7:29*. https://doi.org/10.1186/1742-9994-7-29.

Rivera-Alba, M., Vitaladevuni, S.N., Mishchenko, Y., Lu, Z., Takemura, S., Scheffer, L., Meinertzhagen, I.A., Chklovskii, D.B., and de Polavieja, G.G. (2011). Wiring Economy and Volume Exclusion Determine Neuronal Placement in the *Drosophila* Brain. Current Biology *21*, 2000–2005. https://doi.org/10.1016/j.cub.2011.10.022.

Ronacher, B., Gallizzi, S., Wohlgemuth, S., and Wehner, R. (2000). Lateral optic flow does not influence distance estimation in the desert ant *Cataglyphis fortis*. Journal of Experimental Biology *203*, 1113–1121. .

Roper, M., Fernando, C., and Chittka, L. (2017). Insect Bio-inspired Neural Network Provides New Evidence on How Simple Feature Detectors Can Enable Complex Visual Generalization and Stimulus Location Invariance in the Miniature Brain of Honeybees. PLOS Computational Biology *13*, e1005333. https://doi.org/10.1371/journal.pcbi.1005333.

Rössler, W. (2019). Neuroplasticity in desert ants (Hymenoptera: Formicidae) – importance for the ontogeny of navigation. Myrmecological News *29*, 1–20. https://doi.org/10.25849/myrmecol.news_029:001.

Rusanen, J., Vähäkainu, A., Weckström, M., and Arikawa, K. (2017). Characterization of the first-order visual interneurons in the visual system of the bumblebee (*Bombus terrestris*). Journal of Comparative Physiology A *203*, 903–913. https://doi.org/10.1007/s00359-017-1201-9.

Rutkowski, A.J., Miller, M.M., Quinn, R.D., and Willis, M.A. (2011). Egomotion estimation with optic flow and air velocity sensors. Biological Cybernetics *104*, 351–367. https://doi.org/10.1007/s00422-011-0440-z.

Ryan, L.A., Cunningham, R., Hart, N.S., and Ogawa, Y. (2020). The buzz around spatial resolving power and contrast sensitivity in the honeybee, *Apis mellifera*. Vision Research *169*, 25–32. https://doi.org/10.1016/j.visres.2020.02.005.

Sabo, C., Chisholm, R., Petterson, A., and Cope, A. (2017). A lightweight, inexpensive robotic system for insect vision. Arthropod Structure & Development *46*, 689–702. https://doi.org/10.1016/j.asd.2017.08.001.

Sanes, J.R., and Zipursky, S.L. (2010). Design Principles of Insect and Vertebrate Visual Systems. Neuron *66*, 15–36. https://doi.org/10.1016/j.neuron.2010.01.018.

Santschi, F. (1911). Observations et remarques critiques sur le mécanisme de l'orientation chez les fourmis. Revue Suisse de Zoologie *19*, 303–338. .

Saumweber, T., Rohwedder, A., Schleyer, M., Eichler, K., Chen, Y., Aso, Y., Cardona, A., Eschbach, C., Kobler, O., Voigt, A., et al. (2018). Functional architecture of reward learning in mushroom body extrinsic neurons of larval *Drosophila*. Nature Communications *9:1104*. https://doi.org/10.1038/s41467-018-03130-1.

Sayre, M.E., Templin, R., Chavez, J., Kempenaers, J., and Heinze, S. (2021). A projectome of the bumblebee central complex. ELife *10:e68911*. https://doi.org/10.7554/eLife.68911.

Schoenemann, B., Pärnaste, H., and Clarkson, E.N.K. (2017). Structure and function of a compound eye, more than half a billion years old. Proceedings of the National Academy of Sciences *114*, 13489–13494. https://doi.org/10.1073/pnas.1716824114.

Scholtz, G., Staude, A., and Dunlop, J.A. (2019). Trilobite compound eyes with crystalline cones and rhabdoms show mandibulate affinities. Nature Communications *10:2503*. https://doi.org/10.1038/s41467-019-10459-8.

