



DEPARTMENT OF INFORMATICS

TECHNISCHE UNIVERSITÄT MÜNCHEN

Master's Thesis in Informatics: Robotics, Cognition and Intelligence

**Biomimetic Visual Navigation: Understanding
the Visual Ecology of Expert Navigators
through Virtual Reality**

Jose Adrian Vega Vermehren





DEPARTMENT OF INFORMATICS

TECHNISCHE UNIVERSITÄT MÜNCHEN

Master's Thesis in Informatics: Robotics, Cognition and Intelligence

**Biomimetic Visual Navigation: Understanding
the Visual Ecology of Expert Navigators
through Virtual Reality**

**Bionische Visuelle Navigation:
Untersuchungen der Visuellen Ökologie von
Navigationsexperten durch Virtuelle Realität**

Author: Jose Adrian Vega Vermehren
Supervisor: Prof. Dr. Paul Graham | Prof. Dr. Daniel Cremers
Advisor: Dr. Cornelia Bühlmann | Nikolaus Demmel
Submission Date: 10th of August 2021



I confirm that this master's thesis in informatics: robotics, cognition and intelligence is my own work and I have documented all sources and material used.

A handwritten signature in black ink, consisting of several fluid, overlapping loops and strokes, positioned above the printed name.

Munich, 10th of August 2021

Jose Adrian Vega Vermehren

Acknowledgments

First and foremost, I would like to thank Cornelia and Paul for taking me again under their wing, and letting me be part of their exciting research. I can't wait to be back in Brighton to do some more awesome science together.

A big thanks to Roman, Jeremy and Barbara for the additional support and feedback.

This project would not have been possible without the support of Nikolaus. Niko, thanks for believing in me and my project when nobody else did.

Organizing an international research project in the middle of a pandemic is no easy task. I am very grateful to all the people who stood beside me in the process:

My parents and brother, who are always there for me.

Kai, Niki, Pablo and Felix who got me through the weirdest of times and back out into the world. Special thanks to Felix, whose emotional support proved key in planning this project.

Charlie, Bernie and Aaron who let me believe I could achieve my goals.

Dan, Steve, Blaithin, Blake, Kostas and Matthias who are always there when I need an alternate reality.

Munashe and Justine, without whom Brighton wouldn't have been the same.

Finally, I am very thankful for the financial support provided by the DAAD.

Abstract

Wood ants thrive as expert navigators in the same cluttered, dynamic, light-variant environments where most robot navigation algorithms fail [1, 2]. Virtual Reality (VR) is a novel technique to study the visual ecology and navigation strategies of these insects [3, 4]. We evaluate the functionality of a novel treadmill and VR system that allows experiments with untethered walking insects [5, 6]. In a series of four experiments with incrementally complex visual stimuli, we validate the setup and gather insights into the ants' behaviour. The system works remarkably well to study ant navigation, despite a few limitations. On the behavioural side, ants use oscillations as a visuomotor control mechanisms even when deprived of sensory feedback from the environment (rot. close-loop). Furthermore, ants exploit edges as information-rich features in both simple and complex visual scenarios. Both of these behaviours have important implications for the design of robotic navigation architectures. VR opens a new opportunity to learn from the simple yet effective navigation strategies of ants and design new biomimetic navigation systems.

Kurzfassung

Waldameisen gedeihen als erfahrene Navigatoren unter Bedingungen, in denen die meisten Roboternavigationsalgorithmen versagen: überladene, dynamische Umgebungen mit ständig wechselnden Lichtverhältnissen [1, 2]. Virtuelle Realität (VR) bietet eine neue Methodologie um die visuelle Ökologie und Navigation dieser Insekten zu untersuchen [3, 4]. Wir bewerten die Funktionalität eines neuen Systems, ein an ein VR System gekoppeltes Laufband, mit dem laufende Insekten untersucht werden können, ohne sie zu fixieren [5, 6]. In einer Reihe von vier Experimenten mit visuellen Reizen zunehmender Komplexität validieren wir den Versuchsaufbau und sammeln neue Erkenntnisse zum Verhalten der Ameisen. Zur Untersuchung der Navigation von Ameisen eignet sich das System trotz mancher Limitierungen bemerkenswert gut. Wir haben herausgefunden, dass Ameisen trotz mangelnder visueller Rückkopplung (rot. close-loop) ein Oszillationsmechanismus zur visuomotorischen Regelung anwenden. Außerdem nutzen Ameisen die potenziell informationsreichen Eigenschaften von Kanten in einfachen und komplexen Szenen aus. Beide beobachteten Verhaltensweisen können wegweisend sein für das Design von neuartigen Navigationsarchitekturen für Roboter. VR eröffnet neue Möglichkeiten von den einfachen, aber eleganten Strategien von Ameisen zu lernen, um neue bionische Navigationssysteme zu entwickeln.

Contents

Acknowledgments	iii
Abstract	iv
Kurzfassung	v
1. Introduction	1
1.1. Research Goals and Project Outline	2
2. Literature Review	4
2.1. Robot Visual Navigation Using Maps	5
2.1.1. Map-Based Systems	5
2.1.2. Map-Building Systems	6
2.2. Mapless Systems	8
2.2.1. Optic Flow	8
2.2.2. Feature-Based Tracking	9
2.2.3. Appearance-Based Matching	10
2.3. Summary: Technical and Biomimetic Navigation Algorithms	11
3. Methods	14
3.1. Integration of the Treadmill and VR System	14
3.1.1. Motion Compensating Treadmill	14
3.1.2. Virtual Reality System	16
3.1.3. Systems Integration	16
3.2. Experiments	17
3.2.1. Wood Ants	17
3.2.2. Experimental Procedures	18
3.3. Data Analysis	20
3.3.1. Data Preprocessing	20
3.3.2. Chunks and Activity Levels	21
3.3.3. Individual Analysis	21
3.3.4. Statistics	24

4. Results	25
4.1. Treadmill and VR Integration in Open- and Close-Loop	25
4.1.1. Using Ant Activity as a Metric to Filter Behaviour	26
4.1.2. Landmark Attraction in the Open-Loop	29
4.1.3. Ant Behaviour in the Close-Loop	32
4.2. Validating the Display of Natural Images on the VR Setting	36
4.2.1. Display of Complex Artificial Patterns	37
4.2.2. Display of Complex Natural Images	40
5. Discussion	45
5.1. Evaluation of the Virtual Reality Experimental Setting	46
5.1.1. Hardware Limitations	46
5.1.2. Software Limitations	48
5.2. Wood Ant Behaviour While on the Experimental Setting	49
5.2.1. Open- and Close-Loop Oscillations	50
5.2.2. Edge Fixation	51
5.3. Conclusion	53
A. Appendix	54
A.1. Data and Scripts	54
List of Figures	55
List of Tables	59
Bibliography	60

1. Introduction

From single celled organisms to mammals with complex behaviour, visual systems have played an important role in evolution. The analogous development of light sensitive organs has fascinated scientists since Darwin first posed his theory of evolution [7]. And not surprisingly, behaviour is tightly intertwined: visually guided behaviour acts both as the driver as much as the consequence of these evolving systems. Take the notorious "eagle eye" as an example where the predator's hunting behaviour and distinct sharp vision (a vertebrate lens eye) evolved in synchrony [8]. Or, on an independent line of evolution, compound eyes in insects have evolved along with remarkable control and navigation behaviours [9].

Human behaviour has also evolved strongly around vision. In fact, there is significantly more research in vision than in any another sensory modality, up to the point that more general concepts like "perception" or "perceptual memory" are often used synonymously to "visual perception" and "visual memory" [10]. By no coincidence, human design of artificial intelligence and autonomous agents reflects this tendency for vision over other modalities.

One key role for vision in natural and artificial systems is to provide information to navigation systems. Roughly defined as finding a suitable path between the current location and a goal, navigation is a necessary precursor for more complex behaviours [11, 12]. Robot navigation strategies based on vision have received significant attention in the last three decades due to the large application scope and promising results in reaching true autonomy [11]. While contemporary visual navigation algorithms have succeeded in niche applications, navigating the complex, undetermined, and chaotic real world remains a challenge to be solved [13].

Natural agents have clearly found solutions to the problems of navigation through complex environments. Quite notoriously, insects have evolved vision and navigation systems strongly constrained by their tiny brains and low visual sensory resolution. Insects are capable of performing simple yet elegant computations to achieve their navigation tasks, e.g. foragers of many social insect species can travel vast distances and find their way back home in the most complex of environments [9, 14].

Ants belong to this group of expert navigators. Although ants have access to a variety of orientation mechanisms [15], including some forms of social cues, solitary foragers rely foremost on vision to navigate [1, 16]. Foraging ants learn the necessary visual information to guide long and complex routes between their nest and a stable food site [17, 18]. And when compared to artificial systems, ants have taken a different approach to solve the problem of navigation. Contrary to map based technical implementations like visual Simultaneous

Localization and Mapping (vSLAM), ant navigation is thought to be of procedural nature, whereby ants use visual cues to trigger appropriate behaviours [1]. Furthermore, evidence shows ant navigation to rely on the overall appearance of scenes rather than segmenting individual visual features [19].

Although efforts to develop bio-inspired robots that mimic the ants remarkable navigation strategies have come far, plenty of questions regarding the visual ecology and spatial memory of the model foragers still remain unanswered [1, 20]. Research on ant navigation has historically been limited to two approaches: field work and lab experiments. Research using either needs to compromise between portraying foragers in their natural environments and control over the stimuli influencing the research subjects. However, developments in computer vision and computer graphics have created the opportunity to adapt a new technique of experimentation: virtual reality (VR). Using VR to study the visual ecology of navigating ants promises the best of both methodologies: absolute control over the stimuli while simulating the ant's natural scenery.

As a study subject, we choose wood ants. These species of ants thrives as a navigator in exactly the kind of cluttered, dynamic and uncontrolled environment in which most robots fail to navigate. Understanding the underlying behavioural and neuronal mechanisms involved in their navigation could lead to a breakthrough in robot navigation.

Following this logic, this project aims to shed some light into virtual reality as a method to study ant navigation.

1.1. Research Goals and Project Outline

The use of virtual reality as a methodology to study insects is novel, however not unheard off. In general, most of the research in insect behaviour focuses on the common fruit fly. VR has been extensively validated for the study of these organisms [21]. However, most of the validated VR methodologies, including those developed for social insects [4, 3], fixate the study specimens to a tether. This methodology significantly constraints natural behaviour. A less invasive method requires the parallel development of a multilateral treadmill system where walking insects can move "naturally" in response to the VR stimuli without changing location [5]. Goulard et al. [6] recently developed such a treadmill for ant study. Under their supervision, this project contributes to the development and validation of this system. The following research objectives are pursued:

1. Build upon the development of a trackball and VR system for untethered ants.
2. Authenticate the system as a research methodology to study ant visual navigation.
3. Evaluate the influence of a "close-loop" setting in the ant's navigation behaviour.
4. Describe ant behaviour when confronted with natural images in the VR setting.

To put the project into context, I begin by reviewing relevant literature (Chapter 2). The answer to two questions summarize the outcome of this literature review: what are shortcomings of state-of-the-art visual navigation algorithms like vSLAM, and what do previously proposed ant navigation algorithms look like?

The outcomes of the first research goal are mainly described in the methods (chapter 3). Here I detail upon the integration of the treadmill and VR systems as well as the implementation of different experimental settings including a closed-loop rotation system.

The results in chapter 4 are divided into two sections: (i) validation of the system to investigate ant navigation and (ii) ant navigation in complex VR scenes. Research goals 2 and 3 are explored in the first part. Using two sets of experiments with a simple visual cue, I describe and compare ant navigation on the open- and close-loop system. Research goal number 4 is detailed in the second part; here, I describe ant navigation on the novel system with a complex artificial and a complex natural scene.

Finally, a discussion over the results is offered in chapter 5. Here I reflect upon the four research goals and summarize the validity of the novel system to study ant navigation. Furthermore, I describe follow-up experiments that could eventually lead to a new robot navigation algorithm inspired by ants.

2. Literature Review

Both natural and artificial agents rely on spatial memory to navigate, i.e. agents need to remember information about their surroundings in order to be able to return to that specific location. Navigation using vision is the primary modality for both expert animal navigators and state-of-the-art robot navigation algorithms [11, 9]. That being said, visual spatial navigation varies in complexity. We distinguish four levels of cognitive complexity [22]. In increasing order:

1. recognition of a location upon re-encounter
2. visual servoing towards a clear visual cue at the destination
3. visual homing, i.e. comparing a stored "home" view to guide the agent back to its "base"
4. linking visual information to a "map-like" representation of space

Even the simplest of these visual recognition tasks, challenges agents to deal with all sort of visual problems [22, 13]. Breakthroughs in computer science have advanced robot navigation significantly. Even so, navigation in certain locations, e.g. cluttered, dynamic, outdoor environments, remain unconquered. And yet, ants and other social insects thrive as navigators under these conditions; evolution seems to have found clever solutions to many of these visual problems [9].

In order to explore new strategies for visual navigation, it is important to understand the limitation of previous approaches. This chapter systematically explores the most important techniques in visual navigation. Both technical algorithms and biologically inspired models are categorized by their strategy. The main drawbacks of each approach are outlined. Figure 2.1 offers an overview over the categorization.

In literature, we find a distinction between indoor and outdoor navigation [23, 11]. The latter is then divided in structured (e.g. roads) and unstructured environments. However, if we ignore methods involving an external reference frame (e.g. GPS), outdoor navigation shares similar constraints with its indoor counterpart. This thesis focuses on this set of application scenarios and thus treats outdoor navigation as a complex, less controlled version of indoor environments.

Based on the available information and strategy, visual (indoor) navigation algorithms can be roughly divided into three categories: mapless systems [24], map-based systems and map-building systems (Fig. 2.1; [23, 11, 25]). While mapless systems vary in cognitive complexity

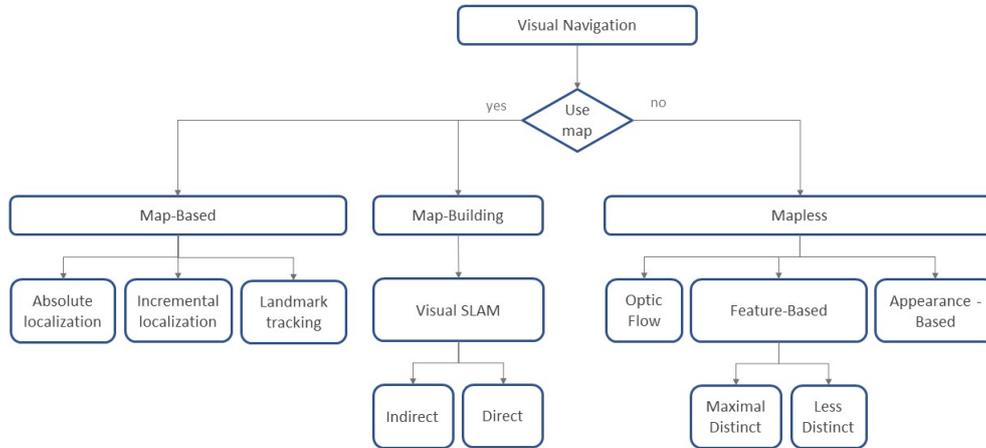


Figure 2.1.: Overview vision-based navigation methodologies. Adapted from [25].

(levels 1-3), the latter two use maps, and thus fall into the highest level of cognitive complexity (level 4).

2.1. Robot Visual Navigation Using Maps

Increasing developments in computer vision have made vision based navigation specially effective to pursue robot autonomy. In the last three decades, countless research efforts have been made to improve their navigation algorithms. Most approaches are based on the use of maps. Depending on the task and the previously available information, agents can employ map-based or map-building systems [23, 25].

2.1.1. Map-Based Systems

Map-based systems require a predefined spatial layout of the environment. The data structure and detail varies from a full CAD model to a simple graph, depending on the system [24]. At its core, the map is used to generate a sequence of expected landmarks while the vision system attempts to locate them. A successful match allows the robot to self-estimate its position in the map relative to the recognized landmark. Hence, navigation using this method is a sequence of four steps: (i) image acquisition, (ii) landmark detection, (iii) matching of expectation, and observation and finally (iv) position estimation.

The complexity of these systems lies mainly in the third step, matching (correspondence problem). Different approaches to solve it have been proposed [24, 23]:

Absolute Methods The system does not know the starting position of the robot. Hence, an exact match between features is required to locate the agent within the map (correspondence problem).

Incremental Methods These algorithms keep track of the localization uncertainty starting from a known position. Hence, an exact match is only needed to recalibrate the error propagation. The matching problem is reduced to a probabilistic analysis.

Landmark Tracking The general idea is to trace known landmarks matched at a known starting position across subsequent frames. Here, landmarks can be either natural or artificial and need to be defined by a human operator.

Map-based systems are quite successful to solve tasks in known and well-defined environments. However, they are strongly affected by certain limitations [24, 23]:

- **Prior knowledge** of the environment in the form of a map is needed. Without an additional obstacle avoidance systems, they are unable to navigate dynamic environments.
- Ambiguities in matching increase the complexity and reduce the systems' **robustness**. Map-based systems are unsuitable to navigate cluttered, chaotic environments.

2.1.2. Map-Building Systems

Map-building systems have received the highest attention in recent years. Compensating one of the main drawbacks of map-based systems, map-building systems navigate the environment, while building a representation of it. Although other methods exist, by far the leading strategy, is Simultaneous Localization and Mapping (SLAM). These systems are able to navigate unknown environments by performing three activities in parallel, navigation, mapping and localization [24, 23, 2].

Vision based SLAM (vSLAM) algorithms have been greatly developed in the last thirty years, and are currently considered the state-of-the-art approach for visual robot navigation. Hence, a more detailed analysis of the strategies employed, and their shortcomings, is in order:

vSLAM Methods

With the rapid development in camera technology, camera sensors have become cheaper, they consume less power and are able to provide robust and highly detailed real time information of the environment. vSLAM algorithms rely solely on cameras as their navigation sensor; significantly increasing the range and flexibility, not to say affordability, of their application domain. Nevertheless, compared to other approaches that use different sensors (e.g. Lidar), vSLAM algorithms come with a higher technical difficulty due to the limited field of view [2,

26].

vSLAM works in four modules. In the first module, the sensor data is retrieved and pre-processed. Next, the Front-End module uses the image motion to generate a first estimate of position. In parallel, a system called Loop Closure uses the preprocessed sensor data to calculate the similarity between the current image and a stored map representation. The third module, Back-End, constraints the initial estimation from the Front End with the information from the Loop Closure. Finally, the Back-End solves an optimization problem between the last module, Map, and its current global estimate of the agent's position [2].

Based on the method employed by the Front-End, vSLAM algorithms are divided in two categories: indirect and direct methods [2, 25]:

Indirect methods Indirect Methods are based on the assumption of geometric consistency. Instead of using entire images, algorithms in this category use geometric features to estimate the image motion. A feature should be invariant to rotation, viewpoint and scale changes, as well as robust against noise, blur and illumination. These feature points need to be extracted, matched and estimated. The first step, feature extraction, is a computational expensive operation, hence the success of the algorithm relies heavily on the speed and quality of feature extraction (feature extraction problem). SIFT is the most used extraction method (see section 2.2.2 for more detail). Like with absolute map-based systems, matching is also an important issue, heavily constrained by mismatches and moving targets (correspondence problem).

Furthermore, the density of distinct points that the algorithm can reconstruct has a strong influence on the application domain; indirect methods do not perform well in cluttered texture-less environments [2, 25].

Direct methods Direct methods are based on the greyscale-value invariance assumption: a pixel has the same greyscale value in subsequent images. Hence, the direct method does not extract individual features, but uses all the information in the image to estimate image motion. Algorithms under this category are more robust to geometric changes. They are in general faster and are better at dealing with dense maps. The achievable density of these maps is proportional to the available computational power. Nevertheless, these algorithms are sensitive to direct sunlight and shadows, illumination variance and specular reflections [2, 25].

