

This is to certify that the thesis entitled

2007

HOW HONEY BEES USE VISUAL LANDMARKS DURING GOAL-DIRECTED NAVIGATION: WAYFINDING STRATEGIES AS SEQUENTIAL DECISION-MAKING PROCESSES

presented by

FRANCIS NORMAN BARTLETT III

has been accepted towards fulfillment of the requirements for the

Doctoral degree in Zoology Major Professor's Signature 2006 Date

MSU is an Affirmative Action/Equal Opportunity Institution



PLACE IN RETURN BOX to remove this checkout from your record. TO AVOID FINES return on or before date due. MAY BE RECALLED with earlier due date if requested.

DATE DUE	DATE DUE	DATE DUE

2/05 p:/CIRC/DateDue.indd-p.1

HOW HONEY BEES USE VISUAL LANDMARKS DURING GOAL-DIRECTED NAVIGATION: WAYFINDING STRATEGIES AS SEQUENTIAL DECISION-MAKING PROCESSES

By

Francis Norman Bartlett III

A DISSERTATION

Submitted to Michigan State University In partial fulfillment of the requirements For the degree of

DOCTOR OF PHILOSOPHY

Department of Zoology

ABSTRACT

HOW HONEY BEES USE VISUAL LANDMARKS DURING GOAL-DIRECTED NAVIGATION: WAYFINDING STRATEGIES AS SEQUENTIAL DECISION-MAKING PROCESSES

By

Francis N. Bartlett III

The impressive wayfinding abilities and unique communication system of foraging honey bees has made them an important model organism for the study of navigation for over seventy years. In this time, researchers have identified and quantified many of the navigational strategies these animals use to discover and exploit remote food sources in their environment. Computer scientists interested in designing intelligent mobile robots often look to natural systems like the honey bee for inspiration when developing the sensory systems and learning algorithms of their artificial agents. In this dissertation I examine several navigational strategies of honeybees and discuss how these behaviors might deal with two of the more difficult, overarching problems encountered by engineers attempting to build effective mobile robots: 1.) how the agent determines what perceptual information provided by the environment is reliable for decision-making and 2.) how the agent uses this information to guide its path through large, continuous realworld environments. Chapter 1 begins by briefly reviewing the components of the sequential decision making framework. Sequential decision-making is a machine learning approach used to address how robots and other artificial systems might deal with difficult real-world decision problems such as navigation. The remainder of chapter 1 applies this framework to what is currently known about insect navigation. The goal of this chapter is to advocate the integration of the sequential decision-making approach

with insect navigation. Chapter 2 experimentally examines how honey bees deal with the problem of perceptual aliasing. This problem, often encountered by mobile robots, arises when two different locations in the environment present identical perceptual information but require different actions to attain the goal. I present a series of experiments designed to determine if bees can use visual memory to solve an aliased y-maze task. The results suggest that honey bees do not use visual memory as a primary strategy for resolving perceptual ambiguity. Chapter 3 identifies and investigates the spontaneous attraction honey bees have to visual landmarks they encounter en-route to a goal. These experiments suggest that this attraction is based on the visual size of the en-route landmark. I discuss how honey bees may use this behavior to aid the selection of useful perceptual information as well as segment a route into a smaller more manageable series of decision points. Chapter 4 is also concerned with the how honey bees use en-route visual landmarks to divide their environment into discrete and manageable number of decision points. The data presented in this chapter address the use of local vectors by honey bees over long distances and characterize a previously unreported strategy associated with en-route landmarks: specifically that bees repeatedly return to the most recently experienced en-route landmark when they are unsuccessful in finding the goal location. Chapter 5 looks closely at the snapshot model of visual navigation which can be regarded as a mechanism for simplifying the task of generating flight control actions based on visual landmarks in large, complex environments. These experiments suggest that the use of visual beacons is the primary navigation strategy of honey bees while snapshot navigation is probably limited to only a small area very near the final goal.

Dedicated to my wife Trisha and my children Ciera, Alexis, Francis and Kenna.

ACKNOWLEDGEMENTS

I would first like to begin by thanking my advisor Fred Dyer for his patience, support enthusiasm throughout my graduate career at MSU. Thanks for sticking by me even when things looked bleak. I also give thanks to my committee; Heather Eisthen, Tom Getty and Erik Altmann for helping me pull the many conceptual threads of my work into a cohesive story.

I also express gratitude to my family; my kids and mother for being so supportive throughout my graduate career but especially to my wife Trisha who shouldered a great deal of our household responsibilities during the writing of this dissertation.

I have had the great fortune to be assisted by several graduate students over the years and I list them in alphabetical order; Lora Bramlett, Lauren Davenport, Jennifer Jones, Yoav Littman, Chris Speilburg and Kourtney Trudgen. Thank you so much for your hard work.