Schultheiss, P., Cheng, K., and Reynolds, A.M. (2015). Searching behavior in social Hymenoptera. Learning and Motivation *50*, 59–67. https://doi.org/10.1016/j.lmot.2014.11.002.

Schultheiss, P., Wystrach, A., Schwarz, S., Tack, A., Delor, J., Nooten, S.S., Bibost, A.-L., Freas, C.A., and Cheng, K. (2016a). Crucial role of ultraviolet light for desert ants in determining direction from the terrestrial panorama. Animal Behaviour *115*, 19–28. https://doi.org/10.1016/j.anbehav.2016.02.027.

Schultheiss, P., Stannard, T., Pereira, S., Reynolds, A.M., Wehner, R., and Cheng, K. (2016b). Similarities and differences in path integration and search in two species of desert ants inhabiting a visually rich and a visually barren habitat. Behavioral Ecology and Sociobiology *70*, 1319–1329. https://doi.org/10.1007/s00265-016-2140-0.

Schultheiss, P., Buatois, A., Avarguès-Weber, A., and Giurfa, M. (2017). Using virtual reality to study visual performances of honeybees. Current Opinion in Insect Science *24*, 43–50. https://doi.org/10.1016/j.cois.2017.08.003.

Schultz, T.R. (2000). In search of ant ancestors. Proceedings of the National Academy of Sciences *97*, 14028–14029. https://doi.org/10.1073/pnas.011513798.

Schürch, R., and Ratnieks, F.L.W. (2015). The spatial information content of the honey bee waggle dance. Frontiers in Human Neuroscience *3*, 7. https://doi.org/10.3389/fevo.2015.00022.

Schwarz, S., Narendra, A., and Zeil, J. (2011a). The properties of the visual system in the Australian desert ant *Melophorus bagoti*. Arthropod Structure & Development *40*, 128–134. https://doi.org/10.1016/j.asd.2010.10.003.

Schwarz, S., Albert, L., Wystrach, A., and Cheng, K. (2011b). Ocelli contribute to the encoding of celestial compass information in the Australian desert ant *Melophorus bagoti*. Journal of Experimental Biology *214*, 901–906. https://doi.org/10.1242/jeb.049262.

Schwarz, S., Wystrach, A., and Cheng, K. (2011c). A new navigational mechanism mediated by ant ocelli. Biology Letters *7*, 856–858. https://doi.org/10.1098/rsbl.2011.0489.

Schwarz, S., Mangan, M., Zeil, J., Webb, B., and Wystrach, A. (2017a). How Ants Use Vision When Homing Backward. Current Biology *27*, 401–407. https://doi.org/10.1016/j.cub.2016.12.019.

Schwarz, S., Wystrach, A., and Cheng, K. (2017b). Ants' navigation in an unfamiliar environment is influenced by their experience of a familiar route. Scientific Reports *7*. https://doi.org/10.1038/s41598-017-14036-1.

Schwarz, S., Mangan, M., Webb, B., and Wystrach, A. (2019). Sequence of view matters in route-following ants. In The Fourth International Conference on Invertebrate Vision, (Bäckaskog Castle, Sweden), p.

Schwarz, S., Clement, L., Gkanias, E., and Wystrach, A. (2020). How do backward-walking ants (*Cataglyphis velox*) cope with navigational uncertainty? Animal Behaviour *164*, 133–142. https://doi.org/10.1016/j.anbehav.2020.04.006.

Seelig, J.D., and Jayaraman, V. (2015). Neural dynamics for landmark orientation and angular path integration. Nature *521*, 186–191. https://doi.org/10.1038/nature14446.

Sheehan, Z.B.V., Kamhi, J.F., Seid, M.A., and Narendra, A. (2019). Differential investment in brain regions for a diurnal and nocturnal lifestyle in Australian *Myrmecia* ants. Journal of Comparative Neurology *527*, 1261–1277. https://doi.org/10.1002/cne.24617.