Hybrid methods Hybrid methods attempt to increase efficiency by exploiting each of the previous methods at its strongest. They first employ indirect methods to initialize feature correspondence and then turn towards direct methods to refine the camera poses [25].

Challenges of vSLAM

Although vSLAM algorithms have come a long way, plenty of challenges remain to be solved. Mainly, how to robustly navigate more complex environments while meeting the requirements of real-time agents [27, 2]:

- vSLAM algorithms are still not very good at dealing with **dynamic environments**. Moving elements in a robot's surrounding significantly complicate feature matching and generate unwilling changes in illumination.
- **Texture richness** is important. Indirect methods perform poorly in texture-less environments. Furthermore, regular patterns can cause mismatches and missing features.
- **Variance in illumination** violates the underlying assumption of direct methods. It has yet to be solved for dynamic scenarios and refined for static ones. Intense sunlight and shadows, as well as reflections common in uncontrolled outdoor spaces, can cause the system to fail.
- The **complex computations** required to match camera poses and find a global optimum, constraints to the size and speed of agents running vSLAM algorithms.
- **Scarcity** influences the precision during navigation and the quality of the arising maps. There is a trade-off between density and computational power.
- Mismatches by the Loop Closure detection have the potential to limit the system's **robustness** and fail navigation in both dynamic and static environments.

2.2. Mapless Systems

As established before, keeping a "map-like" representation of space requires the highest level of cognition (level 4), i.e. relying on maps (either prebuilt or built during navigation) is an expensive strategy with a strong influence on performance. And yet, evidence suggests that social insects are able to perform complex navigational tasks without the use of cognitive maps [28]. Here, I shall hence focus on the next most complex cognitive ability, homing (level 3).

Systems in this category are relatively new compared to the previously described ones. This section reviews the most important technical and biomimetic approaches to mapless navigation. Special emphasis is made on insect-navigation models and algorithms.

2.2.1. Optic Flow

Optic flow (OF) is a strategy employed both by humans and animals. Agents use the motion of the surface elements in the environment to calculate a moving direction and the distance

to obstacles [24]. Speed and accuracy can be regulated according to the task at hand. Bees, expert navigators, are known for their mastery of this strategy. These insects are able to regulate flight direction, speed and height as well as avoid obstacles and calculate odometry with little more than optic flow calculations [9].

Already in 1993 Santos-Victor et al. [29] proposed an optic flow algorithm inspired by the bee's flight strategies. Here, localization is achieved by comparing the image velocity of both eyes (i.e. cameras). If both sides move at the same speed the agent keeps moving forwards, however if there is a difference, the robot steers in the direction of lower speed.

Robots implementing optic flow as a navigation strategy face two important challenges: (1) OF cannot disentangle distance from velocity, and (2) OF is very small and thus less descriptive in the direction of flight [30]. De Croon, De Wagter and Seidl [30] recently proposed a method that treats optic flow as a learned and not innate feature. Their robot achieves smoother landings, better obstacle avoidance and higher speeds by implementing a previous learning process.

Homing based on Optic Flow OF is a powerful tool for steering and object avoidance. It is however often underestimated as a method of spatial memory. Vardy and Möller [31] proposed in 2005 a series of techniques based on insect homing that allow corresponding two images by means of optic flow. They show multiple methods, including block matching and differential OF as plausible models of efficient insect homing. The simplicity and robustness of these methods and their propensity for low-frequency features are ideal for lightweight robot navigation.

2.2.2. Feature-Based Tracking

Equivalent to its map-based and map-building counterparts, mapless feature-based methods also track the relative changes of previously extracted features across subsequent images. Since in this case no map is involved, a learning step, where the agent remembers its surrounding in the form of *snapshots* at the home location, is required [24].

A feature is a landmark that can be clearly segmented from the image's background. First, the agent must extract the same feature in both the current image and snapshot (feature extraction problem). Then, each feature has to be matched in the current and remembered image (correspondence problem) [24]. Extraction approaches vary on the level of landmark uniqueness, while some methods strive for maximal distinctive features, others use less unique features [31]:

Maximal Distinct Features

Algorithms in this category strive to extract maximally distinct features. The ideal feature is unique to the point, that it can be corresponded with 100% accuracy, i.e. there is no other

feature that could look like that. The computer vision algorithm SIFT stands out for its popularity and performance:

Scale Invariant Feature Transform A milestone in feature extraction was the invention of the scale invariant feature transform (SIFT) algorithm [32]. SIFT is nowadays a standard method in landmark detection, both for map-building and mapless technical systems alike. In a series of image operations, SIFT extracts features invariant to scaling, rotation and illumination, significantly increasing robustness to matching upon re-encounter.

In order to establish one-to-one correspondences, distinct feature methods have to search for the landmark in the whole image. Since the entire image is searched through anyway, no preprocessing steps like image aligning are needed. That being said, distinct features are hard to extract in cluttered environments. Furthermore, the extraction and matching process are both computationally expensive and a trade-off between computation speed and onboard load has to be made for autonomous agents [24].

Less Distinct Features

Less distinct features are significantly easier to extract, matching, however, is often ambiguous. One of the first models of insect navigation falls into this category:

The Snapshot Model Most insect navigation research and many robot navigation algorithms relate to the snapshot model. Based on landmark navigation experiments with bees, Cartwright and Collet [33] described a homing method capable of deriving a heading direction from the discrepancies between a single snapshot and the current image. Less unique, dark and bright features are used to navigate. The difference between snapshot and current image is used to generate a relative movement vector. By iteratively lessening the mismatch between features, the agent is able to navigate back home.

When using non-unique features, many correspondences might exist. To solve this problem with reliability, the agent needs to align the current view and the snapshot view to a common coordinate system. The implications are both unpractical and biologically implausible, which is why follow-up research has focused on methods of alignment matching [31].

2.2.3. Appearance-Based Matching

In its core, appearance based methods store representations of the environment and associate them with the correct command to steer the agent to its goal [24]. It is a two-step approach. First, the agent has to learn prominent information of its surrounding and attach the appropriate steering information to it. Afterwards, during navigation, the agent has access to the correct steering command upon re-encounter with a previously learned template. Different

methods employ different sources of information to describe unique locations. Two strategies derived from ant navigation are worth highlighting:

Image Warping The warping method, initially proposed by Franz et al. [34], has proven its value by multiple robotic implementation as a robust homing strategy. The general idea is to compute the set of all possible positions and rotations between the current location and the goal. The current image is distorted (warped) to approximate the view the robot would have, had it moved according to the corresponding parameters. The warped image can then be compared with the snapshot by a similarity measure. Homing is achieved by an exhaustive search for the parameters that generate the most similar warping.

The robustness of this method makes it a very reliable navigation strategy. Nevertheless, the cognitive complexity grows exponentially with the space of possible parameters, making it unsuitable for lightweight navigation in complex environments.

Image Difference Function The homing algorithm presented by Zeil et al. [35] is surprisingly simple yet effective. The home location is saved as an omnidirectional image. The Image Difference Function (IDF) calculates the pixel-wise similarity (root-mean-square) between the current image and snapshot. In natural images, the difference increases monotonically with distance from the home location, so gradient descent can be used to determine a heading direction. To address the image-align problem, compass information can be abstracted from an equivalent gradient in the rotation Image Difference Function (rIDF).

Although successful in both indoor and outdoor experiments, this strategy is limited by a certain distance from the home location, within which, a global minimum of the function exists (often called catchment area) [36]. Further research has shown, that the catchment area can be significantly expanded by stitching views together and forming routes [37, 38, 39].

2.3. Summary: Technical and Biomimetic Navigation Algorithms

Navigation is a complex task that agents need to master to achieve true autonomy. This chapter reviews some of the most influential approaches to visual navigation proposed in the last 30 years. Table 2.1 shows an overview of the different methodologies and their main drawbacks.

Certain ideas and strategies stretch across the presented categories. A comparison of the limitation of each methodology leads to certain conclusions:

- Navigation with a map is an anthropomorphic notion. Social insects are an example of how complex navigational tasks can be achieved without them. Prebuild maps are just not practical, and building a map during navigation is an expensive luxury. The procedural nature of mapless strategies is simple yet effective.

Table 2.1.: Overview of the reviewed algorithms.

Method	Category	Strategy	Main Drawback(s)
Absolute	Map-based	Exact feature match	Map needed, robustness
Incremental	Map-based	Uncertainty propagation	Map needed
Landmark Tracking	Map based	Known landmark match	Landmark design
Indirect	vSLAM	Geometric consistency	Texture, map scarcity
Direct	vSLAM	Greyscale-value invariance	Lightning variability
Hybrid	vSLAM	Direct and Indirect vSLAM	Comp. complexity
Correspondence OF	Mapless	Optic Flow	Robustness
SIFT	Feature-Based	Maximal distinct features	Comp. complexity
Snapshot Model	Feature-Based	Indistinct Features	View alignment
Image Warping	Appearance-Based	Image Distortion	Comp. complexity
IDF	Appearance-Based	Pixel-Wise Similarity	Catchment area

- Across categories, methodologies relying on geometric features (indirect vSLAM, feature-based mapless) are faced with the feature extraction problem. Robust navigation requires these features to be scale, rotation and illumination invariant. Feature extraction algorithms like SIFT are computational expensive and perform poorly in cluttered, texture-poor environments. Systems that employ less distinct features (Snapshot Model) perform better under these conditions, but are unpractical due to the image-alignment prerequisite.
- Using the information encoded in the entire image instead of single features leads to more robust algorithms (direct vSLAM, appearance-based mapless). A good metric for image similarity are pixel-wise comparisons. The image attribute used to describe a location varies across methods. One of them, photometric comparison, is susceptible to extreme lighting conditions. Optic flow is another source of information that can be exploited to describe a location. The IDF strategy shows promising results, and it has been shown that bigger catchment areas can be achieved by stitching multiple views together.

Mapless, appearance-based strategies are both biologically plausible models of insect navigation and efficient solutions for robot navigation. Current robotic implementations of these techniques still struggle under certain environmental conditions: cluttered, dynamic, light-variant sceneries. Wood ants have adapted to navigate under precisely these conditions.

Yet, we know little about the visual ecology of these expert navigators. In other words, we have a good model on **HOW** ants (and robots) can use visual information to navigate, but still need to understand **WHICH** information in their natural panorama plays a decisive role. New research methodologies are needed to explore this question.

3. Methods

Our VR/treadmill system provides a novel technique to investigate the visual ecology of wood ants. This chapter deals with the methodologies involved, and is divided in 3 sections. First, I describe the integration between the treadmill and VR systems. I report the functionality and limitations of the experimental setup. Next, I describe the experiments we performed. We collected four datasets, divided in two groups: (i) with a simple visual cue, and (ii) with a complex scene. These groups correspond to the two main sections of the results (Chapter 4). Finally, I describe the tools and methods used to analyse the data on each of the four datasets.

3.1. Integration of the Treadmill and VR System

The system is composed of two subsystems, the motion compensating treadmill and the VR setting. In a previous publication, Goulard et al. [6] documented the treadmill subsystem extensively, however at that stage of the systems' development, the VR system had not yet been developed. In this section the functionality of the experimental setup and the new integration between both subsystems is described. Figure 3.1 offers an overview of the complete experimental setting.

3.1.1. Motion Compensating Treadmill

The motion compensating treadmill works through an information loop (3.1A blue, [6]). The ants are positioned on top of a white foam sphere with a diameter of 120 mm. Aligned to the centre of the sphere, a high speed zenithal camera (Basler ace acA640-750 μm – monochrome) tracks the movements of the ant. The camera is positioned 150 mm above the sphere and is equipped with a 12 mm focal lens objective. External lighting is provided by four lights, each consistent of three white LEDs and a diffusive cover. The lights are supported by a white cardboard plane that blocks external visual stimuli. The camera communicates with a Raspberry Pi (4 Model B) microcontroller (running Ubuntu 17), which serves as the computation unit of the treadmill system. Here, the motion of the ant is computed into a compensating motor signal and forwarded to the motor controller via USB. An Arduino Uno microcontroller connected to 3 motor drivers (STMicroelectronics, ULN-2064B) controls the rotation of the rotors. The three stepper motors (SANYO SY42STH38 - 0406A) are equally distributed around the sphere and tilted by 60 deg. Each one is equipped with a dual disc

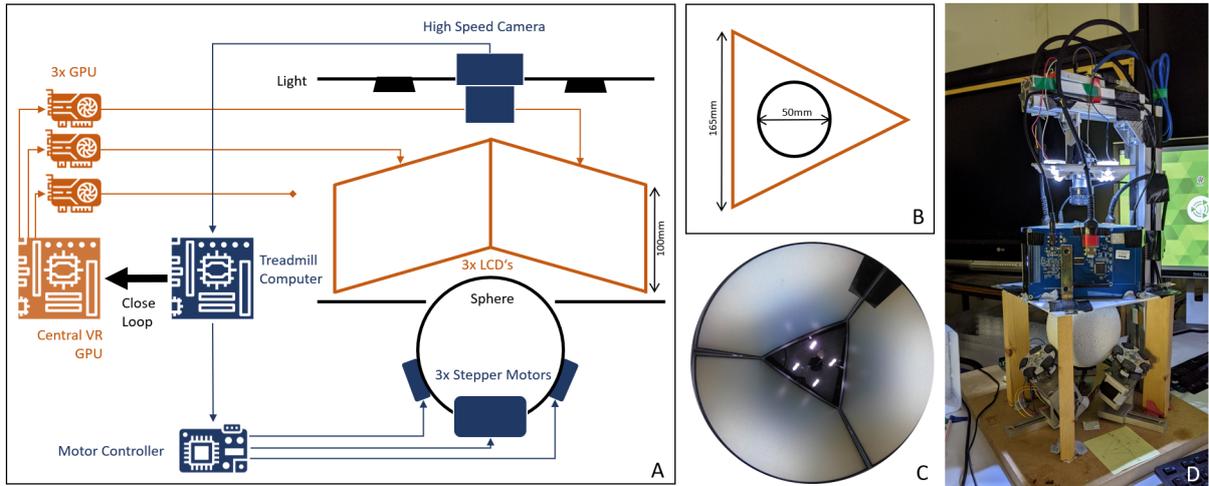


Figure 3.1.: Experimental Setting. (A) Information flow between systems: treadmill subsystem (blue) and VR subsystem (orange). The black arrow symbolizes communication in the close-loop condition. (B) Top view. VR system is composed of three screens arranged in an equilateral triangle around the trackball. (C) Panoramic picture from inside the VR chamber. The horizontal lines of the displayed visual elements are distorted to accommodate the system's geometry and appear as projected equidistantly from the ants' perspective. (D) Picture of the experimental setup.

omni wheel (aluminium, diameter: 60 mm, 2 discs, 5 rollers/ disc) to avoid creating friction when the ball has to rotate perpendicularly to the axis of the rotor. The compensation signal is carried out by the three motors to rotate the sphere and keep the ant at the centre of the ball. A white board covers the space around the top of the sphere (5 cm) and supports the three screens of the VR system. A small webcam is additionally fitted to the cardboard ceiling and allows the experimenter to look inside the chamber without disturbing the ant.

Motion Tracking

The position and orientation of the ants are tracked every 700 frames^{-1} . The exposure time is set to $1000 \mu\text{s}$. A custom Python program developed by Goulard et al. [6] and slightly modified by me is used. Each frame is binarized using a threshold of 50% and reversed such that the ants appears as a white "blob" in a black background. Based on previous frames and size, an ellipse is fitted around the contour of the blob that most likely represents the ant. The centre and orientation of the ellipse are calculated. The orientation is approximated with a $\pm 180 \text{ deg}$ uncertainty, which is resolved in post-processing by assuming the ant moves forwards, and kept consistent by minimizing the heading change between two consecutive frames. The initial orientation is determined by the experimenter by looking at the ant though the webcam and "flipping" the orientation in case of a mismatch. Sudden movements of the

ants might cause a mismatch in orientation further into the trial. These are addressed in the data analysis.

Controller

The x and y position of the ant, estimated from the centre of the tracked ellipse, are used to actuate the motors and keep the ant in the centre of the sphere. The system employs two proportional derivative (PD) controllers, one for each coordinate. For a detailed implementation of the controller, see [6]. The only change from the previous implementation of the system are the control parameters. After mounting the VR and treadmill system together, I manually optimized the proportional and derivative gains ($K_p = 1$, $K_d = 0.05$). The control feedback loop is kept at a frequency of approx. 500 Hz.

3.1.2. Virtual Reality System

The VR system (Fig. 3.1A and B orange) is composed of three LCDs (WINSTAR TFT-LCD Module 7") with a resolution of 800x480 px. The screens (165x100 mm) are mounted on the white board around the sphere, forming an equilateral triangle. Each screen is powered independently and connected via HDMI to its individual graphics process unit (GPU). All three screen-GPUs connect via micro-USB to a central GPU running Ubuntu 17. The VR computer runs a custom python program developed by Goulard et al. [40] and slightly modified by me. The communication between the GPUs is coordinated by an instance of the Robot Operating System (ROS); the central VR GPU serves as master and the individual screen-GPU's subscribe to it.

The VR system supports the display of simple visual shapes and natural panoramic images. Regardless of the input, the displayed visual elements are distorted to fit the screen geometry, i.e. as the image elements move towards the corners, the horizontal lines are scaled disproportionately to accommodate for the change in perspective. Hence, from the centre of the sphere, the visual elements all look as if projected equidistantly. Figure 3.1 shows the result of the distortion on the example of a black rectangle on one of the corners.

The system is also able to display panoramic images. These are first unwrapped (black pixels are used to fill up empty space) and then distorted to accommodate for the change in perspective. The system is able to display full colour images, nevertheless for the experiments presented in this thesis, we transform the images into black and white (Fig. 3.3).

3.1.3. Systems Integration

The treadmill and VR systems communicate using ROS. The treadmill computer serves as ROS-master, and the VR computer subscribes as a listener. This way both systems stay in synchrony. The experimental conditions, i.e. type and duration of visual display, are defined

by the VR system. This one remains in standby until a new "trial" is started globally by the treadmill system. Communication happens at a rate of approx. 65 Hz. At each time step, the treadmill system communicates the estimated heading direction of the ant to the VR system. Two important data outputs are generated. The treadmill outputs the angular velocity of the sphere at a rate of approx. 113 Hz and the VR system outputs the heading direction of the ant at a rate of approx. 65 Hz. The timestamp of each data entry is synchronized by the ROS.

System Open- and Close-Loop

The integration of both systems allows to manipulate the feedback loop between the ant and its environment. From a systems' perspective, the natural condition of the world is open-loop. This means that the ant has an influence on how she perceives the environment by moving, e.g. when the ant rotates clockwise, the visual features appear to rotate anti-clockwise. This is an example of a rotational open-loop. In translational open-loop, when the ant moves towards an object, the projection of the object on the ants' retina grows.

In its current stage of development, the VR system is not yet capable of simulating a translational open-loop, i.e. when the ant moves towards a black rectangle, the display size of the rectangle stays the same. Hence, our system is constantly in translational close-loop.

For rotation, the system offers the possibility to change between open- and close-loop. If a visual stimulus stays at the same position on the screen, the system is simulating the natural open-loop condition. However, if we rotate the visual stimuli on the screens by the same amount and direction as the ant rotates, then the system is in close-loop and the ant has no control over what she sees. This is a powerful tool to analyse the neuronal circuitry of the insects. To achieve this condition, the VR system takes the estimated orientation of the ant communicated by the treadmill and rotates the displayed stimulus accordingly.

3.2. Experiments

To test the functionality of the system and learn more about the visual strategies employed by ants, we designed two groups of experiments. All the experiments were carried out at the University of Sussex in Brighton, UK.