I would also like to thank my colleague Aaron Pearson for collaborating on several projects with me. I cannot remember having more fun in the field with anyone else. For helping me learn to program in MatLab® I give many thanks to both Mike Mack and Dirk Colbry. I would like to thank Mike Mack additionally for generating the elegant computer simulations of honey bee behavior presented in chapter 5 of this dissertaion. I would also like to thank Steven Fry and Rudiger Wehner for lending me their Trackit® system, a powerful tool that made data collection and analysis so much simpler. I also thank Steven Fry for his continued support with our tracking software as well as for great conversations about life as well as honey bee navigation. For advice, good laughs and

v

the occasional shoulder to cry on, I give thanks to the other past and present members of the bee laboratory; John Townsend-Mehler, Katie Wharton, Cindy Wei, Mathew Collett, Micah Gill, Dina Grayson and Kevin Guse.

Finally I am very grateful to the National Science Foundation IGERT program for their financial support without which none of this would have been possible.

TABLE OF CONTENTS

List of tables	x
List of figures	xi
Chapter 1	
Navigation as a sequential decision-making problem: insects as a model system	1
Introduction	1
Sequential decision-making	5
Basic framework	5
Exploration vs. Exploitation	8
Defining state	9
Some difficulties for applying the sequential decision-making framewo	rk
to animal behavior	. 16
Defining navigational state for honey bees	. 19
Path integration	. 21
Visual landmarks	. 21
Olfactory Cues	. 23
Motivational state	. 24
Exploratory behaviors and selective attention	. 25
Summary	. 29
Prospects for an integrative research program	. 31
Chapter 2 The use of visual memory by honeybees to solve a perceptually aliased maze task Introduction	. 33 . 33
Path integration	. 36
Time of day	. 36
Hunger state	.37
Trip order	.37
Visual and olfactor memory	.38
Methods	. 40
General training and testing procedures	. 40
Experimental variations	45
Condition 1	.45
Condition 2	45
Condition 3	. 46
Condition 4	46
Condition 5	. 46
Condition 6	. 47
Condition 7	. 48
Simple discrimination	. 48
Results	. 50

Overall	50
Dependence of second choices	
Discussion	55
Training methods	57
Odor control	
Indoors vs. outdoors	
Cylinder maze chambers vs. boxes	59
Baffles	59
A local strategy for solving the maze task	60
Puzzling results	
Conclusions	62

Chapter 3

How honeybees use an off-route landmarks near the goal	64
Introduction	64
Methods	67
Apparatus	67
Training and testing procedures	69
Data analysis	
Results	
Experiment 1	
Headings	
Path length	
Experiment 2	
Experiment 3	
Discussion	

Chapter 4

Honeybee use an en-route landmark as a resetting point for search flights	
Introduction	
Methods	
Apparatus	
Training and testing	
Data analysis	
Results	
Search distribution	
Survivorship analysis	
Flight direction	100
Discussion	101

Chapter 5

How honey bees use landmarks near the goal: an examination of the role of		
snapshot guidance	106	
Introduction	106	
Methods	111	
Apparatus	111	

Training and testing	112
Results	114
Experiment 1	114
Experiment 2	116
Approaching the goal	117
Landmark removals	118
Pushing bees off course	121
Array configuration	122
Body axis analysis	123
Experiment 3	126
Discussion	127
Support for the snapshot model	128
Beacons	130
Scanning flight and TBLs	131
Conclusion	132
References	134

List of Tables

Table 1-2. Condition variations across maze experiments49
Table 2-2. Choice Dependencies: The tables show the relationship of honey bee decisions on two consecutive trips to the same stimulus configuration in the three training conditions where such behavior was permitted
Table 5-1. Body axis headings near the goal. Angular means and confidence intervals of headings shown in Figure 5-8

LIST OF FIGURES

Figure 1-1. Agent/Environment interaction dialog
Figure 1-2. Hypothetical state assessment and action selection of a honey bee30
Figure 2-1 . <i>MSU maze apparatus with tunnel</i> : The approach tunnel was 100cm long by 22cm wide and covered in random Julesz pattern pixel size of 1 cm ² . The height of the apparatus was 22 cm throughout. Entrance holes, 4cm in diameter, were cut through the center of the stimuli. The tops of the maze chambers were covered with custom fitted Plexiglas® lids. Both horizontal and vertical grating patterns had a period of 20 degrees of visual angle when viewed from the center of the choice point
Figure 2-2. ANU cylinder maze apparatus: The tops of the maze chambers were covered with custom fitted Plexiglas® lids. Entrance and exit holes were 3 cm in diameter. All stimuli were 18 cm x 18 cm. Both horizontal and vertical grating patterns had a period of 18 degrees of visual angle when viewed from the center of the choice point
Figure 2-3 . <i>MSU maze apparatus without a tunnel</i> : The height of the apparatus was 22 cm throughout. Entrance holes 4cm in diameter were cut through the center of the stimuli. The tops of the maze chambers were covered with custom fitted Plexiglas® lids. Both horizontal and vertical grating patterns had a period of 20 degrees of visual angle when viewed from the center of the choice point
Figure 2-4. Simple discrimination over time: Percent correct choices based on the number of visits to the apparatus for color (solid line) and grating (dashed line) discrimination tasks. First bees we trained bees to one of two stimuli followed by a reversal stage where the other alternate stimulus predicted reward
Figure 2-5. Complex association performance across experimental conditions: Percent correct choices across bees for each condition. Numbers inside the bars are the sample sizes for each condition. Data for condition 6 was collecting at ANU. All other experimental took place at MSU
Figure 2-6 . Complex association performance of individual bees trained and tested at <i>ANU</i> : Percent correct decisions for each individual bee. Numbers inside the bars are the total number of decisions performed by each individual