Smith, D., Wessnitzer, J., and Webb, B. (2008). A model of associative learning in the mushroom body. Biological Cybernetics *99*, 89–103. https://doi.org/10.1007/s00422-008-0241-1.

Snyder, A.W., Menzel, R., and Laughlin, S.B. (1973). Structure and function of the fused rhabdom. Journal of Comparative Physiology *87*, 99–135. https://doi.org/10.1007/BF01352157.

Spalthoff, C., Gerdes, R., and Kurtz, R. (2012). Neuronal representation of visual motion and orientation in the fly medulla. Frontiers in Neural Circuits *6:72*, 1–13. https://doi.org/10.3389/fncir.2012.00072.

Srinivasan, M.V., Zhang, S.W., Lehrer, M., and Collett, T.S. (1996). Honeybee Navigation *en route* to the Goal: Visual Flight Control. The Journal of Experimental Biology *199*, 237–244. .

Srinivasan, M.V., Zhang, S.W., and Bidwell, N.J. (1997). Visually Mediated Odometry in Honeybees. The Journal of Experimental Biology 2513–2522. .

Stavenga, D. (2003). Angular and spectral sensitivity of fly photoreceptors. I. Integrated facet lens and rhabdomere optics. Journal of Comparative Physiology A *189*, 1–17. https://doi.org/10.1007/s00359-002-0370-2.

Stavenga, D.G., and Hardie, R.C. (2011). Facets of vision (Springer-Verlag Berlin Heidelberg).

Stavenga, D.G., and van Hateren, J.H. (1991). Focusing by a high-power, low-Fresnel-number lens: the fly facet lens. Journal of the Optical Society of America A *8*, 14–19. https://doi.org/10.1364/JOSAA.8.000014.

Sterling, P., and Laughlin, S. (2015). Principles of neural design (Cambridge, Massachusetts: The MIT Press).

Stöckl, A.L., O'Carroll, D.C., and Warrant, E.J. (2020). Hawkmoth lamina monopolar cells act as dynamic spatial filters to optimize vision at different light levels. Science Advances *6*. https://doi.org/10.1126/sciadv.aaz8645.

Stone, T., Mangan, M., Ardin, P., and Webb, B. (2014). Sky segmentation with ultraviolet images can be used for navigation. In Robotics: Science and Systems X, (Robotics: Science and Systems Foundation), p.

Stone, T., Differt, D., Milford, M., and Webb, B. (2016). Skyline-based localisation for aggressively manoeuvring robots using UV sensors and spherical harmonics. In 2016 IEEE International Conference on Robotics and Automation (ICRA), (Stockholm, Sweden: IEEE), pp. 5615–5622.

Stone, T., Webb, B., Adden, A., Weddig, N.B., Honkanen, A., Templin, R., Wcislo, W., Scimeca, L., Warrant, E., and Heinze, S. (2017). An Anatomically Constrained Model for Path Integration in the Bee Brain. Current Biology *27*, 3069-3085.e11. https://doi.org/10.1016/j.cub.2017.08.052.

Stowers, J.R., Fuhrmann, A., Hofbauer, M., Streinzer, M., Schmid, A., Dickinson, M.H., and Straw, A.D. (2014). Reverse Engineering Animal Vision with Virtual Reality and Genetics. Computer *47*, 38–45. https://doi.org/10.1109/MC.2014.190.

Stowers, J.R., Hofbauer, M., Bastien, R., Griessner, J., Higgins, P., Farooqui, S., Fischer, R.M., Nowikovsky, K., Haubensak, W., Couzin, I.D., et al. (2017). Virtual reality for freely moving animals. Nature Methods *14*, 995–1002. https://doi.org/10.1038/nmeth.4399.

Strausfeld, N.J. (2002). Organization of the honey bee mushroom body: Representation of the calyx within the vertical and gamma lobes. The Journal of Comparative Neurology *450*, 4–33. https://doi.org/10.1002/cne.10285.