3.2.1. Wood Ants

We work with wood ants (*Formica rufa*) that are kept in the laboratory (Fig. 3.2). The colony was collected from woodland (Broadstone Warren, East Sussex, UK) and is housed in large tanks in a lab with regulated temperature of 20-25°C. Water, sucrose and dead crickets are provided *ad libitum* on the surface of the nest. Ants are kept under a 12 h light to 12 h dark cycle. During the experiments, the food supply is kept at a minimum in order to increase



Figure 3.2.: Wood ant. Image shows a forager of the species *Formica rufa* standing on the trackball while facing a black bar as a unique visual cue in an otherwise white environment.

foraging motivation. Water access remains constant.

During an experiment, we are mostly interested in ants with a strong motivation to forage. To only select "motivated" ants, we position a small plastic box without lid on the surface of the nest. Motivated foragers climb inside the box. The inner walls are painted with a fluon solution, so the ants can not climb back out. Only ants that pass the "box-test" are tested on the VR system.

Ant Marking

In some experiments, individual ants are tracked across multiple trials. Between each trial, the ants are returned to their nest momentarily. In order to differentiate them, ants in these experiments are marked. We use acrylic paint in different colours to mark the thorax and abdomen of the ants with a tiny drop of paint. Through a colour code, we are then able to identify individuals.

3.2.2. Experimental Procedures

The general procedure to test the ants is constant across all trials. After depositing the ant on the surface of the trackball with a plastic flag, the experiment is started. Within the first 30sec, the experimenter determines if the heading orientation is being tracked correctly, and otherwise flips it with a keyboard command. The ant stays on the treadmill until (i) the time is over (either 3 or 10 min), (ii) the experimenter notices that the ant tracking flipped the heading direction, or (iii) the ant escapes the treadmill. In any case, the tracking is stopped, and the ant is collected. Trials that incidentally stop in the first 30% of the total time are not saved. In these cases, the same ant is tested again.

Experiment	Conditions	Trials	Time Frame
Open-loop	C20: s0, s1, s2, s3	43	27-29.04.2021
Open-loop	C20: s1, s2, s3	26	13.07.2021
Open-loop	C40: s1, s2, s3	45	14-16.07.2021
Close-loop	0, ± 30 offset	69	21-22.04.2021
Close-loop	± 20 offset	31	12-13.07.2021
Close-loop	± 10 , 40 offset	64	17-19.07.2021
Artificial Pattern	A1, B2, A3, B4	95	09-11.06.2021
Natural Image	A1, B2, A3, B4	74	13-14.05.2021
Natural Image	A5	13	18.05.2021
Natural Image	B1, C1	35	13-14.05.2021

Table 3.1.: Overview of recording days and sample sizes for all four datasets.

Experiments with a Simple Visual Cue

Two datasets with a simple visual cue were collected: open- and close-loop. Both follow the same experimental procedures. A simple black rectangle is used as a unique visual landmark in an otherwise white environment (Fig. 3.2). The size of the cue is of 35 deg high and 20 deg wide. The close-loop dataset uses an additional experimental condition where the cue dimensions are altered to 35x40 deg. All the experiments in this group last a maximum of 10 min. In this group, only naive ants are tested. Table 3.1 provides an overview of the test dates and sample sizes.

Experiments with a Complex Pattern or Image

In order to test the VR’s ability to display complex visual stimuli, as well as explore the visual ecology of wood ants, two datasets were collected: artificial pattern and natural image. Both have similar experimental procedures. In this group, we test individual ants multiple times (see Ant Marking). The marked ants are returned to their nest after each trial for approx. 3 hours. In the natural image experiment, the fifth and last trial is conducted four days later. Table 3.1 provides an overview of the test dates and sample sizes.

Artificial Pattern The visual stimulus used in this dataset is a panoramic pattern of black rectangles (Fig. 4.11). The rectangles are organized in groups. There are three groups: one wide rectangle, two middle-sized rectangles and four thin rectangles. The height of all bars is the same. The panoramic pattern was designed in Cartesian coordinates and wrapped into polar coordinates using an external function (MATLAB 2021a, *PolarToIm(img, 0, 1, 1024, 1024)* [41, 42]). The polar pattern is then fed into the VR system. This process is followed for each rotation of the pattern.

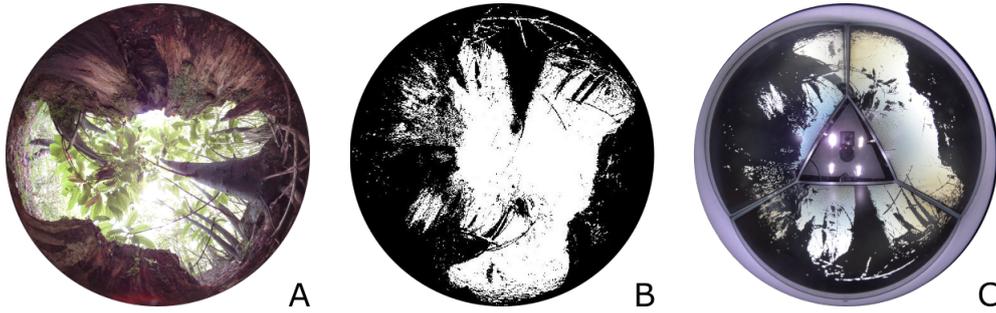


Figure 3.3.: Natural woodland image. (A) Raw image in original rotation, (B) image from A rotated by +90 deg and transformed to black and white. (C) Picture from inside the VR chamber with condition C, which is image A rotated by -90 deg.

Natural Image For this dataset we use a panoramic picture taken in the woods (Abbots Woods, East Sussex, UK) using a fish-eye camera (Kodak SP360 4K). We selected the image for its texture richness. The raw picture is rotated according to the desired condition (A, B, C). We binarize the pictures before feeding them into the VR system. Figure 3.3 shows each of the steps.

3.3. Data Analysis

For the entire data analysis, I use MATLAB [41]. The data outputs (see 3.1.3) for each trial are saved in *.csv* format and imported into MATLAB. In total, four datasets were collected, namely: open-loop, close-loop, artificial pattern and natural image. Each is analysed independently.

3.3.1. Data Preprocessing

As described before, one of the shortcomings of the tracking system is the inability to differentiate between the head and tail of the ant, leading to an ± 180 inaccuracy. To deal with this, at the beginning of each trial the experimenter flips the heading direction manually in case of a wrong head-tail match (happens in 50% of the trails). During the trial, the treadmill system assumes a minimal change in heading direction between two frames and hence keeps this one constant. Nevertheless, if an ant makes a very sudden turn or stands on its hind legs (reducing the ellipse shape to a circle), the ant-tail matching might fail again. These are ant-induced flips (happens in about 10% of the trials), and the data afterwards is not useful any more.

All datasets undergo the same preprocessing step. Based on the recorded heading direction, I calculate the angular velocity each 0.2 seconds. A threshold in the angular velocity is manually fitted to each dataset to identify strong sudden changes in heading direction, i.e. a flip either manually induced by the experimenter at the beginning or by the ant due to a strong sudden

movement (close-loop: threshold = 455 deg/second, all others: threshold = 750 deg/second). A correction induced by the experimenter at the beginning needs to happen within the first 60 sec. Data before this flip is cut off. Any subsequent flip outside this time frame is considered ant-induced, and the data after the flip is cut off. Hence, only data after the human correction and before an ant-induced-flip, is considered further. Finally, the total recording time after the cuts needs to be greater than 10% of the experiment time, trails that do not match this criterium are discarded.

Afterwards, all heading angles are rotated by 90 deg to accommodate to the system's coordinate system, which has a zero value at the centre of screen one and positive angular velocities in counterclockwise rotations.

3.3.2. Chunks and Activity Levels

I use the activity of the motors as a metric to segment ant behaviour. The underlying assumption is the following: ants are attracted to visual landmarks. When an ant moves towards a landmark, it will decrease its angular speed and increase its forwards speed. A high forward speed on the treadmill is reflected in a high (counter) speed of the trackball. Hence, the level of ant attraction is proportional to the rotation velocity of the trackball.

To exploit this, I segment the motor angular velocity data into chunks of uninterrupted activity. Figure 3.4 illustrates the chunk extraction process. First, I sum the x (A) and y angular velocities ($\dot{\Theta}_x$ and $\dot{\Theta}_y$) and take the absolute (B). Each data entry is then classified into a quadratic scale, assigning each entry a value between one and four based on the strength of the signal. To encode the density of the signal, I calculate the average value over a 2 sec window (C). Here, all values under 0.05 are rounded down to zero. The resulting step function encodes both the density and the amplitude of the signal. I smooth this function using a moving median and a window of 1000 entries (D).

The smoothed step function is divided into chunks. A chunk is a period of time within which the step function never reaches zero. For each chunk, I determine the integral under the step function. Integrals under a value of 100 are discarded as noise. All other integral values (one per chunk) in a dataset are pooled together, and the quartiles are calculated. I use the quartiles as boundaries and assign each chunk an activity level between one and four (E).

Since the treadmill and VR data are synchronized, I can take the time interval of each chunk and apply the same division to the ant's heading data. The activity level, which is solely based on motor activity, is kept. As an end result, I have the ant's heading direction data segmented and categorized into activity chunks (F).

3.3.3. Individual Analysis

Although the overall analysis of the data followed the same methodologies, in each dataset I performed some specific manipulations in the data. This section gathers these techniques.

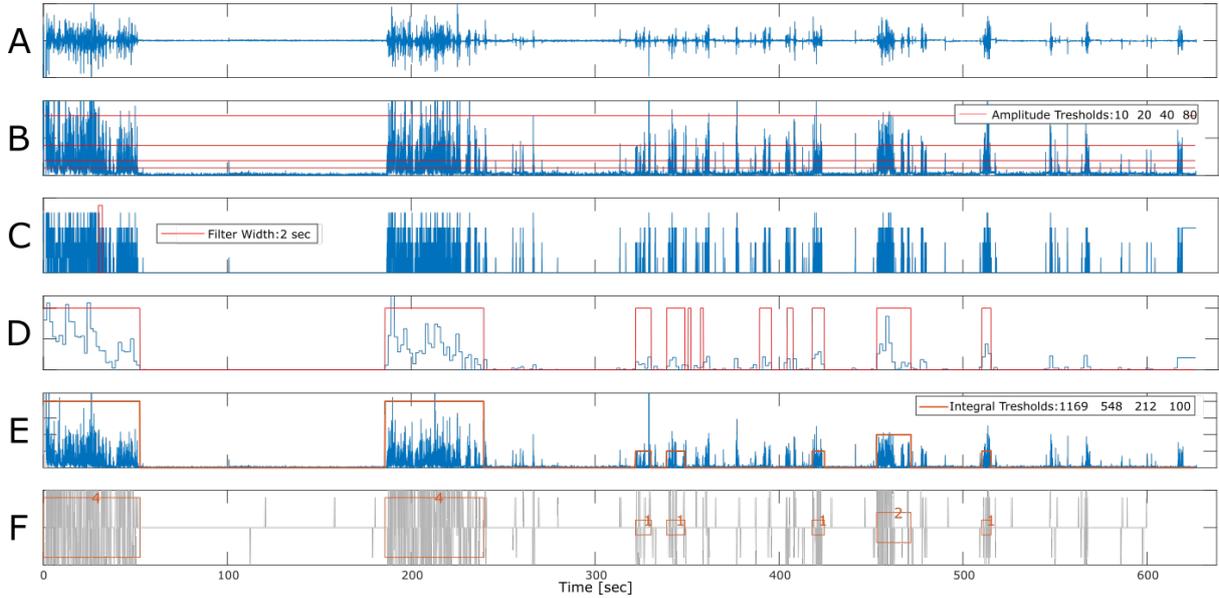


Figure 3.4.: Activity chunks extraction process. Motor angular velocity in blue, ant angular velocity in grey. (A) Raw motor angular velocity in x direction ($\dot{\Theta}_x$). (B) Absolute sum of $\dot{\Theta}_x$ and $\dot{\Theta}_y$. Thresholds used to sort data in next step. (C) Motor activity divided in four levels according to amplitude. Density of the signal is encoded by the average amplitude in a 2 sec window. (D) Contour of the smoothed motor activity. Data is divided into chunks within which the motor activity does not reach zero. (E) Motor data (same as B) with the extracted chunks. The integral under the contour of each chunk is calculated and divided into one of four quartiles of activity. (F) Angular velocity of the ant with the chunk divisions and labels calculated on the base of motor activity.

Circular Histograms

For all circular histograms presented in the results, I use the `polarhistogram()` function with a bin size of 5 deg and a normalization set to probability. The mean and variance values of all circular data are calculated using the circular statistics toolbox [43]. All histograms within a dataset are scaled by the same factor. The dots on the outer rim of the plot have the mean angle for an individual ant or chunk and are colour coded to represent the relative variance of that mean. To do this, I pool together the variance value of all ants in that condition, calculate the quartiles and use these to assign each dot a fill-strength: fully-filled = low variance, no-fill = high variance.

Moments of Fixation

For the experiment presented in section 4.1.2, I pool all screen variations together. To achieve this, I rotate the heading direction of the ants by the offset of the respective condition to

screen 1. Then I define "moments of fixation". An ant is considered to fixate a visual cue if its heading direction stays within ± 45 deg of the centre of the cue for at least 5 seconds.

Mean Heading Direction in a Moment of Fixation A qualitative analysis of the data led to the assumption that ants fixate the edges of the cue rather than the centre. To test for this hypothesis, the mean heading direction of the fixations is calculated. Most of the ants switch between cue edges within one fixation once or even twice. In these cases, the mean heading direction balances down to the cue centre (by definition 0). To account for this, fixations that cross the centre of the cue for longer than 30% of the entire duration of the fixation are divided at the crossing point. If this is the case, two individual mean values are calculated, one for the values below zero and one for all values above zero. In subsequent analysis, I weight these means by the amount of data entries they are calculated from. The 30% threshold is determined through a qualitative analysis of the splits.

Turning Points, Frequency and Amplitude

To analyse the oscillations in heading direction, I extract the "turning points", i.e. points where the ant starts turning in a new direction. To do this, I first smooth the data using the *smoothdata()* function. In the case of the open-loop, I only calculate the turning points of data within a ± 45 deg range centred to zero by definition, so the data can be smoothed directly with a factor of 0.15. In the case of the close-loop, the angle range is much greater, so I first unwrap the data (*unwrap()*), smooth the data by a factor of 0.015 and wrap it back to $\pm\pi$ (*wrapToPi()*). For both conditions, I then calculate the differential of the smoothed heading direction, i.e. the angular velocity between every data entry. A turning point is a point where the angular velocity switches sign or is exactly zero.

The frequency and amplitude are calculated as the difference in time and heading direction between two subsequent turning points. I filter out turning points (and hence amplitudes and frequencies) where the corresponding amplitude is smaller than 0.5 deg in the open-loop and smaller than 5 deg in the close-loop. This serves as a noise filter.

Angular Velocities The angular velocities presented in the open-loop results are calculated using the *diff()* function. The range of the data allows for the normal differential function. Since the data here is split in small fixations, the angular velocity has to be calculated before smoothing to avoid losing the expression on the edges. The angular velocities in the close-loop are calculated after smoothing using the *circ_dist()* function [43]. This is necessary due to the step-like character and range of the data.

Image Difference Function

In section 4.2.2 I evaluate the similarity between the display of the three different rotations of the same image. To do so, I took a panoramic picture from inside the experimental chamber (we removed the trackball and placed a fish-eye lens, Kodak SP360 4K instead). This is an approximation of what the ant sees inside the chamber. I unwrap the pictures using the external function *ImtoPolar*(*img*, 0.35, 0.85, 700, 3072) [42]. I then binarize the resulting image and flip it vertically to accommodate to our perception.

The rotation image difference function calculates the pixel wise root-mean-square difference between each 10 deg rotation of image A against all images, A, B and C.

3.3.4. Statistics

All statistics are calculated using MATLAB and the Statistics and Machine Learning Toolbox [41]. For all circular data, I use an external toolbox [43].

Given the high frame rate of the treadmill system, the data is quite dense, e.g. a 10 min recording of ant heading generates around 50,000 entries. The statistical tests employed to analyse the data, namely circular Rayleigh-Test (RT), circular V-Test (VT), circular K-Test (KT) and Kolmogorov-Smirnov Test (KS-T) are not build to handle such high sample sizes [44].

To deal with this, when performing the statistical tests, I take only every x data entry, where x is the data length divided by a factor of 200. This value was manually adjusted based on the dimension of the p values ($10^{-20} < p < 1$) and is kept constant throughout the entire analysis. Given the arbitrariness of this adjustment, all statistical results presented in this thesis should be taken as an orientation. To represent the magnitude of very low p values, if p is smaller than 10^{-4} , I use the notation $p \ll 0.05$.

4. Results

The use of VR in behavioural experiments opens new opportunities to investigate the visual ecology of ants. The novel system described in chapter 3 allows ants to move freely (un-tethered) without changing location and gives the experimenter full control over the visual stimuli the ants are subject to. In this chapter, I describe the performance of the overall system in two parts.

First, I explore how the treadmill and VR system integrate by presenting a single visual bar either in open or close-loop. From a systems' perspective (engineering perspective), open-loop is when the agent has control of how the environment looks like, i.e. if the ant moves towards something or rotates in one direction, its visual perception of an object grows, or rotates in the opposite direction. This is how an animal naturally perceives the world. Contrarily, in close-loop the agent has no power over their perception of the environment, if the ant moves, the retinal image stays the same. It is worth mentioning, that in biology these terms are used inversely; from the perspective of the animal, close-loop is the natural state of the environment. Here we stick to the engineering definition.

At its current state of development, only rotational open-loop is supported by the VR system. This means, that when the ant moves towards a visual cue, the retinal projection does not grow as it would in the real world. Hence, for all the experiments described in this thesis, the system is in translational close-loop, and we can only change between rotational open- and close-loop.

The second part of this chapter deals with the representation of more complex patterns and images on the VR system. I explore the limitations in representation and the orientation behaviour of ants in response to complex visual scenes.

4.1. Treadmill and VR Integration in Open- and Close-Loop

Wood ants are innately attracted to conspicuous visual cues [45]. Most insects are attracted to dark visual objects, and bars and cylinders are commonly used in behavioural experiments [33, 46].

We start testing the open-loop VR-system with the display of a simple black bar against a white background. We test 111 naive ants in seven different conditions. The visual cue is displayed at the centre of one of the three screens (S1, S2 and S3) or not at all (control = S0). The height of the bar remains constant (35 deg), however two different widths, 20 deg (C20)

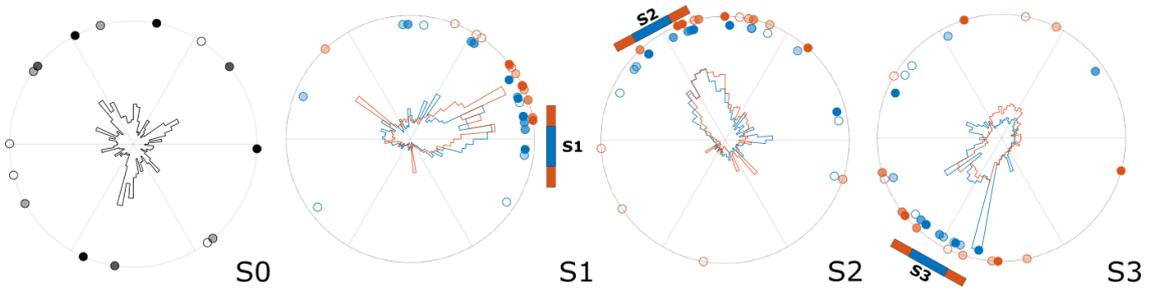


Figure 4.1.: Polar histograms of heading directions in response to a simple black cue. The width of the cue is 20 deg (blue) and 40 deg (orange) respectively. The cue is displayed in either of the three screens (S1, S2, S3) or not at all (S0). Rectangles on the edge of each histogram show the position of the respective cue. The circular histograms show the probability distribution for the entire group, bin size 5 deg. The circles show the mean heading direction for each ant in the group. The circle shade of colour encodes the variance in quartiles, fully-filled = low variance, no-fill = high variance.

and 40 deg (C40), are tested. We use both widths in all three screen variations. Each of the resulting seven groups has approx. 15 ants (S0: $n=15$; C20: S1 $n=18$, S2 $n=16$, S3 $n=16$; C40: S1-S3 $n=15$). Each ant is tested for a maximum of 10 min. Figure 4.1 shows the ants' heading directions in response to the simple stimulus in each condition.