Strausfeld, N.J. (2012). Arthropod brains: evolution, functional elegance, and historical significance (Cambridge, Mass: Harvard University Press).

Strausfeld, N.J., and Hirth, F. (2013). Deep Homology of Arthropod Central Complex and Vertebrate Basal Ganglia. Science *340*, 157–161. https://doi.org/10.1126/science.1231828.

Strausfeld, N.J., Hansen, L., Li, Y., Gomez, R.S., and Ito, K. (1998). Evolution, Discovery, and Interpretations of Arthropod Mushroom Bodies. Learning & Memory *5*, 11–37. .

Strausfeld, N.J., Sinakevitch, I., Brown, S.M., and Farris, S.M. (2009). Ground plan of the insect mushroom body: Functional and evolutionary implications. The Journal of Comparative Neurology *513*, 265–291. https://doi.org/10.1002/cne.21948.

Strauss, R., Schuster, S., and Götz, K.G. (1997). Processing of artificial visual feedback in the walking fruit fly. Journal of Experimental Biology 1281–1296. .

Stürzl, W., and Zeil, J. (2007). Depth, contrast and view-based homing in outdoor scenes. Biological Cybernetics *96*, 519–531. https://doi.org/10.1007/s00422-007-0147-3.

Stürzl, W., Cheung, A., Cheng, K., and Zeil, J. (2008). The information content of panoramic images I: The rotational errors and the similarity of views in rectangular experimental arenas. Journal of Experimental Psychology: Animal Behavior Processes *34*, 1–14. https://doi.org/10.1037/0097-7403.34.1.1.

Stürzl, W., Grixa, I., Mair, E., Narendra, A., and Zeil, J. (2015). Three-dimensional models of natural environments and the mapping of navigational information. Journal of Comparative Physiology A *201*, 563–584. https://doi.org/10.1007/s00359-015-1002-y.

Stürzl, W., Zeil, J., Boeddeker, N., and Hemmi, J.M. (2016). How Wasps Acquire and Use Views for Homing. Current Biology *26*, 470–482. https://doi.org/10.1016/j.cub.2015.12.052.

Sun, X., Yue, S., and Mangan, M. (2020). A decentralised neural model explaining optimal integration of navigational strategies in insects. ELife *9:e54026*. https://doi.org/10.7554/eLife.54026.

Takalo, J., Piironen, A., Honkanen, A., Lempeä, M., Aikio, M., Tuukkanen, T., and Vähäsöyrinki, M. (2012). A fast and flexible panoramic virtual reality system for behavioural and electrophysiological experiments. Scientific Reports *2:324*. https://doi.org/10.1038/srep00324.

Takemura, S., Aso, Y., Hige, T., Wong, A., Lu, Z., Xu, C.S., Rivlin, P.K., Hess, H., Zhao, T., Parag, T., et al. (2017). A connectome of a learning and memory center in the adult *Drosophila* brain. ELife *6:e26975*. https://doi.org/10.7554/eLife.26975.

Taylor, G.J., Luu, T., Ball, D., and Srinivasan, M.V. (2013). Vision and air flow combine to streamline flying honeybees. Scientific Reports *3, 2614*. https://doi.org/10.1038/srep02614.

Touretzky, D.S., Redish, A.D., and Wan, H.S. (1993). Neural Representation of Space Using Sinusoidal Arrays. Neural Computation *5*, 869–884. https://doi.org/10.1162/neco.1993.5.6.869.

Turner-Evans, D.B., and Jayaraman, V. (2016). The insect central complex. Current Biology *26*, R453–R457. https://doi.org/10.1016/j.cub.2016.04.006.

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

Van De Poll, M.N., Zajaczkowski, E.L., Taylor, G.J., Srinivasan, M.V., and van Swinderen, B. (2015). Using an abstract geometry in virtual reality to explore choice behaviour: visual flicker preferences in honeybees. Journal of Experimental Biology *218*, 3448–3460. https://doi.org/10.1242/jeb.125138.

Varela, F.G. (1970a). The Optics of the Compound Eye of the Honeybee (*Apis mellifera*). The Journal of General Physiology *55*, 336–358. https://doi.org/10.1085/jgp.55.3.336.

Varela, F.G. (1970b). Fine Structure of the Visual System of the Honey Bee (*Apis mellifera*) II. The Lamina. Journal of Ultrastructure Research 178–194. .

Varga, A.G., Kathman, N.D., Martin, J.P., Guo, P., and Ritzmann, R.E. (2017). Spatial Navigation and the Central Complex: Sensory Acquisition, Orientation, and Motor Control. Frontiers in Behavioral Neuroscience *11:4*. https://doi.org/10.3389/fnbeh.2017.00004.

Vasas, V., Peng, F., MaBouDi, H., and Chittka, L. (2019). Randomly weighted receptor inputs can explain the large diversity of colour-coding neurons in the bee visual system. Scientific Reports *9:8330*. https://doi.org/10.1038/s41598-019-44375-0.

Via, S.E. (1977). Visually mediated snapping in the bulldog ant: A perceptual ambiguity between size and distance. Journal of Comparative Physiology - A *121*, 33–51. https://doi.org/10.1007/BF00614179.

Vickerstaff, R.J., and Cheung, A. (2010). Which coordinate system for modelling path integration? Journal of Theoretical Biology *263*, 242–261. https://doi.org/10.1016/j.jtbi.2009.11.021.

Viollet, S., and Zeil, J. (2013). Feed-forward and visual feedback control of head roll orientation in wasps (*Polistes humilis*, Vespidae, Hymenoptera). Journal of Experimental Biology *216*, 1280–1291. https://doi.org/10.1242/jeb.074773.

Vogt, K., Schnaitmann, C., Dylla, K.V., Knapek, S., Aso, Y., Rubin, G.M., and Tanimoto, H. (2014). Shared mushroom body circuits underlie visual and olfactory memories in *Drosophila*. ELife *3:e02395*, 1–22. https://doi.org/10.7554/eLife.02395.

Vogt, K., Aso, Y., Hige, T., Knapek, S., Ichinose, T., Friedrich, A.B., Turner, G.C., Rubin, G.M., and Tanimoto, H. (2016). Direct neural pathways convey distinct visual information to *Drosophila* mushroom bodies. ELife *5:e14009*, 1–13. https://doi.org/10.7554/eLife.14009.

Walløe, S., Pakkenberg, B., and Fabricius, K. (2014). Stereological estimation of total cell numbers in the human cerebral and cerebellar cortex. Frontiers in Human Neuroscience *8:508*, 1–9. https://doi.org/10.3389/fnhum.2014.00508.

Warrant, E.J. (2017). The remarkable visual capacities of nocturnal insects: vision at the limits with small eyes and tiny brains. Philosophical Transactions of the Royal Society B: Biological Sciences *372*, 20160063. https://doi.org/10.1098/rstb.2016.0063.

Warrant, E., and Dacke, M. (2016). Visual Navigation in Nocturnal Insects. Physiology *31*, 182–192. https://doi.org/10.1152/physiol.00046.2015.

Warrant, E., Frost, B., Green, K., Mouritsen, H., Dreyer, D., Adden, A., Brauburger, K., and Heinze, S. (2016). The Australian Bogong Moth *Agrotis infusa*: A Long-Distance Nocturnal Navigator. Frontiers in Behavioral Neuroscience *10:77*. https://doi.org/10.3389/fnbeh.2016.00077.

Warren, W.H. (2019). Non-Euclidean navigation. The Journal of Experimental Biology *222*, jeb187971. https://doi.org/10.1242/jeb.187971.

Webb, B. (2007). Insect Behaviour: Controlling Flight Altitude with Optic Flow. Current Biology *17*, R124–R125. https://doi.org/10.1016/j.cub.2006.12.008.