The ants show the expected (innate) attraction towards the visual cue in all but one conditions of the bar (S1 C20 and C40: Rayleigh-Test (RT) $p \ll 0.05$, V-Tests (VT) $p \ll 0.05$; S2 C20 and C40: RT $p \ll 0.05$, VT $p \ll 0.05$; S3 C20: RT $p \ll 0.05$, VT $p \ll 0.05$; S3 C40: RT $p = 0.39$). The ants in the control are not significantly directed (S0: RT $p = 0.4$), and the variation in cue width does not significantly influence the heading distribution (K-Test (KT) C20 vs. C40: S1 $p=0.69$, S2 $p=0.32$, S3 $p=0.18$).

The results validate the ants' ability to express their natural behaviour despite the novel, unnatural experimental setting. The display of visual information in either of the three screens is equivalent. Furthermore, the control condition lets us rule out any directional biases caused by the setup. Nevertheless, the presented (raw) data is quite turbulent, and the ants' behaviour is shadowed by noise introduced by the experimental setup and settings.

4.1.1. Using Ant Activity as a Metric to Filter Behaviour

Although the experimental setting allows the ants to move untethered, the movement and vibrations of the treadmill introduce some noise into the data. Furthermore, as discussed before, the visual stimuli presented to the ants does not change in response to the ants' translational movements, i.e. the cue does not "grow" as it should if the ant would approach it naturally. I assume this to have an effect on the ants' motivation to continue moving towards a landmark they do not seem to ever get close to. This section investigates a method to filter

Activity	S0	S1	S2	S3	Σ S1-S3
75-100%	13	13/13	14/13	11/09	38/35
50-75%	21	12/12	9/03	18/11	39/26
0-50%	59	28/17	34/08	20/07	82/32

Table 4.1.: Chunk sample sizes divided by activity level and conditions. Values for C20 and C40 before and after the slash respectively (C20/C40).

the ants' behaviour.

Roughly, the ants' behaviour while on the treadmill can be divided into one of three activities: (i) resting, (ii) exploration and (iii) attraction. Most of the ants engage in resting behaviour (i) at least once while on the treadmill. Here, the ants stop running and keep the same heading direction for a longer period of time. Video recordings from inside the VR chamber often show the ants rubbing their antenna with their front legs during these resting periods. An extreme example of this is seen in Figure 4.1 S3, just at the edge of the cue, the blue histogram is significantly bigger for a single bin.

In the exploration activity (ii), the ants constantly change their heading direction. The control condition is a good example of how the heading distribution looks liked during this activity (Fig. 4.1 S0). During exploration, the ants might rotate through the cue without acknowledging it (if there is one) or change into attraction (iii) behaviour upon encounter. This last activity is when the ants display a clear attraction for the visual cue and attempt to run towards it. A high activity level and low heading variance are typical for this activity. The overall directness of the S1 to S3 conditions is due to this activity.

As discussed in detail in chapter 3.3.2, I developed a system to categorize the ants' activity during the experiment. Based on the movement of the ball (i.e. angular speed induced by the motors) I divide the recording in so-called "activity chunks". Each chunk is defined by an integral value that encodes the duration and the amplitude of the activity; a high integral is a relatively long and strong moment of ant activity. The integral is completely arbitrary as an absolute value, but functions as a relative comparison.

I analyse the activity for the experiment with the simple black cue. Chunks with an integral under a threshold (100) are initially filtered out as noise from the motor vibrations. A total of 454 chunks across all conditions are extracted. The integrals of all other chunks are pooled and the quartiles calculated. The data is then split into 4 levels of activity using the integral quartiles as boundaries (see Tab. 4.1 for sample sizes). Figure 4.2 shows the heading direction data split into activity levels. The lower two levels are pooled together (0-50%) as they displayed the same characteristics in the distribution.

The heading direction is not directed for any cue location or cue width in the lowest three activity levels (0-75%; S0-S3, C20 and C40: RT $p > 0.05$ for all). In the highest activity level (75-100%), all but one conditions are directed and have a mean towards the centre of the

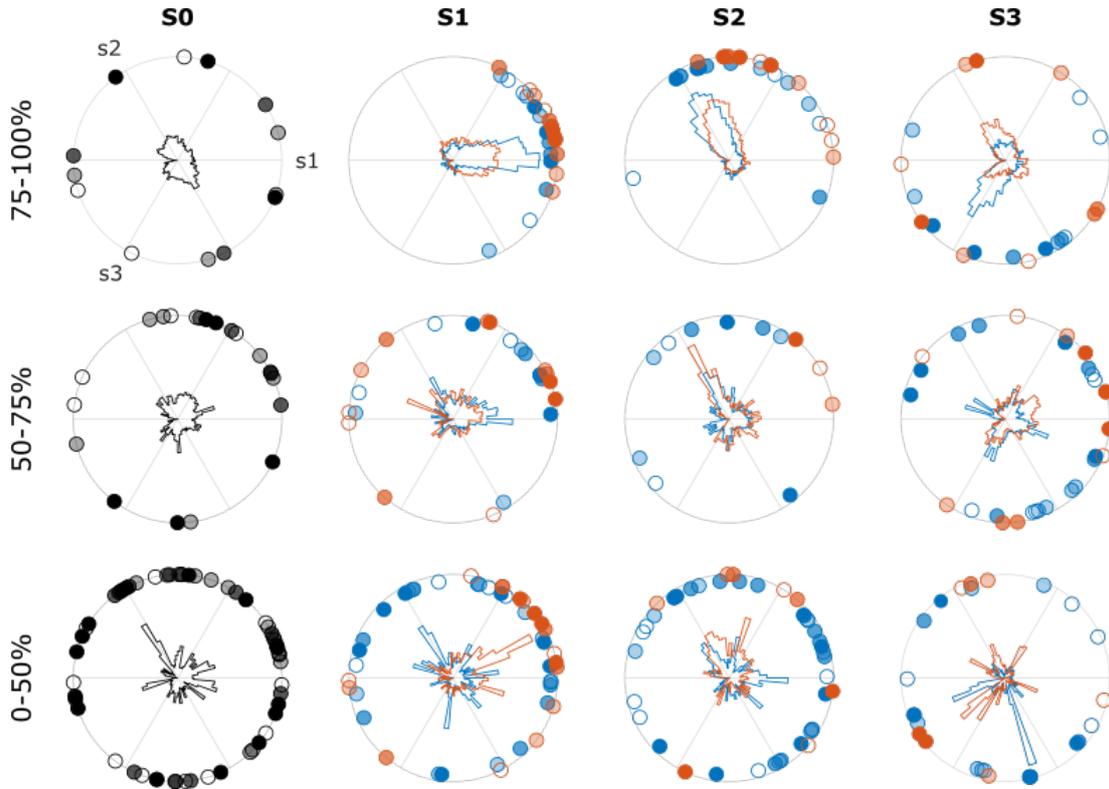


Figure 4.2.: Polar histograms of heading directions divided by activity levels. Each row shows the heading data within a level of activity. The width of the visual cue is 20 deg (blue) and 40 deg (orange) respectively. The cue is displayed on either of the three screens (S1, S2, S3) or not at all (S0). The angle of the centre of each screen is marked in the top left plot. The polar histograms show the probability distribution for the entire group, bin size 5 deg. The circles show the mean heading direction for each ant in the group. The circle shade of colour encodes the variance in quartiles, fully-filled = low variance, no-fill = high variance.

respective visual cue (S0: RT $p=0.8$; S1 C20 and C40: RT and VT $p \ll 0.05$; S2 C20 and C40 RT and VT $p \ll 0.05$; S3 C20: RT and VT $p < 0.05$, C40: RT $p=0.08$).

The level of activity is a good metric to separate the ants' behaviour. This separation of the data roughly corresponds to the three behaviours described above. In the lowest two activity levels, we can see single, very concentrated bins characteristic for resting behaviour (i) for all four groups. The angle of the resting moments is randomly distributed, hence the visual cue has no effect during resting. Activity between 50 and 75% is equally spread all around the circle, but less punctuated. This matches the characteristics of exploration behaviour (ii): high activity and heading variance, with a slight influence by the visual cue. Exploration can also be seen in all conditions. Attraction behaviour (iii) can be clearly seen in the highest activity level. The histograms are (mostly) directed towards the visual cue, and the motivation of the

ant to pursue the cue is mirrored in the strong motor rotation.

Interestingly, the control condition shows some exploration behaviour in the top level (Fig. 4.2 S0 75 - 100%). This suggests that this might be the case for the other conditions as well, and might explain the lack of directness in the S3's conditions. We also need to consider that the ants in the control condition could still be engaging in low attraction behaviour to any random element of the visual environment.

The ants are able to express some of their natural behaviour on the treadmill. The experimental setup introduces some noise into the data, yet I have developed a methodology to filter our noise and extract the ants' behaviour. This serves as a foundation for a more in depth analysis of behaviour.

4.1.2. Landmark Attraction in the Open-Loop

Instead of moving straight forward, wood ants are known to oscillate when they are navigating towards a landmark. This strategy allows them to adjust their heading direction and extract parallax information [46]. In this section, I investigate how ants fixate a simple, conspicuous visual landmark on the treadmill. I start by describing the path details of ants attracted to the visual cue in the rotational open-loop experiment.

Parting from the results of section 4.1.1, we are only interested on the attraction behaviour and thus only consider the chunks in the top level of activity (75-100%). Furthermore, I pool together conditions S1 to S3 by rotating the data such that all screen centres align (S2 rotated by -120 deg and S3 by +120 deg). The control is also not considered any further. Hence, we have two conditions of the visual cue, C20 and C40 with 38 and 35 chunks of strong activity respectively.

To further investigate cue attraction, I define "moments of fixation" within the activity chunks. An ant is considered to fixate the landmark if its heading direction is within ± 45 deg of the cue centre for longer than 5 seconds. I extracted 138 moments of fixation for C20 and 115 for C40, with an overall mean duration of 11 sec. Figure 4.3 shows the characteristics of a fixation with two examples.

Edge Fixation

Ants on the treadmill oscillate to fixate the visual cue. Instead of fixating the centre of the cue, they alternate between fixating the left and the right edge of the cue (Fig. 4.4). In figure 4.3, the blue path shows a typical fixation of the left C20 edge (10 deg). The mean heading angle along the path (blue dashed line) is 14.2 deg which roughly corresponds to the left edge.

However, a lot of ants do not only fixate one of the edges, but alternate between them during a single fixation (orange path). In these cases, the mean of the entire path would balance out to approx. zero. To investigate the actual fixation angle, I separate the fixations at the crossing point and calculate two separate means. In the orange path, the ant oscillates around

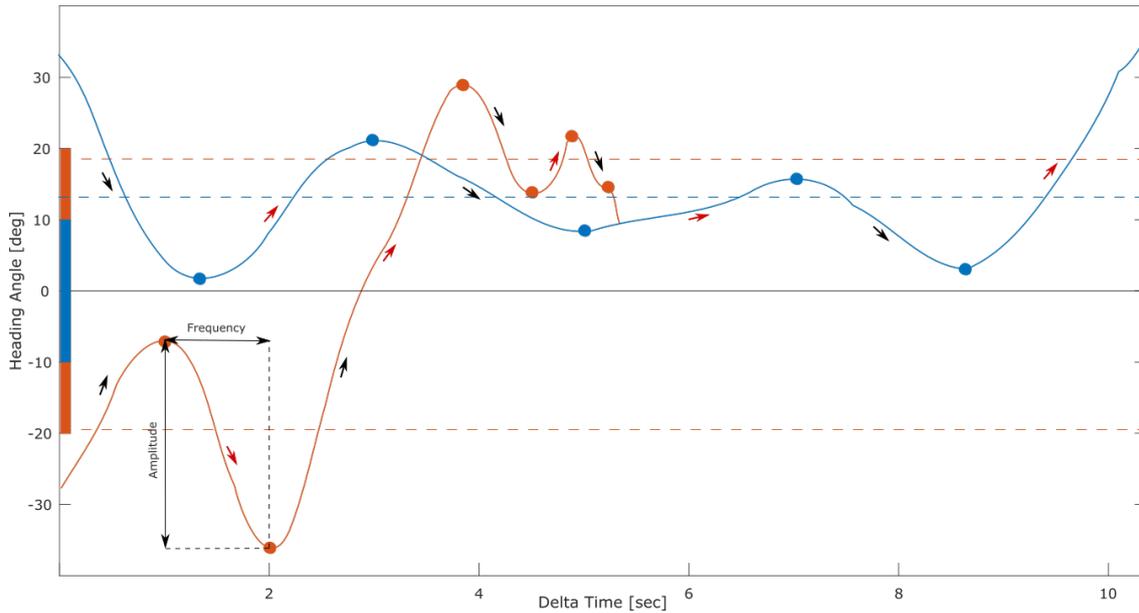


Figure 4.3.: Heading direction during a fixation. The plot shows two examples of a fixation, one for the C20 condition (blue, ant 7 fixation 11) and one for C40 (orange, ant 75 fixation 10). Paths smoothed for representation. Dots represent the turning points. Frequency and amplitude are calculated between two neighbouring turns. Each segment is categorized by its direction to the cue centre: towards the cue (black arrows) or away from the cue (red arrows). Dashed lines show the mean for each path respectively. Since the orange path crosses the cue centre, two mean values are calculated. Coloured bars on the left show the position of the cue.

the right C40 edge (mean: -19.48 deg) before crossing over and oscillating around the left C40 edge (second mean: 18.53 deg).

Figure 4.4 gathers the mean angles for all fixations. The histograms are weighted by the amount of data points represented by each mean value. In the C20 condition, we can see a clear peak in probability around both of the edges, -10 deg and +10 deg. In comparison, in the C40 condition, the probability is higher further away from the cue centre, however not quite at the expected ± 20 deg. A possible explanation is that the impulse to switch edges is governed by the distance between them: the higher the distance, the higher the need to switch. Hence, in the C40 position the ant sees both edges 40 deg apart, and has a stronger impulse to switch to the other edge. This spreads the distribution of heading directions and moves the means towards the centre. In fact, the paths in C40 have around 10% more crossings than C20.

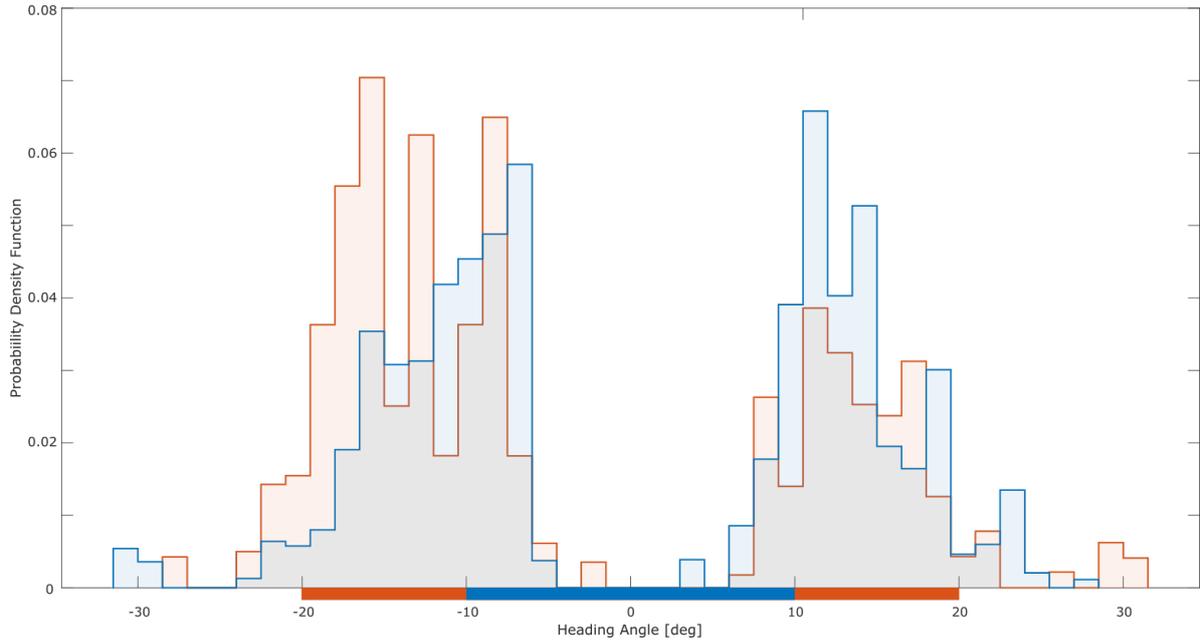


Figure 4.4.: Distribution of weighted mean headings. C20 in blue and C40 in orange. Bin size 1.5 deg. For each fixation, either two or one mean heading values are calculated, depending on the ant crossing the cue centre for more than 30% of the path or not. Each mean value is weighted by the amount of data points it represents.

Amplitude and Frequency of Oscillation

For each fixation, I extract the turning points, i.e. the moments where the ant switches moving direction. I calculate the frequency and amplitude of the oscillation as the time and angle difference between each neighbouring turning point (Fig. 4.3 illustrates this process). Figure 4.5 pools together all the fixations. The frequency is relatively constant across all fixations, with a mean value of 1.6 Hz in both conditions (Kolmogorov–Smirnov test (KS-T): C20vsC40 $p > 0.05$). The amplitude has a mean value of 11 deg in C20 and 12.8 deg in C40 (KS-T: C20vsC40 $p > 0.05$), it is however more spread towards the high amplitudes, specially in the C40 condition. This can be explained by the fact that the oscillations while changing cue edge tend to be of higher amplitude, specially in C40 where the ant needs to travel a longer distance between edges.

Angular Velocity During Fixations

Lastly, I investigate the angular velocity during a fixation. I break down the oscillation and categorize each turn according to its direction: (i) towards the cue or (ii) away from the cue (Fig. 4.3 illustrates this with black and red arrows respectively). Figure 4.6 shows the angular velocity for ants moving towards (bottom) and away (top) from the cue at different angular

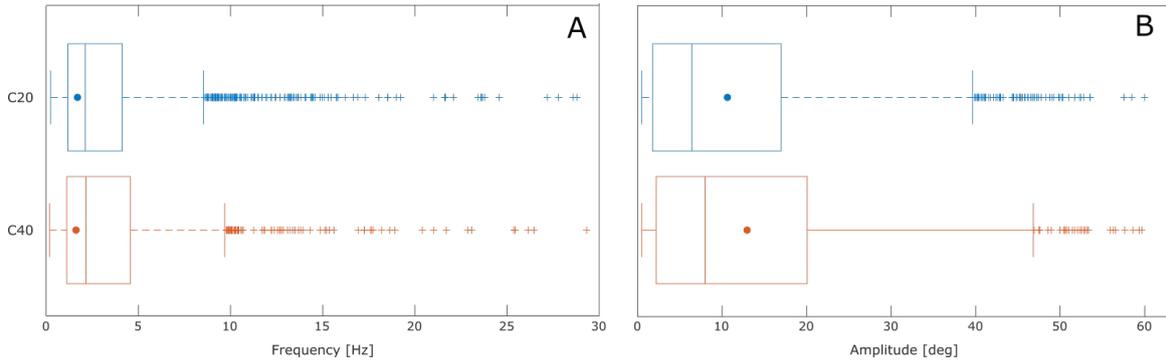


Figure 4.5.: Frequency (A) and amplitude (B) of the fixations. Box plots show the distribution of values for both conditions C20 (blue) and C40 (orange). The middle line shows the median, the box frame the 25th and 75th percentiles. Whiskers show all extreme values not considered outliers (crosses). Filled dots represent the mean values.

distances from the cue centre. In general, ants on the treadmill move at a speed of approx. 13 deg per second, which is consistent with our frequency and amplitude calculations.

In the towards speeds, there is a drop in velocity around 35 and 40 deg which rises again towards 45 deg. This is the same angle upon which Goulard et al. [40] measured ants to first react to the presence of the visual cue on a previous version of the VR system. Hence, when moving towards the cue, the ants slow down upon first encounter with the black rectangle. The edges of the respective cues are roughly marked by a speed decrease, the C20 angular velocity falls slightly between 10 and 15 deg in a similar dimension as the C40 does between 20 and 25 deg. Furthermore, the C40 speed is higher than the C20 for the lower distances (20-0 deg). Again, we can explain this with the fact that the ants need to travel a higher distance between cue edges and hence keep a higher velocity as they move past the cue centre.

Ants in this experiment are in rotational open-loop. They express the same oscillation behaviour [46] and edge attraction [47] previously described in ant research. Both behaviours rely on the visual feedback the ants gathers while navigating. To better understand the neuronal regulations of these behaviours, we exploit the experimental setup and test the ants in rotational close-loop.

4.1.3. Ant Behaviour in the Close-Loop

One of the remarkable opportunities that come with the VR system, is the possibility to close the loop between the ant and its environment. This section seeks to validate the close-loop condition with a comparison between the results in section 4.1.2 and a set of new close-loop experiments. Furthermore, I explore the implications of the closed-loop on the ants' behaviour.

4. Results

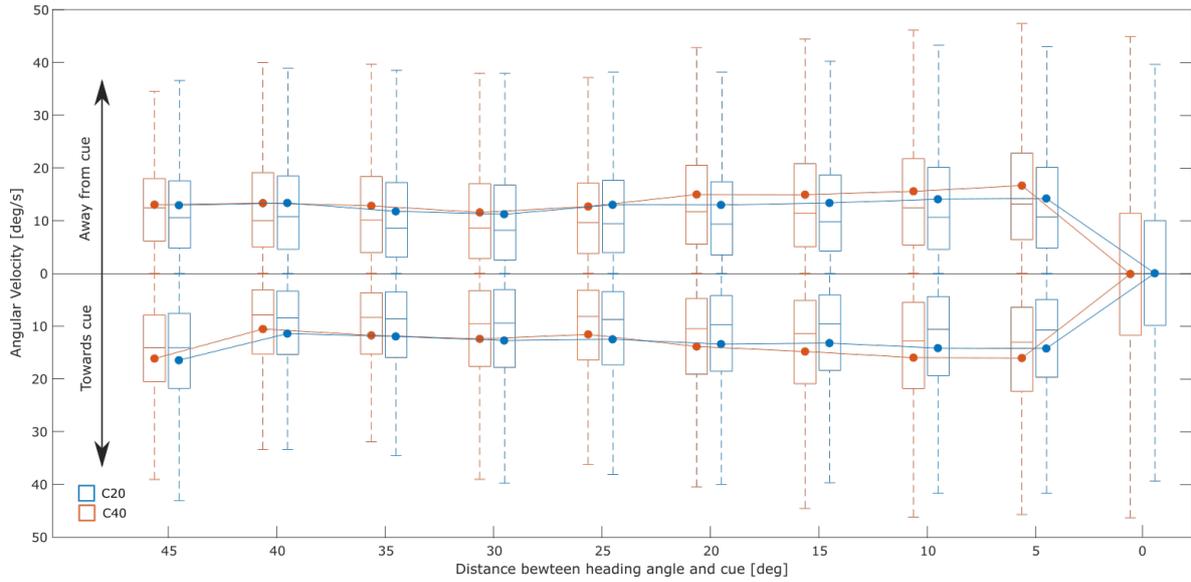


Figure 4.6.: Angular velocity during periods of fixation. Box plots show the distribution of angular velocities in relation to the distance of the ant from the cue centre. Each box plot gathers the velocities in ± 5 deg range. The middle line shows the median, the box frame the 25th and 75th percentiles. Whiskers show all extreme values, not considered outliers (not displayed). The dots show the corresponding mean for each group. Bottom half of the plot shows velocities while approaching the cue, top half while moving away from it. Condition C20 in blue and C40 in orange.

Using the same simple visual cue (35×20 deg, equivalent to the C20 condition in open-loop), we test 150 ants on the treadmill while on close-loop. The visual cue is rotated by the system to always be directly in front of the ant (offset = 0 deg, angular distance between cue centre and ant heading) or to have an offset of $\pm 10, 20, 30$ and 40 deg for a total of 9 different conditions (-40: $n=16$, -30: $n=22$, -20: $n=15$, -10: $n=15$, 0: $n=18$, +10: $n=13$, +20: $n=15$, +30: $n=22$, +40: $n=14$). Each ant is tested for a maximum of 10 min.

Following the assumption that ants have an innate attraction to the visual cue, I use the activity level methodology previously described to filter behaviour (Fig. 4.7). A negative offset between heading direction and cue centre should generate mostly negative angular velocities, and the opposite for positive offsets. The integral boundaries are similar in dimension to the open-loop experiments. Since by design the ant is always in viewpoint of the visual cue, in theory the ants should always be in attraction behaviour. This is also the case with every ant at the very start of the trial (all ants have a chunk of strong activity at the beginning). However, as the trial progresses, most ants temporally abandon the attempt to approach the cue and also engage in exploratory and resting behaviour.

I filter out unwanted behaviour in the lower levels of activity (0-68%). The top level of activity

4. Results

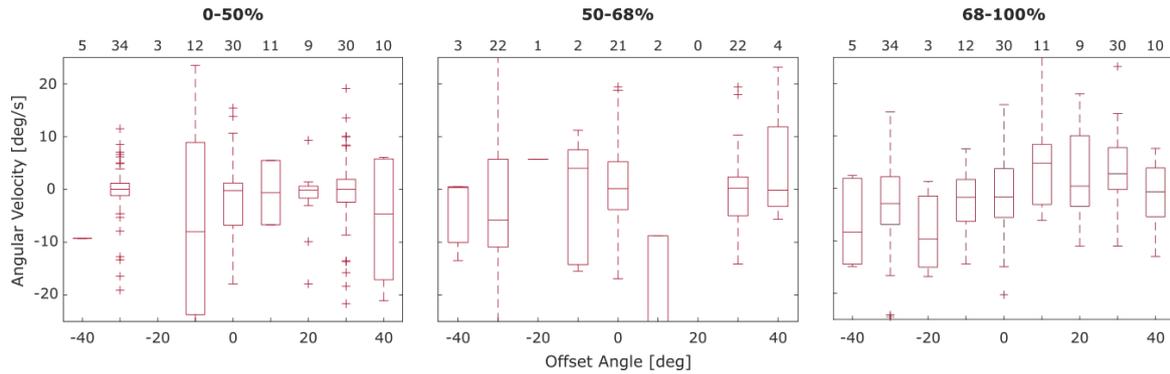


Figure 4.7.: Angular velocity for each offset condition and activity level. Box plots show the distribution of angular speeds for each offset of the cue. Each plot stands for one level of activity: 0-50%, 50-68% and 68-100%. Sample sizes displayed on top. The middle line shows the median, the box frame the 25th and 75th percentiles. Whiskers show all extreme values, not considered outliers (not displayed).

(68-100%) gathers the chunks where the ants move mostly towards the cue and are hence considered to be engaged in attraction behaviour. The boundary is moved down from the original 75% to 68% in order to increase the sample sizes of the top activity level. The sample sizes are displayed at the top of Figure 4.7. For all further analysis, I only consider the chunks in the 68-100% level.

Close-loop Oscillation

Despite the unnatural close-loop setting, the ants stick to their innate oscillation behaviour while attempting to approach the cue. Figure 4.8 shows, on the example of a chunk in the no-offset condition, a typical oscillation in the close-loop. The heading angle corresponds to the general coordinate system of the VR setting, it is however completely arbitrary for the ant, since the only visual stimulus is always just in front of her. A qualitative inspection of the heading angles along the paths shows no difference between the different offset conditions.

Compared to the paths in the open-loop (Fig. 4.3), the spread of angles is much higher. In the open-loop the fixations are limited to a ± 40 deg range by definition, here the distances covered by the ants in the close-loop are significantly higher. This is reflected in significantly higher amplitudes (KS-T with Bonferroni correction: open-loop vs. close loop for all offsets $p < 0.0125$; all offsets vs. all offsets: $p > 0.01$). Figure 4.9B shows the distribution of amplitudes for the different offsets. Here, the two different directions (towards and away from the cue centre) for each offset value are pooled together.

In the close-loop, the amplitudes are around four times as long as in the open-loop (close-loop 0: mean (mn)=40 deg, 10: mn=43 deg, 20: mn=46 deg, 30: mn=39 deg, 40: mn=55 deg; open-loop: mn=13 deg). The amplitude remains constant in the 0, ± 10 , ± 20 and ± 30 conditions, it

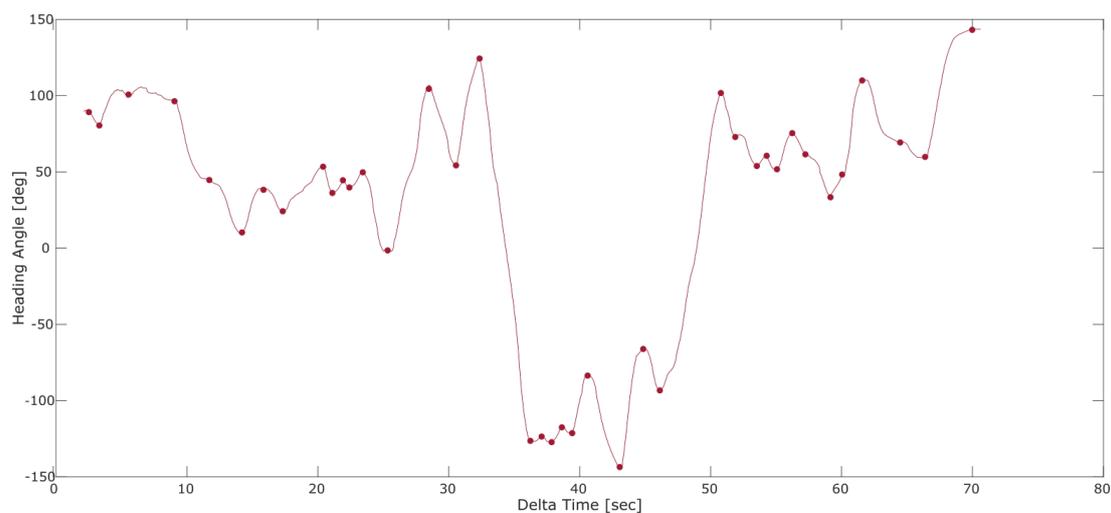


Figure 4.8.: Close-loop oscillation. Example oscillation from ant 5 chunk 7 in the condition with no-offset. Path smoothed for representation. Dots show turning points, which are extracted for further analysis.

increases however at an offset of ± 40 . The frequency behaves anti-proportionally (Fig. 4.9A). The distribution is not very spread, however it is approx. four times smaller than in the open-loop (KS-T with Bonferroni correction: open-loop vs. close loop for all offsets $p < 0.0125$; all offsets vs. all offsets: $p > 0.01$; close-loop 0: $mn = 0.5$ Hz, 10: $mn = 0.4$ Hz, 20: $mn = 0.3$ Hz, 30: $mn = 0.4$ Hz, 40: $mn = 0.4$ Hz; open-loop: $md = 1.6$ Hz).

While in close-loop, the ants still oscillate when attracted to a conspicuous visual cue. Without the visual feedback of moving past the cue, the amplitude of the oscillations quadruples. The underlying frequency of the oscillations drops anti-proportionally, i.e. longer oscillations require more time. The angular velocity of the oscillations should thus remain constant across the open- and close-loop.

Angular Velocity in the Close-Loop

As expected, the general angular velocity remains constant in the open- and close-loop (Fig. 4.10). We have a mean angular velocity of 14 deg per second (compared to 13 deg/sec. in the open-loop).

In the close-loop, there is an increase in speed at a 40 deg offset. It corresponds to an increase in amplitude with a constant frequency. There is also a slight drop in speed in the away 20 deg close-loop distribution. This correlates with an increased spread in amplitudes at this offset. Given the relative low sample size of this condition, this could be attributed to noise in the data.

4. Results

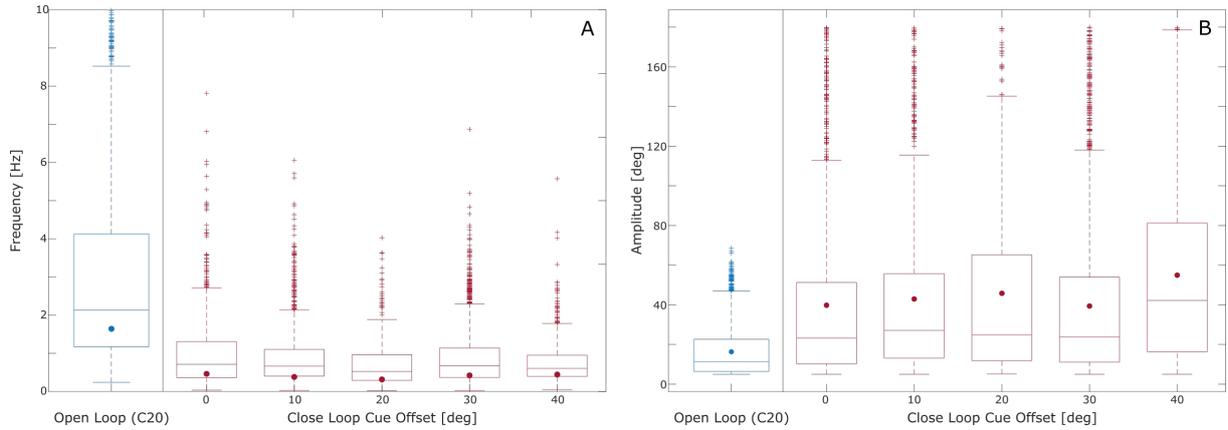


Figure 4.9.: Frequency (A) and amplitude (B) in the open and close-loop. Box plots show distribution of values for all chunks in the respective condition. For the close-loop (red) data separated by offset (plus and minus pulled together). For comparison, open-loop (blue) in the corresponding C20 condition. The middle line shows the median, the box frame the 25th and 75th percentiles. Whiskers show all extreme values, not considered outliers (crosses). Filled dot represents the mean value of each group.

Despite the simplicity of the visual stimuli we present to the ants in these two sets of experiments, the ants display some complex mechanisms of navigation. Ants use the edges of the visual cue as fixation points during attraction. Furthermore, they are able to adapt their oscillation mechanisms to the close-loop condition, where they are deprived of parallax information. This offers us new insights into their underlying neuronal mechanisms. To gather more insights into their behaviour on the treadmill, we expand the complexity of the visual stimuli.

4.2. Validating the Display of Natural Images on the VR Setting

The first part of the results dealt with validating the integration of treadmill and VR setting in the two loop conditions. The system works well as a whole and the ants are able to express simple innate behaviour while in the experimental setting. In this section I focus on the VR setting. I evaluate the feasibility to display more complex patterns and natural images. Furthermore, I research how ants interact with more complex natural and artificial scenes. As described in section 3.1.2, the VR setting consists of three screens arranged around the treadmill in an equilateral triangle. The system takes panoramic images as input and distorts them to accommodate them to the screen geometry. The image distortion as well as visual features of the setting, like the screen edges, introduce some noise into the VR setting. I evaluate to what degree this noise constraints experiments on the VR system. Furthermore, I address a simple behavioural question, do ants as a population have a preferred direction in

4. Results

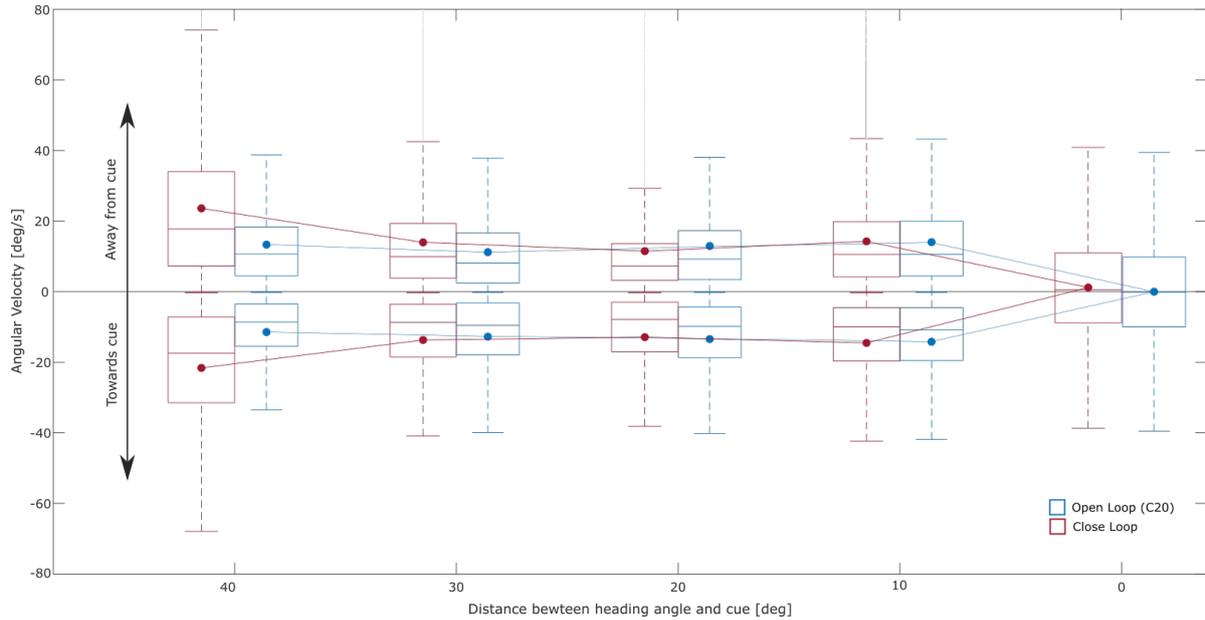


Figure 4.10.: Angular velocity in the close and open-loop. Box plots show the distribution of angular velocities in relation to the distance of the ant from the cue centre. For the close-loop (red) the distance is equal to the offset between heading direction and cue centre. In the open-loop (blue), each box plot bins ± 2.5 deg. The middle line shows the median, the box frame the 25th and 75th percentiles. Whiskers show all extreme values, not considered outliers (not displayed). The dots show the corresponding mean for each group. Bottom half of the plot shows velocities while approaching the cue, top half while moving away from it.

a complex natural scene?

4.2.1. Display of Complex Artificial Patterns

We expand the visual stimuli to include a whole range of black bars with different widths. Figure 4.11A shows the geometry of the pattern. Using this pattern, we test 55 different ants in three rotations of the visual stimulus: (A) the four thin bars aligned at the centre of screen one (Fig. 4.11A), (B) pattern from A rotated by 120 deg such that the four thin bars are at the centre of screen two and (C) pattern from A rotated by +90 deg such that the thin bars fall into a corner between two screens. Figure 4.11C illustrates the perspective of the ants with a panoramic picture from inside the VR chamber for condition C.