Webb, B. (2019). The internal maps of insects. The Journal of Experimental Biology *222*, jeb188094. https://doi.org/10.1242/jeb.188094.

Webb, B., and Wystrach, A. (2016). Neural mechanisms of insect navigation. Current Opinion in Insect Science *15*, 27–39. https://doi.org/10.1016/j.cois.2016.02.011.

Wehner, R. (2020). Desert navigator: The journey of an ant (Cambridge, Mass. USA: Harvard Univerity Press).

Wehner, R., and Menzel, R. (1990). Do Insects Have Cognitive Maps? Annual Review of Neuroscience *13*, 403–414. .

Wehner, R., and Muller, M. (2010). Piloting in desert ants: pinpointing the goal by discrete landmarks. Journal of Experimental Biology *213*, 4174–4179. https://doi.org/10.1242/jeb.050674.

Wehner, R., and Wehner, S. (1986). Path integration in desert ants: Approaching a long-standing puzzle in insect navigation. Monitore Zoologico Italiano *20*, 309–331. .

Wehner, R., Michel, B., and Antonsen, P. (1996). Visual Navigation in Insects: Coupling of Egocentric and Geocentric Information. Journal of Experimental Biology *199*, 129–140. .

Wehner, R., Meier, C., and Zollikofer, C. (2004). The ontogeny of foragwehaviour in desert ants, *Cataglyphis bicolor*. Ecological Entomology *29*, 240–250. https://doi.org/10.1111/j.0307-6946.2004.00591.x.

Wehner, R., Boyer, M., Loertscher, F., Sommer, S., and Menzi, U. (2006). Ant Navigation: One-Way Routes Rather Than Maps. Current Biology *16*, 75–79. https://doi.org/10.1016/j.cub.2005.11.035.

Wehner, R., Cheng, K., and Cruse, H. (2014). Visual navigation strategies in insects: lessons from desert ants. In The New Visual Neurosciences, (MIT Press), pp. 1153–1163.

Wehner, R., Hoinville, T., Cruse, H., and Cheng, K. (2016). Steering intermediate courses: desert ants combine information from various navigational routines. Journal of Comparative Physiology A *202*, 459–472. https://doi.org/10.1007/s00359-016-1094-z.

Wenner, A.M., Wells, P.H., and Rohlf, F.J. (1967). An Analysis of the Waggle Dance and Recruitment in Honey Bees. Physiological Zoology *40*, 317–344. https://doi.org/10.1086/physzool.40.4.30158452.

Williams, R.W., and Herrup, K. (1988). The Control of Neuron Number. Annual Review of Neuroscience 423–453. .

Wittlinger, M., Wehner, R., and Wolf, H. (2006). The Ant Odometer: Stepping on Stilts and Stumps. Science *312*, 1965–1967. https://doi.org/10.1126/science.1126912.

Wittmann, T., and Schwegler, H. (1995a). Path integration—a network model. Biological Cybernetics *73*, 569–575. .

Wittmann, T., and Schwegler, H. (1995b). Path integration - a network model. Biological Cybernetics *73*, 569–575. .

Wystrach, A., and Graham, P. (2012). What can we learn from studies of insect navigation? Animal Behaviour *84*, 13–20. https://doi.org/10.1016/j.anbehav.2012.04.017.

Wystrach, A., Beugnon, G., and Cheng, K. (2011a). Ants might use different view-matching strategies on and off the route. Journal of Experimental Biology *215*, 44–55. https://doi.org/10.1242/jeb.059584.

Wystrach, A., Schwarz, S., Schultheiss, P., Beugnon, G., and Cheng, K. (2011b). Views, landmarks, and routes: how do desert ants negotiate an obstacle course? Journal of Comparative Physiology A *197*, 167–179. https://doi.org/10.1007/s00359-010-0597-2.