Different to the experiments introduced in section 4.1, here we are not only interested in the ants' innate behaviour but on the behavioural consistency of the individual ants and population. Hence, some ants are tested multiple times. 20 ants are tested, alternating conditions A and B for up to 4 times. Between each trial, the marked ants are returned to their nest for approx. 3 hours. The amount of ants decreased during the experiment as some

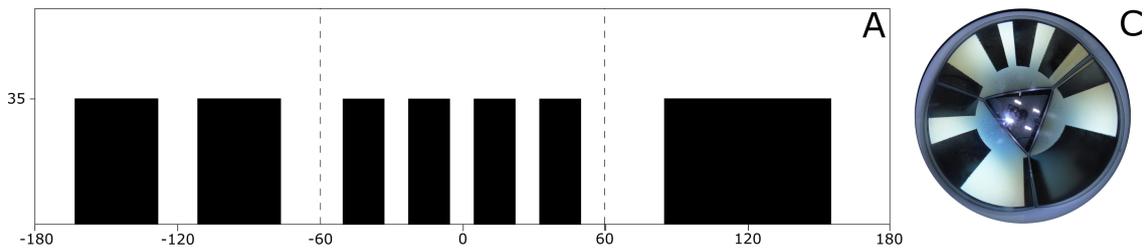


Figure 4.11.: Complex Artificial Image. The visual stimulus is composed of three groups of bars with different widths. Three rotations of the image are used, condition A as in (A), condition B as a +120 deg rotation of A and condition C as a -90 deg rotation of A (C). (A) Planar pattern as by design. (C) Panoramic image from inside the experimental chamber.

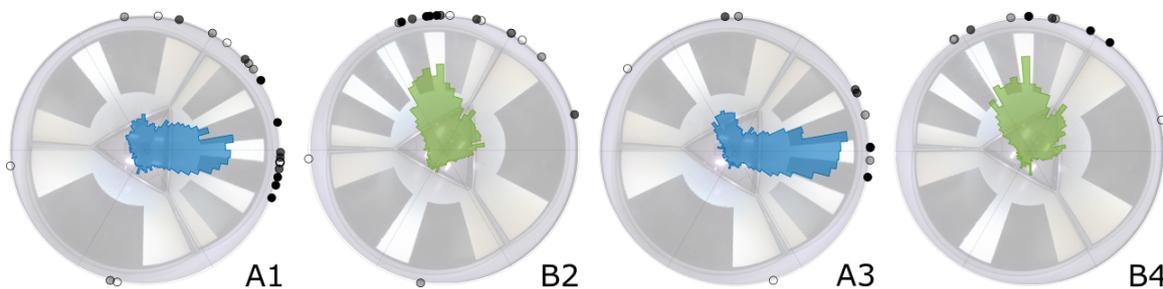


Figure 4.12.: Heading distributions against complex artificial pattern. The same ant is tested up to 4 times with alternating conditions A and B. B is rotated by 120 deg in respect to A. Background images show the corresponding orientation of the pattern. Circular histograms show the probability distribution of heading directions with a bin width of 5 deg. The circles show the mean heading direction for each ant in the group. The circle fill encodes the variance in quartiles, fully-filled = low variance, no-fill = high variance.

individuals did not emerge again from the nest (A1: $n=20$, B2: $n=18$, A3: $n=12$, B4: $n=11$). Each trial lasts a maximum of 3 min. Figure 4.12 shows the development in heading direction for the entire group.

The ants' paths in all four groups are directed (RT: A1, B2, A3 and B4 $p < 0.05$), and all four distributions have a mean direction towards the centre of the screen displaying the four thin bars (VT: A1, B2, and A3 $p < 0.05$, B4 $p < 0.05$). The four distributions are not significantly different from each other (KT with Bonferroni correction A1-B4 vs A1-B4 $p > 0.017$ for all combinations). In all repetitions, the ants seem to have a clear attraction for the pattern with the four thin bars, regardless of the orientation of the visual pattern.

To investigate if the attraction to the patterns is learned during this experiment or innate in the population, we test additional naive ants with condition B for the first time (B1: $n=16$) and with condition C for the first time (C1: $n=18$). Figure 4.13 shows the distribution of heading directions in conditions A, B, C where all the ants are exposed for the first time to the setting.

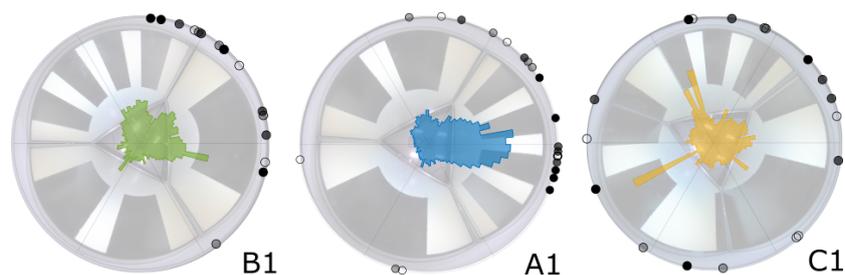


Figure 4.13.: Heading distributions against complex artificial pattern. Naive ants are tested for the first time in one of 3 orientations of the same pattern. With A as default, B is rotated by +120 deg and C is rotated by +90 deg. Background images show the corresponding rotation of the pattern. Circular histograms show the probability distribution of heading directions with a bin width of 5 deg. The circles show the mean heading direction for each ant in the group. The circle fill encodes the variance in quartiles, fully-filled = low variance, no-fill = high variance.

From the three distributions, only A1 is directed (RT: B1 $p > 0.05$, A1 $p \ll 0.05$, C1 $p > 0.05$). While in condition A, the population seems to have a great interest for the four thin bars, in B and C there seems to be no preferred direction.

Although not directed, the ants in B1 do show some interest in the thin bars as well as the edge between screens one and two. There are also some ants interested in the single wide bar, something not seen again in other conditions. In condition C1 the distribution is quite uniform around the circle. The results in Figure 4.13 suggest that the attraction for the thin bars in conditions B2, A3 and B4 might be learned from the first exposure during A1. Furthermore, a preferred direction might be more consistent for individual foragers as for the entire population.

If we evaluate the progression of individual ants in the A-B-A-B sequence, certain behaviour patterns emerge:

- Out of the 20 tested ants, only 7 showed a clear attraction towards the four thin bars on their first trial. Out of these 7, 3 did not emerge again, and the remaining four showed a learned attraction towards the four bars in subsequent trials.
- A group of 6 ants showed no clear attraction to any feature in the patterns on their first trial (2 of them also in the second), however in subsequent trials each of them focused on the four thin bars pattern. It seems as if the ants first explored the panorama and learned the "preferred" pattern for subsequent trials.
- 6 of the ants showed no preference in a pattern, but rather have a preferred external direction regardless of the rotation of the pattern. Out of the 6, 2 kept this attraction in at least three trials, 2 abandoned it after the second trial and then attracted towards the

four thin bars pattern and 2 did not record for more than two trials.

- 1 ant showed no preferred direction in any of the four trials.

The general trend seems to suggest that the first trial is not very definitive. Although many of the ants recognize the four thin bars as the preferred direction immediately, most either explore the whole panorama first or keep an external preferred direction at the beginning. Once the four bars have been acknowledged as a preferred direction, the ant will keep heading there in all subsequent trials. Most of the ants that did not do so in the beginning, eventually recognize and head towards the four bars. Hence, the four thin bars pattern seems to be an advantageous direction to follow as a population, and most ants need some exposure before recognizing it.

Ants show learning behaviour of a complex visual pattern in the treadmill. For one, this validates the effectivity of our system. With increasing exposure, ants finally settle on a preferred heading direction as a population. To better understand this behaviour, we expose the ants to a woodland scene and explore their behaviour while on a simulation of their natural habitat.

4.2.2. Display of Complex Natural Images

In chapter 3 we concluded that most artificial navigation systems fail to navigate in complex, outdoor, dynamic environments. This is exactly the kind of environment in which wood ants thrive as expert navigators. To investigate the visual ecology of the ants in their natural habitat, the VR setting needs to be able to robustly display panoramic images from the woods. In this section I evaluate the display of a complex woods scene on the VR setting.

The panoramic image was taken using a fish-eye camera in the same habitat where the ants kept in the lab were previously collected (Fig. 3.3). We selected the image for its texture richness, i.e. multiple distinct elements. Roughly, the selected scene has three distinct features: a small hill, a wide hill and a tree. The image is transformed to black and white and distorted to fit the screen geometry. We exposed the ants to three orientations of the image: (A) original, (B) image A rotated by +90 deg and (C) image A rotated by -90 deg.

Image Distortion in the VR system

The panoramic image was taken in the woods with a fish-eye lens. In order to display it in the VR setting, the original image is distorted to accommodate the geometry of the setting. Hence, a feature that falls into the edge between two of the screens has to be scaled up so that, from an ant's perspective, it seems as if it were closer to the ant and not far in the corner. Furthermore, in the unwrapping process of the fish-eye pictures, black pixels have to be introduced to fill the screens. In order to evaluate the distortion of the VR setting, we

took a panoramic picture form inside the VR chamber for each of the rotations. This is an approximation of what the ant sees during the recordings.

I take the pictures from inside the VR chamber, unwrap them and convert to black and white. I then calculate the similarity between the images using a rotation Image Difference Function (rIDF): image A is rotated every 10 deg and the pixel wise root-mean-square difference towards each of images A, B and C is calculated. Figure 4.14 shows the unwrapped VR chamber pictures and the image difference function between A and the respective rotation.

The difference towards image A at a 0 deg rotation of Image A is by design 0 (Fig. 4.14A).

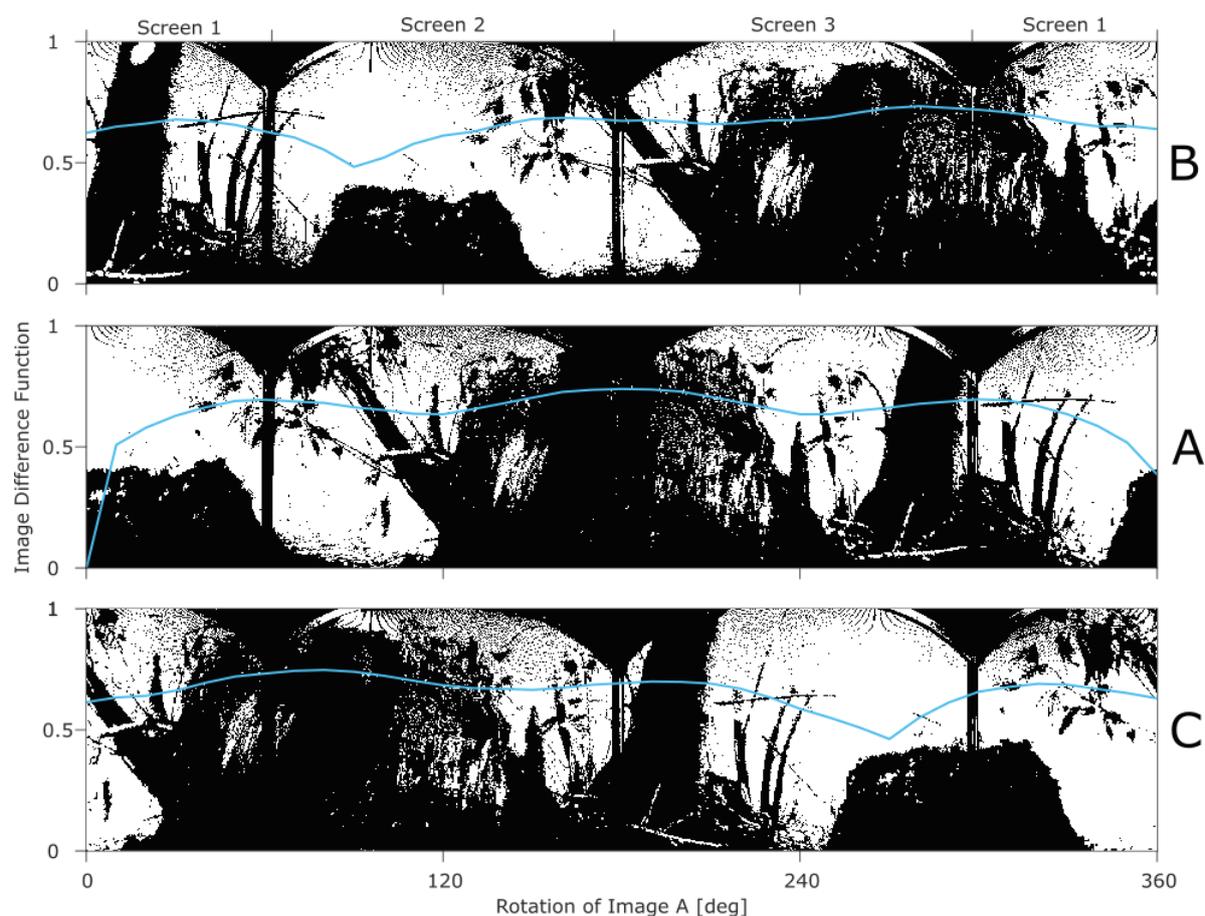


Figure 4.14.: Image difference function of the woodland image. (A) Original image, (B) image A rotated by +90 deg and (C) image A rotated by -90 deg. Difference function (blue): image A is rotated every 10 deg and the pixel wise root-mean-square difference towards each of images A, B and C is calculated.

The rIDF of A against A lets us evaluate how distinct the elements in the chosen woods picture are. The function in A fluctuates around 0.7 and has no distinct minima, hence the features of the image are quite distinct. In B and C, we see that the IDF has a clear minimum

at the edge of the small hill. If the display of the images on the VR screens were perfect, the function should reach 0 at these positions, instead it only drops to approx. 0.5.

The image distortion and the added setting features like the screen edges introduce quite some noise into the VR setting. Nevertheless, the same position on the panorama is clearly identifiable in both rotations inside the VR chamber.

Ant Preferred Direction in the Complex Natural Environment

To evaluate the display of the woods image, we test a total of 56 ants in the 3 different rotations of the image. 21 ants are tested alternating conditions A and B for up to 5 times. Between each trial, the marked ants are returned to their nest for approx. 3 hours. Trial five happened one weekend after four. The amount of ants decreased during the experiment as some individuals did not emerge back from the nest (A1: n=21, B2: n=20, A3: n=18, B4: n=15, A5: n=13). Each trial lasted a maximum of 3 min. Figure 4.15 shows the development in heading direction for the entire group.

In general, there seems to be an attraction for one of the elements, the big hill, as a visual

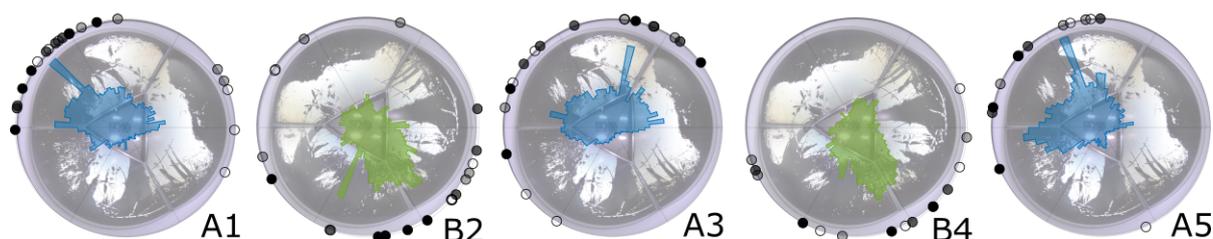


Figure 4.15.: Heading distributions against panoramic woods image. The same ant is tested up to 5 times with alternating conditions A and B. B is rotated by 90 deg in respect to A. Background images show the corresponding rotation of the image. Circular histograms show the probability distribution of heading directions with a bin width of 5 deg. The circles show the mean heading direction for each ant in the group. The circle fill encodes the variance in quartiles, fully-filled = low variance, no-fill = high variance.

feature. The distribution of heading directions is quite constant in subsequent trials within the same rotation. Nevertheless, the fixation point varies slightly between A and B. All five conditions are directed (RT: A1 $p \ll 0.05$, B2 $p < 0.05$, A3 $p < 0.05$, B4 $p < 0.05$, A5 $p \ll 0.05$). In all the conditions, the ants seem to be directed towards one of the corners between the screens. All but one of the distributions is directed towards the corner over which the big hill falls (VT for 180 deg: A1 $p < 0.05$, A3 $p = 0.45$, A5 $p \ll 0.05$; VT for -60 deg: B2 $p \ll 0.05$, B4 $p \ll 0.05$). The distribution in A3 breaks the pattern, it is, however, more spread and directed towards the other corner (VT for -60 deg with Bonferroni correction: A3 $p < 0.025$).

There seems to be an attraction towards the wide hill in the natural image. This is probably the most conspicuous element of the environment, it almost resembles the black rectangle we

used in previous experiments. As we established in section 4.1.2, ants do not fixate the centre of such a rectangle, but fixate on the edges. In this case, the most conspicuous edge to fixate is not the side edges of the wide hill, but the corner of the screens which falls within the visual element. The ants seem to clearly prefer the wide hill as a visual stimulus and either (i) the impulse to fixate an edge is stronger than to be consistent with learned behaviour or (ii) they do not see the slight difference in the position of the edge relative to the wide hill and hence believe they are being consistent.

To control for learned behaviour, we test additional naive ants with condition B for the first time (B1: $n=20$) and with condition C for the first time (C1: $n=15$). Figure 4.13 shows the distribution of heading directions in conditions A, B, C where all the ants are exposed for the first time to the setting. All three distributions are directed (RT: B1 $p<0.05$, A1 $p\ll 0.05$, C1 $p<0.05$). They are also all directed towards the corner that falls over the wide hill (VT for -60 deg: B1 $p<0.05$, VT for 180 deg: A1 $p\ll 0.05$, VT for 60 deg: C1 $p<0.05$).

At a population level, the ants seem to be naively attracted to the edge in the wide hill. The spread of the distributions falls with subsequent trials, suggesting that the ants learn the visual feature of a strong edge inside a wide hill. On subsequent trials, they attempt to match the learned pattern and probably do not realize there is a slight offset between the screen corner and the wide hill.

We validate the system ability to simulate the natural habitat of the ants. Here we see the

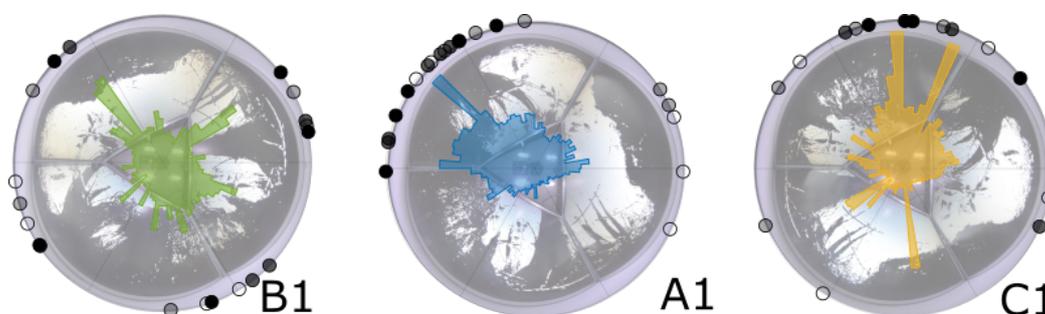


Figure 4.16.: Heading distributions against panoramic woods image. Naive ants are tested for the first time in one of 3 rotations of the same pattern. With A as basis, B is rotated by $+90$ deg and C is rotated by -90 deg. Background images show the corresponding rotation of the pattern. Histograms show the probability distribution of heading directions with a bin width of 5 deg. The circles show the mean heading direction for each ant in the group. The circle fill encodes the variance in quartiles, fully-filled = low variance, no-fill = high variance.

same tendency towards a preferred direction as we saw with the artificial pattern. The corners between the screen appear to slightly limit our display of panoramas, given their edge like character and the ants' attraction towards them.

The analysis of four experiments with increasingly complex visual stimuli allows us to gather valuable insights on the limitations of our experimental setting. We expose some complex navigation mechanisms employed by ants to navigate different conditions and scenarios. In the next chapter, I summarize these results and offer some explanations for the observed behaviours.

5. Discussion

Navigation is an essential behaviour any truly autonomous agent needs to master [24]. Approximately thirty years of research in robot navigation has explored multiple approaches to provide artificial agents with spatial cognition [23, 11, 25]. Although these algorithms have come far, robot navigation still fails to perform well in cluttered, dynamic, light variant environments [2, 27].

It is in these sceneries where wood ants thrive as expert navigators. Despite their tiny brains and low sensory resolution, ant foragers are capable of extraordinary navigation feats [1]. Contrary to the state-of-the-art approach for robot navigation, vSLAM, ant navigation does not require the use of maps or the extraction of features from the visual panorama [28, 35]. To better understand the visual ecology of wood ants and the underlying neuronal mechanisms, a new research methodology is required.

Developments in computer vision and graphics have opened up virtual reality (VR) as a novel technique to study insect behaviour [21]. VR offers the experimenter the ability to simulate the natural environment of the study subject, while having complete control over the visual stimuli the animal is exposed to. Furthermore, VR allows us to invert the loop that naturally exists between the environment and an animal's behaviour, which in turn opens the possibility to study the behaviour to brain loop of the animal in real time [48].

Recently, Goulard et al. [6] initiated the development of such a system. Different to previous approaches, where the insects are tethered [4, 3], the proposed setup combines a 3DoF motion compensation treadmill for walking insects [5], with a VR system. The success of the treadmill was already documented [6], however the integrated treadmill and VR systems have yet to be validated for the study of insect navigation.

This thesis deals with the integration of the VR and treadmill system. We designed a series of behavioural experiments to test the limitation of the new experimental setup, and gather new knowledge about mechanisms of visual navigation in ants. In this chapter, I gather the results and discuss possible implications. In the first part, I evaluate the experimental setup and explore new features for future development. In the second part, I reflect upon the observed ant behaviour and what this means for future research.

5.1. Evaluation of the Virtual Reality Experimental Setting

On a general scale, the proposed experimental setup works remarkably well as a novel technique to study ant navigation. We designed four experiments of increasing complexity, recorded a total of 498 ants and evaluated more than 50 hours of recordings. During the experiments, we saw the wood ants express some of the same natural behaviours that have been documented in the past: innate landmark attraction [45], oscillatory movement towards a goal [46], edge fixation [47] and pattern learning [33]. A more in depth discussion on these behaviours is offered in section 5.2.

That being said, we also observed certain shortcomings and limitations of the experimental setting, mainly:

- The treadmill controller is imperfect, causing unwanted oscillations and vibrations of the trackball during the recording.
- The corners between the screens act as unwanted visual features in the VR setting.
- The system does not support the simulation of a translational open-loop. Hence, the natural behaviour of the ant can only be expressed to a certain limit.
- The ant tracking has a ± 180 degree uncertainty, which has to be thoroughly compensated by the experimenter and data analysis.

In this section I address these limitations, evaluate them within the greater context and offer some modifications.

5.1.1. Hardware Limitations

Motor Control

One of the main sources of noise in the experimental setting is the treadmill itself. The trackball often vibrates and oscillates inadequately. Unwanted movements of the trackball can be followed back to two causes: (i) tracking failure and (ii) control. The first one shall be addressed in the Software Limitations (Sec. 5.1.2).

The treadmill is controlled in a loop. The x and y position of the ant are determined by the tracking subsystem and used to calculate a counteracting set of rotation vectors for the trackball. These are then transformed into three motor signals, which will in turn move the ant back to the centre of the setup. The new position is tracked, and the loop is closed. Two PD controllers serve this purpose. The respective gains were optimized by the experimenter using a trial and fail strategy over a range of sensible values. One possible tweak would be to revisit these parameters in an attempt to reduce the noise.

A second option involves expanding the controllers to include an integral component. The

oscillations of the trackball are most likely due to a lack of damping, and this can be addressed with a PID controller. Additionally, the system could be expanded to include an optical encoder on each of the motors. The additional sensory data would increase the control accuracy independently of the tracking subsystem and compensate for the PID with a second, shorter feedback loop.

Although the treadmill introduces some noise into the data, I present and validate a methodology to filter out most of this noise (Sec. 4.1.1). I use the motor angular velocities to identify moments of high ant activity. The damping and vibration problem described above is stronger when the ant stops moving, hence filtering for high ant activity lets us get away with the most undisturbed chunks of data. This methodology is tested on two independent datasets and proves to be a valuable tool to extract ant behaviour.

Unintentional Visual Stimuli

One of the first steps I take, is to control for any setup or room stimuli which could be influencing the ant. The results of the open loop experiment (Fig. 4.1 control) show none such clear effects. Nevertheless, pictures from inside the VR setting show that the lights are not as diffuse as we hoped for, and could potentially be aiding the ants in their orientation. A simple fix would be to replace the four individual lamps by a ring light around the high-speed camera.

A much bigger source of noise is the geometry of the screens. Each of the screens is framed, and at the corners where two screens meet, the frames build an edge that interrupts the image display. The effect this has on the ants is made clear in the natural scene experiment, where the edges prove to be strong enough to influence the ants' heading direction (Sec. 4.2.2).

One possible concern is the size of the angle between the two screens. Right now, the screens are arranged in an equilateral triangle and the angle is acute (60 deg). An ideal VR setting would have no edges at all, but a cylindrical screen around the trackball. This could be approximated by increasing the number of screens, the more screens, the greater the angle between them. Although this seems like a logical approximation to make, it could also have negative side effects. The distortion of the images to accommodate the screen geometry has no visible effect in our results (Figs. 3.1C, and 4.14), i.e. the manipulations performed by the VR setting on the visual display, already quite accurately approximate a cylindrical display. However, increasing the amount of screens would also increase the number of conspicuous edges between them. These, on the other hand, I show to influence the ants (Sec. 4.2.2). A better investment would hence involve frameless screens to remove any external, conspicuous edges.

5.1.2. Software Limitations

Translational Open-Loop

Roughly, an ant can translate, i.e. change position but keep orientation, or rotate, i.e. change orientation while on the same location. Open-loop is the natural state between an ant and its environment, i.e. any movement the ant makes, will influence her perception of the environment. If she rotates, the environment rotates in an opposite direction, if she translates, the objects in the direction of translation grow. By design, the treadmill constraints all translation movements but keeps rotations. This means that an ant on the treadmill is in translation close-loop, i.e. when she moves towards the black bar, this one stays the same size. Although this is clearly an unnatural state for the ant, the repercussions on our experiments are not so obvious. For one, in the natural open-loop, there are two situations where a visual stimulus does not change size upon translation. The first one involves big elements that are very far away and can clearly be segmented from the skyline, e.g. a mountain does not change size in the ants' immediate perception. The second involves a very big, textureless element that is very close to the ant and cannot be segmented from the skyline, i.e. the ant is so close she only sees "black" growing.

Hence, all the visual stimuli we present to the ant while in translational close-loop, are probably perceived as either very big, and far or very big and close objects. This implicates, that the ant does not "notice" she is in an unnatural situation but is rather expressing her natural behaviour towards very far and very close features. In fact, this might provide an explanation on why the ants in the natural scene experiment (Sec. 4.2.2) preferred to head towards the "wide hill". This is the element on the image closest to the screen top, and, given the translational close-loop, could be perceived as the closest feature of the panorama (very big very close) and thus the preferred feature to explore.

To further study the perception effect of close and far away objects, our system needs to be able to simulate the translational open-loop. Luckily, the constraints of the treadmill can be compensated by the VR. That is, a future version of the system should scale the size of a visual cue displayed on the VR screen depending on the forward speed (trackball speed) in that direction.

Mismatch in Ant Heading Direction

As detailed in chapter 3.1.2, the ant is tracked by fitting an ellipse around the contour of a white blob (against a black background) that most likely represents the ant. This method is quite robust to determine the position of the ant, however does not support the differentiation between the head and the tail of the ant. Once the system has a heading direction, it minimizes the change between frames to keep it constant. However, an initial failure to correct the mismatch would have catastrophic effects on the results and their interpretation.

At the beginning of a recording, there is a 50% chance that the head tail match is incorrect. Therefore, the experimenter needs to manually fix this by looking through the webcam. A failed flip by the experimenter, or a flip to far into the trial, are unlikely, yet additional sources of noise into our dataset. Furthermore, in around 10% of the trials, ants provoke an accidental flip of the tracking. This is often due to a sudden rotation or because the ants stand on their hind legs and the tracking loses track of their heading direction.

In the experiments presented in this thesis, I overcame these limitations in the data analysis, at the expense of losing some valuable data. A future version of the system might address this issue in real time. It could include a more robust algorithm to estimate the heading direction of the ant given the moving direction of the trackball, i.e. the head is most likely to be in the direction the ant is moving forward in.

Despite its current limitations, the treadmill and VR system is a great and novel tool to study the visual ecology and navigation strategies of ants. In the four experiments we perform, we were able to collect valuable new insights into the mechanisms of visual navigation in ants.

5.2. Wood Ant Behaviour While on the Experimental Setting

We use the novel treadmill and VR system to study visual navigation in ants while they are exposed to increasingly complex visual stimuli. First, we test the ants' innate response to a simple black rectangle. We compare the effects between rotational open- and close-loop. We then increase the complexity and present an artificial visual pattern of rectangles in different widths. Here, we explore the preferred direction on a population and individual level. Finally, we expose the ants to a natural woodland scene. Across all four datasets, certain behavioural trends emerge. In this section I explore their implications.

We chose to study ants because of their extensive (shared) navigational toolkit, impressive spatial cognition and learning abilities [1]. Furthermore, ants are central place foragers and as such we have additional access to their motivations, behaviours and neuronal mechanisms. Extensive research has shown insects to be innately attracted to black bars [45]. We take advantage of this attraction to study the navigation strategies of wood ants. Although plenty of research with trackballs and VR has been done with fruit flies [21] and even ants [4, 3], our integrated motion compensation treadmill and VR system opens new research possibilities. For one, this system allows for the first time to study untethered walking insects.

5.2.1. Open- and Close-Loop Oscillations

Plenty of research has documented oscillations in the paths of ants. Oscillations provide the insects with a robust visuomotor control strategy [49, 46]. With each turn, the ant can extract parallax information and regulate its heading direction towards a goal. In our experiments, we observe these oscillations during periods when the ants show a clear attraction towards a conspicuous visual landmark. Furthermore, I document the difference between open- and close-loop navigation on the oscillations' morphology.

Stuck in The Close Loop

The angular speed of oscillations remains relatively constant in both the open- and close-loop. While on open-loop, we document higher angular speeds when the cues are wider. This is consistent with previous results that show ants to regulate the speed of their oscillation by approximating the magnitude of the required turn [49, 47]. Furthermore, we observe that an ant moving towards a visual cue slows down when the cue first appears at an angle of (± 40 deg), only to then increase its speed as it rotates towards it.

This effect is marked strongest in the close-loop experiment, where the cue is constantly kept at an offset of (± 40 deg). Here the ants show significantly higher speeds compared to the open-loop and to other offsets. A possible explanation is, that the ants are stuck with the cue on the position in their visual field, that triggers a constant neuronal response to speed up. Despite the "hacked" neuronal response, the ants do not just follow the cue around the arena in circles, but the underlying frequency of their oscillations kicks in. Hence, the increased speed only causes higher amplitudes in their natural oscillation.

This is a good example of the potential of VR as an experimental methodology. In this experiment, we close the environment to behaviour loop and are able to draw conclusions about the behaviour to neuronal circuit implications. To draw more conclusions around this, a future experiment might record the switch between open- and close-loop. Here, the ants would navigate in the open loop until they fixate a visual cue, upon which the system switches to close-loop.

Amplitude and Frequency of Oscillations

As detailed above, the angular velocity does not change between open- and close-loop (other than in the ± 40 deg offset). Nevertheless, the magnitude of the oscillations between these two conditions changes by a factor of 4: the amplitudes in the close-loop are four times higher and the frequency is four times smaller. Given the anti-proportionality between these two variables, one of them is causing the other. Is the change frequency or amplitude driven?

Most likely, the change is driven by a change towards larger amplitudes. As discussed above, ants use oscillations to regulate their heading direction. They set up their angular speed

before each turn in response to the approx. distance toward the visual cue [49, 47]. As they pass the cue, they rely on parallax information to regulate the amplitude of the oscillation. In the close-loop, they are deprived of this feedback, hence the factor four in amplitude. Longer amplitudes require more time, hence the drop in frequency.

Oscillations appear to be a fundamental element of ant navigation. For one, they are a powerful tool to regulate the heading direction and extract parallax information. Ants show oscillation behaviour, even when in rotational close-loop, where they are deprived of sensory feedback.

In engineering terms, ant oscillations are a simple yet elegant technique for active sampling. Active sampling is a paramount activity any navigation system needs to engage in. The efficiency of an agent's sampling method will determine the quality of the sensory data and directly impact the computation power required to analyse it. Most technical applications differentiate between a sampling and a motion activity (or subsystem). Yet, I show how the ants' oscillation behaviour is deeply embedded into the agent's natural motion. Further research into the underlying frequencies and amplitudes of these oscillations will shed some light into possible imitations of this tight relationship between sampling and motion for robot navigation.

5.2.2. Edge Fixation

Edges play a decisive role in both artificial and natural navigation [23]. For one, they are easy to extract features of an image. Even strategies that do not rely on feature extraction greatly benefit from the high contrast of edges. Ants have been shown to use edges as a strategy to align and match retinotopic views [50, 47].

Edge Switching

In our experiments, we see the effect of edges multiple times. In the open-loop experiment, ants alternate between fixating the left and right edge of the simple visual cue. This is an interesting approach. Ants have been shown to navigate parallel to edges as a strategy for economizing visual recognition along routes [50], however the advantage of keeping track of both edges is not obvious.

Our recordings show, that the frequency of change between edges is relatively high (up to two times in 5 sec) and that the impulse to change is related to the distance between edges (the higher the distance the higher to need to change). A possible explanation is, that this is a compromise between edge extraction and centre of mass calculations [47]. The edges are easy to keep track off and provide the strongest parallax information to regulate the oscillations, and by alternating between them, the path will balance towards the centre of mass of the landmark.

In other to draw more meaningful concussions, a few additional conditions of the experiment would be useful, namely a 30 deg wide cue and a cue where the distance between the edges is grater than >40 deg, which is the angle at which I showed first signs of attraction.

Edge Density

In the experiment with the artificial pattern, the ants show a preferred direction towards the four thin bars pattern (Sec. 4.2.1). The density of edges is highest at this position, and this could cause this pattern to be the most attractive to follow. For one, this is the position in the panorama with the highest density of visual information, and hence the easiest to match upon re-encounter. In fact, we observe that the interest towards the four bars is learned, i.e. the ants do not always prefer them at first, however most of them eventually do, and once they have "selected" the four bars as a preferred direction, they stick to it.

It is possible that the ants relate the four thin bars to the shape of trees, a common visual feature in their natural habitat. This however raises the question on why, while exposed to a natural scene, the ants go for the wide hill and completely ignore the real tree in the image.

Edge Focus

Finally, we also saw a clear attraction towards the unintentional edges formed by the corners between the screens. In the natural image experiment (Sec. 4.2.2) there is a clear preference towards the wide hill visual feature in all the rotations. I theorize that the translational close-loop manipulation of perspective could make the wide hill appear as the closest feature in the panorama, and hence the easiest to explore first.

The hill is however quite wide (approx. 130 deg), and the exact focus point within the hill changes based on the orientation of the image within the setup. I show there to be a correlation towards the corners between the screens, i.e. the ants' heading direction is directed towards the corner that falls within the wide hill. I already established the benefits of edges in terms of information richness. It is possible then, that the corner only serves as an aid to better keep track of the hill, and the ants do not really notice a difference in the relative position of the corner to the hill when exposed multiple times to the image in changing orientations. A fourth condition in this experiment, where the wide hill falls exactly within two of the corners, might help further explore their influence.

In general, edges seem to have a strong influence in the strategies undertaken by the ants. This is not surprising given their richness in information, and is consistent with the findings of years of computer vision research [51]. Most computer science approaches have however focused on the active extraction of these features. Ant research has not found such active extraction of features [47]. Instead, ants benefit from edges through an indirect use of their information richness.

On a neuronal level, edge attraction indicates towards a control loop with plasticity. Here, the ants' behaviour is innately regulated by the presence of an edge-like visual stimuli. Hence, the response is "automatic". Furthermore, we see evidence of a distributed control system. Multiple edges put parallel control loops into conflict, and complex behaviours emerge as a result of their resolution. Robotic navigation systems have a lot to learn about these neuronal architectures. For one, edges are such important sources of visual information, that edge recognition should allocate special hardware resources and free up costly software calculations. Combined with a distributed control system, robots could benefit from the information richness of edges without the need for complex extraction and matching algorithms.

5.3. Conclusion

We use virtual reality as a novel technique to study the visual ecology and navigation of wood ants. Four research objectives are pursued:

Firstly, we successfully integrate the VR and treadmill system. This combination allows us to test untethered ants in complex VR sceneries. We are also able to test ants in rotational close loop. Secondly, we tested the system extensively and authenticated its functionality as a novel technique to study ant navigation. A careful analysis of the limitations of the system serves us as scaffolding for future development.

Thirdly, we gained useful insights into the navigation strategies of ants. Ant oscillation appears to be an important strategy in their navigation, both in open- and close-loop. We also see the strong influence of edges as an indirect visual feature rich in information. Both behaviours have important implications for the design of robot navigation architectures. Finally, we are able to expose ants to complex VR scenes. We see a learned preference for a specific direction in the panorama.

We shed some light into all four research objectives and provide a strong foundation for the future research of ant navigation through virtual reality.