Wystrach, A., Beugnon, G., and Cheng, K. (2011c). Landmarks or panoramas: what do navigating ants attend to for guidance? Frontiers in Zoology *8*, 21. https://doi.org/10.1186/1742-9994-8-21.

Wystrach, A., Mangan, M., Philippides, A., and Graham, P. (2013). Snapshots in ants? New interpretations of paradigmatic experiments. Journal of Experimental Biology *216*, 1766–1770. https://doi.org/10.1242/jeb.082941.

Wystrach, A., Philippides, A., Aurejac, A., Cheng, K., and Graham, P. (2014). Visual scanning behaviours and their role in the navigation of the Australian desert ant *Melophorus bagoti*. Journal of Comparative Physiology A *200*, 615–626. https://doi.org/10.1007/s00359-014-0900-8.

Wystrach, A., Mangan, M., and Webb, B. (2015). Optimal cue integration in ants. Proceedings of the Royal Society B: Biological Sciences *282*, 20151484. https://doi.org/10.1098/rspb.2015.1484.

Wystrach, A., Dewar, A., Philippides, A., and Graham, P. (2016a). How do field of view and resolution affect the information content of panoramic scenes for visual navigation? A computational investigation. Journal of Comparative Physiology A *202*, 87–95. https://doi.org/10.1007/s00359-015-1052-1.

Wystrach, A., Lagogiannis, K., and Webb, B. (2016b). Continuous lateral oscillations as a core mechanism for taxis in Drosophila larvae. ELife *5:e15504*. https://doi.org/10.7554/eLife.15504.

Wystrach, A., Schwarz, S., Graham, P., and Cheng, K. (2019). Running paths to nowhere: repetition of routes shows how navigating ants modulate online the weights accorded to cues. Animal Cognition *22*, 213–222. https://doi.org/10.1007/s10071-019-01236-7.

Wystrach, A., Buehlmann, C., Schwarz, S., Cheng, K., and Graham, P. (2020). Avoiding pitfalls: Trace conditioning and rapid aversive learning during route navigation in desert ants. Current Biology *30*, 1927–1933. https://doi.org/10.1101/771204.

Yukizane, M., Kaneko, A., and Tomioka, K. (2002). Electrophysiological and morphological characterization of the medulla bilateral neurons that connect bilateral optic lobes in the cricket, *Gryllus bimaculatus*. Journal of Insect Physiology *48*, 631–641. https://doi.org/10.1016/S0022-1910(02)00091-4.

Zahedi, S.M., and Zeil, J. (2018). Fractal dimension and the navigational information provided by natural scenes. PLOS ONE *13(5)*, e0196227. https://doi.org/10.1371/journal.pone.0196227.

Zeil, J. (2012). Visual homing: an insect perspective. Current Opinion in Neurobiology *22*, 285–293. https://doi.org/10.1016/j.conb.2011.12.008.

Zeil, J., and Fleischmann, P.N. (2019). The learning walks of ants (Hymenoptera: Formicidae). Myrmecological News *29*, 93–110. .

Zeil, J., Hofmann, M.I., and Chahl, J.S. (2003). Catchment areas of panoramic snapshots in outdoor scenes. Journal of the Optical Society of America A *20*, 450–469. .

Zeil, J., Ribi, W.A., and Narendra, A. (2014a). Polarisation Vision in Ants, Bees and Wasps. In Polarized Light and Polarization Vision in Animal Sciences, G. Horváth, ed. (Berlin, Heidelberg: Springer Berlin Heidelberg), pp. 41–60.

Zeil, J., Narendra, A., and Stürzl, W. (2014b). Looking and homing: how displaced ants decide where to go. Philosophical Transactions of the Royal Society B: Biological Sciences *369*, 20130034. https://doi.org/10.1098/rstb.2013.0034.

Zollikofer, C., Wehner, R., and Fukoshi, T. (1995). Optical scaling in conspecific *Cataglyphis* ants. Journal of Experimental Biology *198*, 1637–1646. .