A. Appendix

A.1. Data and Scripts

All collected data, as well as all the python and MATLAB scripts used, can be found at <https://app.box.com/folder/136369815441>. For access, please contact the administrator.

```
Treadmill 2021
├── Data
│   ├── artificial_scenes_prefered_direction_aspd_[95] ....artificial pattern dataset
│   ├── close_loop_offset_compare_cloc_[164] .....close-loop dataset
│   ├── natural_scenes_prefered_direction_nsdp_[126] ..... natural image dataset
│   ├── open_loop_fixed_cue_olfc_[113] .....open-loop dataset
│   └── Data Structure and File Nomenclature.txt
├── Treadmill and VR Scripts
│   ├── V5_07.06.2021
│   │   ├── Master_Offset.py ..... main VR script
│   │   └── TestWidget_tkinter.py .....main treadmill script
├── Matlab Scripts
│   ├── Res_ASPD ..... main artifical pattern script
│   ├── Res_CLOC ..... main close-loop script
│   ├── Res_NSPD ..... main natural scene script
│   ├── Res_OLFC ..... main open-loop script
│   ├── image_difference ..... IDF script
│   ├── CircStat2012a ..... circular statistics toolbox [43]
│   ├── support_functions .....secondary analysis functions
│   ├── plot_functions .....plot functions
│   └── image_functions ..... image transformations
├── Natural Images
│   ├── ASPD_Images ..... artificial patterns
│   └── NSPD_Images ..... natural images
└── Lab Book
```

List of Figures

2.1. Overview vision-based navigation methodologies. Adapted from [25].	5
3.1. Experimental Setting. (A) Information flow between systems: treadmill subsystem (blue) and VR subsystem (orange). The black arrow symbolizes communication in the close-loop condition. (B) Top view. VR system is composed of three screens arranged in an equilateral triangle around the trackball. (C) Panoramic picture from inside the VR chamber. The horizontal lines of the displayed visual elements are distorted to accommodate the system's geometry and appear as projected equidistantly from the ants' perspective. (D) Picture of the experimental setup.	15
3.2. Wood ant. Image shows a forager of the species <i>Formica rufa</i> standing on the trackball while facing a black bar as a unique visual cue in an otherwise white environment.	18
3.3. Natural woodland image. (A) Raw image in original rotation, (B) image from A rotated by +90 deg and transformed to black and white. (C) Picture from inside the VR chamber with condition C, which is image A rotated by -90 deg.	20
3.4. Activity chunks extraction process. Motor angular velocity in blue, ant angular velocity in grey. (A) Raw motor angular velocity in x direction ($\dot{\Theta}_x$). (B) Absolute sum of $\dot{\Theta}_x$ and $\dot{\Theta}_y$. Thresholds used to sort data in next step. (C) Motor activity divided in four levels according to amplitude. Density of the signal is encoded by the average amplitude in a 2 sec window. (D) Contour of the smoothed motor activity. Data is divided into chunks within which the motor activity does not reach zero. (E) Motor data (same as B) with the extracted chunks. The integral under the contour of each chunk is calculated and divided into one of four quartiles of activity. (F) Angular velocity of the ant with the chunk divisions and labels calculated on the base of motor activity.	22

4.1. Polar histograms of heading directions in response to a simple black cue. The width of the cue is 20 deg (blue) and 40 deg (orange) respectively. The cue is displayed in either of the three screens (S1, S2, S3) or not at all (S0). Rectangles on the edge of each histogram show the position of the respective cue. The circular histograms show the probability distribution for the entire group, bin size 5 deg. The circles show the mean heading direction for each ant in the group. The circle shade of colour encodes the variance in quartiles, fully-filled = low variance, no-fill = high variance.	26
4.2. Polar histograms of heading directions divided by activity levels. Each row shows the heading data within a level of activity. The width of the visual cue is 20 deg (blue) and 40 deg (orange) respectively. The cue is displayed on either of the three screens (S1, S2, S3) or not at all (S0). The angle of the centre of each screen is marked in the top left plot. The polar histograms show the probability distribution for the entire group, bin size 5 deg. The circles show the mean heading direction for each ant in the group. The circle shade of colour encodes the variance in quartiles, fully-filled = low variance, no-fill = high variance. . .	28
4.3. Heading direction during a fixation. The plot shows two examples of a fixation, one for the C20 condition (blue, ant 7 fixation 11) and one for C40 (orange, ant 75 fixation 10). Paths smoothed for representation. Dots represent the turning points. Frequency and amplitude are calculated between two neighbouring turns. Each segment is categorized by its direction to the cue centre: towards the cue (black arrows) or away from the cue (red arrows). Dashed lines show the mean for each path respectively. Since the orange path crosses the cue centre, two mean values are calculated. Coloured bars on the left show the position of the cue.	30
4.4. Distribution of weighted mean headings. C20 in blue and C40 in orange. Bin size 1.5 deg. For each fixation, either two or one mean heading values are calculated, depending on the ant crossing the cue centre for more than 30% of the path or not. Each mean value is weighted by the amount of data points it represents.	31
4.5. Frequency (A) and amplitude (B) of the fixations. Box plots show the distribution of values for both conditions C20 (blue) and C40 (orange). The middle line shows the median, the box frame the 25th and 75th percentiles. Whiskers show all extreme values not considered outliers (crosses). Filled dots represent the mean values.	32

4.6.	Angular velocity during periods of fixation. Box plots show the distribution of angular velocities in relation to the distance of the ant from the cue centre. Each box plot gathers the velocities in ± 5 deg range. The middle line shows the median, the box frame the 25th and 75th percentiles. Whiskers show all extreme values, not considered outliers (not displayed). The dots show the corresponding mean for each group. Bottom half of the plot shows velocities while approaching the cue, top half while moving away from it. Condition C20 in blue and C40 in orange.	33
4.7.	Angular velocity for each offset condition and activity level. Box plots show the distribution of angular speeds for each offset of the cue. Each plot stands for one level of activity: 0-50%, 50-68% and 68-100%. Sample sizes displayed on top. The middle line shows the median, the box frame the 25th and 75th percentiles. Whiskers show all extreme values, not considered outliers (not displayed).	34
4.8.	Close-loop oscillation. Example oscillation from ant 5 chunk 7 in the condition with no-offset. Path smoothed for representation. Dots show turning points, which are extracted for further analysis.	35
4.9.	Frequency (A) and amplitude (B) in the open and close-loop. Box plots show distribution of values for all chunks in the respective condition. For the close-loop (red) data separated by offset (plus and minus pulled together). For comparison, open-loop (blue) in the corresponding C20 condition. The middle line shows the median, the box frame the 25th and 75th percentiles. Whiskers show all extreme values, not considered outliers (crosses). Filled dot represents the mean value of each group.	36
4.10.	Angular velocity in the close and open-loop. Box plots show the distribution of angular velocities in relation to the distance of the ant from the cue centre. For the close-loop (red) the distance is equal to the offset between heading direction and cue centre. In the open-loop (blue), each box plot bins ± 2.5 deg. The middle line shows the median, the box frame the 25th and 75th percentiles. Whiskers show all extreme values, not considered outliers (not displayed). The dots show the corresponding mean for each group. Bottom half of the plot shows velocities while approaching the cue, top half while moving away from it.	37
4.11.	Complex Artificial Image. The visual stimulus is composed of three groups of bars with different widths. Three rotations of the image are used, condition A as in (A), condition B as a +120 deg rotation of A and condition C as a -90 deg rotation of A (C). (A) Planar pattern as by design. (C) Panoramic image from inside the experimental chamber.	38

- 4.12. Heading distributions against complex artificial pattern. The same ant is tested up to 4 times with alternating conditions A and B. B is rotated by 120 deg in respect to A. Background images show the corresponding orientation of the pattern. Circular histograms show the probability distribution of heading directions with a bin width of 5 deg. The circles show the mean heading direction for each ant in the group. The circle fill encodes the variance in quartiles, fully-filled = low variance, no-fill = high variance. 38
- 4.13. Heading distributions against complex artificial pattern. Naive ants are tested for the first time in one of 3 orientations of the same pattern. With A as default, B is rotated by +120 deg and C is rotated by +90 deg. Background images show the corresponding rotation of the pattern. Circular histograms show the probability distribution of heading directions with a bin width of 5 deg. The circles show the mean heading direction for each ant in the group. The circle fill encodes the variance in quartiles, fully-filled = low variance, no-fill = high variance. 39
- 4.14. Image difference function of the woodland image. (A) Original image, (B) image A rotated by +90 deg and (C) image A rotated by -90 deg. Difference function (blue): image A is rotated every 10 deg and the pixel wise root-mean-square difference towards each of images A, B and C is calculated. 41
- 4.15. Heading distributions against panoramic woods image. The same ant is tested up to 5 times with alternating conditions A and B. B is rotated by 90 deg in respect to A. Background images show the corresponding rotation of the image. Circular histograms show the probability distribution of heading directions with a bin width of 5 deg. The circles show the mean heading direction for each ant in the group. The circle fill encodes the variance in quartiles, fully-filled = low variance, no-fill = high variance. 42
- 4.16. Heading distributions against panoramic woods image. Naive ants are tested for the first time in one of 3 rotations of the same pattern. With A as basis, B is rotated by +90 deg and C is rotated by -90 deg. Background images show the corresponding rotation of the pattern. Histograms show the probability distribution of heading directions with a bin width of 5 deg. The circles show the mean heading direction for each ant in the group. The circle fill encodes the variance in quartiles, fully-filled = low variance, no-fill = high variance. . . 43

List of Tables

- 2.1. Overview of the reviewed algorithms. 12
- 3.1. Overview of recording days and sample sizes for all four datasets. 19
- 4.1. Chunk sample sizes divided by activity level and conditions. Values for C20 and C40 before and after the slash respectively (C20/C40). 27

Bibliography

- [1] P. Graham and A. Philippides. “Vision for navigation: What can we learn from ants?” In: *Arthropod Structure & Development* 46.5 (2017), pp. 718–722. DOI: 10.1016/j.asd.2017.07.001.
- [2] A. Li, X. Ruan, J. Huang, X. Zhu, and F. Wang. “Review of vision-based Simultaneous Localization and Mapping”. In: *2019 IEEE 3rd Information Technology 2019*, pp. 117–123.
- [3] Z. Kócsi, T. Murray, H. Dahmen, A. Narendra, and J. Zeil. “The Antarium: A Reconstructed Visual Reality Device for Ant Navigation Research”. In: *Frontiers in Behavioral Neuroscience* 14 (2020), p. 599374. DOI: 10.3389/fnbeh.2020.599374.
- [4] P. Schultheiss, A. Buatois, A. Avarguès-Weber, and M. Giurfa. “Using virtual reality to study visual performances of honeybees”. In: *Frontiers in Behavioral Neuroscience* 24 (2017), pp. 43–50. DOI: 10.1016/j.cois.2017.08.003.
- [5] S. Shigaki, S. Fukushima, D. Kurabayashi, T. Sakurai, and R. Kanzaki. “A novel method for full locomotion compensation of an untethered walking insect”. In: *Bioinspiration & Biomimetics* 12.1 (2016), p. 016005. DOI: 10.1088/1748-3190/12/1/016005.
- [6] R. Goulard, C. Bühlmann, J. E. Niven, P. Graham, and B. Webb. “A motion compensation treadmill for untethered wood ants (*Formica rufa*): evidence for transfer of orientation memories from free-walking training”. In: *Journal of Experimental Biology* 223.Pt 24 (2020). DOI: 10.1242/jeb.228601.
- [7] D.-E. Nilsson. “The evolution of eyes and visually guided behaviour”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 364.1531 (2009), pp. 2833–2847. DOI: 10.1098/rstb.2009.0083.
- [8] D. Miller and R. Stegmann. “The Eye of the Eagle”. In: *European Journal of Implant and Refractive Surgery* 3.1 (1991), pp. 71–73. DOI: 10.1016/S0955-3681(13)80155-4.
- [9] M. V. Srinivasan. “Visual control of navigation in insects and its relevance for robotics”. In: *Current Opinion in Neurobiology* 21.4 (2011), pp. 535–543. DOI: 10.1016/j.conb.2011.05.020.
- [10] F. Huttmacher. “Why Is There So Much More Research on Vision Than on Any Other Sensory Modality?” In: *Frontiers in Psychology* 10 (2019), p. 2246. DOI: 10.3389/fpsyg.2019.02246.

- [11] F. Bonin-Font, A. Ortiz, and G. Oliver. “Visual Navigation for Mobile Robots: A Survey”. In: *Journal of Intelligent and Robotic Systems* 53.3 (2008), pp. 263–296. DOI: 10.1007/s10846-008-9235-4.
- [12] F. Zeng, C. Wang, and S. S. Ge. “A Survey on Visual Navigation for Artificial Agents With Deep Reinforcement Learning”. In: *IEEE Access* 8 (2020), pp. 135426–135442. DOI: 10.1109/ACCESS.2020.3011438.
- [13] Y. D. V. Yasuda, L. E. G. Martins, and F. A. M. Cappabianco. “Autonomous Visual Navigation for Mobile Robots: A Systematic Literature Review”. In: *ACM Computing Surveys* 53.1 (2020), pp. 1–34. DOI: 10.1145/3368961.
- [14] R. Wehner, B. Michel, and P. Antonsen. “Visual navigation in insects: coupling of egocentric and geocentric information”. In: *Journal of Experimental Biology* 199.1 (1996), pp. 129–140. DOI: 10.1242/jeb.199.1.129.
- [15] C. Bühlmann, M. Mangan, and P. Graham. “Multimodal interactions in insect navigation”. In: *Animal Cognition* 23.6 (2020), pp. 1129–1141. DOI: 10.1007/s10071-020-01383-2.
- [16] T. S. Collett, P. Graham, R. A. Harris, and N. Hempel-de-Ibarra. “Navigational Memories in Ants and Bees: Memory Retrieval When Selecting and Following Routes”. In: *Advances in the Study of Behavior*. Vol. 36. Academic Press, 2006, pp. 123–172.
- [17] M. Kohler and R. Wehner. “Idiosyncratic route-based memories in desert ants, *Melophorus bagoti*: How do they interact with path-integration vectors?” In: *Neurobiology of Learning and Memory* 83.1 (2005), pp. 1–12. DOI: 10.1016/j.nlm.2004.05.011.
- [18] M. Mangan and B. Webb. “Spontaneous formation of multiple routes in individual desert ants (*Cataglyphis velox*)”. In: *Behavioral Ecology* 23.5 (2012), pp. 944–954. DOI: 10.1093/beheco/ars051.
- [19] J. Zeil. “Visual homing: An insect perspective”. In: *Current Opinion in Neurobiology* 22.2 (2012), pp. 285–293. DOI: 10.1016/j.conb.2011.12.008.
- [20] B. Webb. “Chapter 1 Using Robots to Understand Animal Behavior”. In: *Advances in the Study of Behaviour*. Vol. 38, pp. 1–58.
- [21] S. N. Fry, N. Rohrseitz, A. D. Straw, and M. H. Dickinson. “TrackFly: virtual reality for a behavioral system analysis in free-flying fruit flies”. In: *Journal of Neuroscience Methods* 171.1 (2008), pp. 110–117. DOI: 10.1016/j.jneumeth.2008.02.016.
- [22] T. Stone, M. Mangan, A. Wystrach, and B. Webb. “Rotation invariant visual processing for spatial memory in insects”. In: *Interface focus* 8.4 (2018), p. 20180010. DOI: 10.1098/rsfs.2018.0010.

- [23] G. N. DeSouza and A. C. Kak. "Vision for Mobile Robot Navigation: A Survey". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24.2 (2002), pp. 237–267. doi: 10.1109/34.982903.
- [24] M. S. Güzel. "Autonomous Vehicle Navigation Using Vision and Mapless Strategies: A Survey". In: *Advances in Mechanical Engineering* 5 (2013), p. 234747. doi: 10.1155/2013/234747.
- [25] Y. Lu, Z. Xue, G.-S. Xia, and L. Zhang. "A survey on vision-based UAV navigation". In: *Geo-spatial Information Science* 21.1 (2018), pp. 21–32. doi: 10.1080/10095020.2017.1420509.
- [26] T. Taketomi, H. Uchiyama, and S. Ikeda. "Visual SLAM algorithms: a survey from 2010 to 2016". In: *IPSP Transactions on Computer Vision and Applications* 9.1 (2017). doi: 10.1186/s41074-017-0027-2.
- [27] J. K. Makhubela, T. Zuva, and O. Y. Agunbiade. "A Review on Vision Simultaneous Localization and Mapping (VSLAM)". In: *2018 International Conference on Intelligent and Innovative Computing Applications (ICONIC)*. IEEE, 122018, pp. 1–5. ISBN: 978-1-5386-6477-3. doi: 10.1109/ICONIC.2018.8601227.
- [28] A. Cheung, M. Collet, T. S. Collett, A. D. Dewar, F. Dyer, P. Graham, M. Mangan, A. Narendra, A. Philippides, W. Stürzl, B. Webb, A. Wystrach, and J. Zeil. "Still no convincing evidence for cognitive map use by honeybees". In: *Proceedings of the National Academy of Sciences* 111.42 (2014), E4396–E4397. doi: 10.1073/pnas.1413581111.
- [29] J. Santos-Victor, G. Sandini, F. Curotto, and S. Garibaldi. "Divergent stereo for robot navigation: learning from bees". In: *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*. 1993, pp. 434–439. ISBN: 1063-6919. doi: 10.1109/CVPR.1993.341094.
- [30] G. C. H. E. de Croon, C. De Wagter, and T. Seidl. "Enhancing optical-flow-based control by learning visual appearance cues for flying robots". In: *Nature Machine Intelligence* 3.1 (2021), pp. 33–41. doi: 10.1038/s42256-020-00279-7.
- [31] A. Vardy and R. Möller. "Biologically plausible visual homing methods based on optical flow techniques". In: *Connection Science* 17.1-2 (2005), pp. 47–89. doi: 10.1080/09540090500140958.
- [32] D. G. Lowe. "Distinctive Image Features from Scale-Invariant Keypoints". In: *International Journal of Computer Vision* 60.2 (2004), pp. 91–110. doi: 10.1023/B:VISI.0000029664.99615.94.
- [33] B. A. Cartwright and T. S. Collett. "Landmark maps for honeybees". In: *Biological Cybernetics* 57.1-2 (1987), pp. 85–93. doi: 10.1007/BF00318718.

- [34] M. O. Franz, B. Schölkopf, H. A. Mallot, and H. H. Bülthoff. “Where did I take that snapshot? Scene-based homing by image matching”. In: *Biological Cybernetics* 79.3 (1998), pp. 191–202. DOI: 10.1007/s004220050470.
- [35] J. Zeil, M. I. Hofmann, and J. S. Chahl. “Catchment areas of panoramic snapshots in outdoor scenes”. In: *Journal of the Optical Society of America. A* 20.3 (2003), pp. 450–469. DOI: 10.1364/JOSAA.20.000450.
- [36] W. Stürzl and J. Zeil. “Depth, contrast and view-based homing in outdoor scenes”. In: *Biological Cybernetics* 96.5 (2007), pp. 519–531. DOI: 10.1007/s00422-007-0147-3.
- [37] B. Baddeley, P. Graham, P. Husbands, and A. Philippides. “A model of ant route navigation driven by scene familiarity”. In: *PLOS Computational Biology* 8.1 (2012), e1002336. DOI: 10.1371/journal.pcbi.1002336.
- [38] P. B. Ardin, M. Mangan, and B. Webb. “Ant Homing Ability Is Not Diminished When Traveling Backwards”. In: *Frontiers in Behavioral Neuroscience* 10 (2016), p. 69. DOI: 10.3389/fnbeh.2016.00069.
- [39] P. Graham, A. Philippides, and B. Baddeley. “Animal cognition: multi-modal interactions in ant learning”. In: *Current Biology* 20.15 (2010), R639–40. DOI: 10.1016/j.cub.2010.06.018.
- [40] R. Goulard, C. Bühlmann, J. E. Niven, P. Graham, and B. Webb. (personal communication August 2021).
- [41] The Mathworks Inc. *MATLAB 2021a*. 2021. URL: <https://uk.mathworks.com/products/matlab.html>.
- [42] P. Manandhar. *Polar To/From Rectangular Transform of Images*. 2007. URL: <https://www.mathworks.com/matlabcentral/fileexchange/17933-polar-to-from-rectangular-transform-of-images>.
- [43] P. Berens. *CircStat: A Matlab Toolbox for Circular Statistics*. 2009. URL: <http://www.jstatsoft.org/v31/i10>.
- [44] E. Batschelet. *Circular Statistics in Biology*. London: Academic Press, 1981.
- [45] C. Bühlmann and P. Graham. “Innate visual attraction in wood ants is a hardwired behaviour seen across different motivational and ecological contexts”. In: *bioRxiv* (2021), p. 2021.01.29.428794. DOI: 10.1101/2021.01.29.428794.
- [46] D. D. Lent, P. Graham, and T. S. Collett. “Phase-dependent visual control of the zigzag paths of navigating wood ants”. In: *Current Biology* 23.23 (2013), pp. 2393–2399. DOI: 10.1016/j.cub.2013.10.014.
- [47] D. D. Lent, P. Graham, and T. S. Collett. “Visual Scene Perception in Navigating Wood Ants”. In: *Current Biology* 23.8 (2013), pp. 684–690. DOI: 10.1016/j.cub.2013.03.016.

- [48] J. Lim and T. Celikel. “Real-time contextual feedback for close-loop control of navigation”. In: *Journal of Neural Engineering* 16.6 (2019), p. 065001. doi: 10.1088/1741-2552/ab2ffa.
- [49] D. D. Lent, P. Graham, and T. S. Collett. “Image-matching during ant navigation occurs through saccade-like body turns controlled by learned visual features”. In: *Proceedings of the National Academy of Sciences* 107.37 (2010), pp. 16348–16353. doi: 10.1073/pnas.1006021107.
- [50] S. C. Pratt, S. E. Brooks, and N. R. Franks. “The Use of Edges in Visual Navigation by the Ant *Leptothorax albipennis*”. In: *Ethology* 107.12 (2001), pp. 1125–1136. doi: 10.1046/j.1439-0310.2001.00749.x.
- [51] D. R. Waghule and R. S. Ochawar. “Overview on Edge Detection Methods”. In: *2014 International Conference on Electronic Systems, Signal Processing and Computing Technologies*. IEEE, 1/9/2014 - 1/11/2014, pp. 151–155. ISBN: 978-1-4799-2102-7. doi: 10.1109/ICESC.2014.31.