Figure 3-5. Experiment 1 path length summary. A summary of path length changes when the training array was the panorama (open squares) and 2 cylinders (filled circles). Distances were transformed into multiples of the calculated straight path to the feeder hole to provide a common measure for the panoramic and 2-cylinder training groups. Error-bars show the standard error of each mean. The panorama and 2-cylinder treatments were found to be statistically different in only the 1st trip condition after conducting a Bonferoni correction for multiple comparisons.

Figure 3-7. Experiment 2 headings. The arrows within the semicircle represent the mean heading vector at 50 cm from the entrance **a**.) Training array **b**.) The added landmark was the same size as the training elements and placed far off route beyond the training array. **c**.) The added landmark was the same size as the training elements and placed far off route before the training array. **d**.) The added landmark was one half of the size of the training elements, placed nearer the goal path and before the training array. **81**

Figure 3-8. Experiment 3 angle summary a.) Training array which included an en-route landmark. b.) Flight headings of bees when the off-route training landmark was removed.

Figure 5-3 Model predictions and flight paths to a 3 landmark array. **a**.) Vector field predicted by the snapshot model (implemented in Matlab). This plot predicts the direction a bee should travel from each 10 cm² bin if she is using the template matching algorithm proposed by Cartwright & Collett (1984). **b**.) Actual flight tracks of the approaches of honeybees to the landmark array. Excessive looping behavior was removed in order to provide a clearer picture of approach trajectories. ET designates the location of the entrance tube and FH the feeder hole.

Our snapshot model was implemented in Matlab as follows:

Figure 5-5 Model prediction and flight paths of bees after the removal of the farthest array elements. a & c.) Vector field predicted by the snapshot model. These plots

Chapter 1

Navigation as a sequential decision-making problem: insects as a model system

Introduction

After observing a dance performed by a fellow forager, a honey bee leaves the hive for her first foraging flight of the day. Using the spatial information provided by the dance, she sets a course for a feeding place in the environment which may be as far as 10 kilometers from the hive. Upon arrival at her destination she begins a local search of the ground and elevated foliage for flowers containing nectar. If the bee successfully locates a high quality flower patch she roams from blossom to blossom collecting nectar without revisiting depleted flowers along the way. Once her honey-crop is filled with nectar she lifts off and flies directly back to the nest. Upon return, she drops off her nectar load and before leaving again she may perform a dance of her own to recruit additional hive-mates to the patch. This will be the first of approximately sixty such trips for a single bee in one day.

The problem faced by this bee is also faced by many other mobile organisms and autonomous robots that must travel through their environment in search of resources. During each foray into its environment the agent, whether natural or artificial, needs to determine for a multitude of locations, which behavioral action will most efficiently lead it to the desired goal. Given that many navigating agents at some initial point are naïve and that even relatively experienced individuals face environments that are prone to

change, how might they determine what behavior is most appropriate in each state of the environment they experience? At first glance, linking behaviors with sensory stimuli seems a straight forward problem. However, computer scientists who have investigated this problem have identified two particular features of real-world environments that can severely complicate navigation and other goal-directed sequential tasks.

The first problem arises because most natural environments have patchy resources where the actions necessary to attain the goal do not immediately precede the receipt of reward in time or space. In navigation tasks a mobile agent is usually required to make a series of decisions in a sequence of intermediate environmental locations before it achieves its goal (reviewed in Kaelbling et al., 1996; Mitchell, 1997; Sutton & Barto, 1998). In such cases the agent must somehow associate the sequence of events leading up to a delayed reward.

The second problem originates in the nature of the sensory stimuli found in real-world environments. For even the most simplistic sensory systems the natural environment presents a glut of data. Some of this sensory input may be useful for making decisions but much of it will be either redundant or irrelevant. If we assume that most agents have a limited amount of memory or information processing capability they will need a strategy for filtering useful cues from sensory clutter (McCallum, 1996). Even if an agent possesses sophisticated sensory equipment and an accurate and efficient sensory filter, many environments have locations where the immediately perceptible sensory cues are not informative enough to make a reliable decision. This can happen when a familiar environmental location changes radically between episodes or when multiple locations have very similar perceptual characteristics. In essence the agent can experience

environmental locations that are simultaneously cluttered with useless information *and* sparse with regard to useful cues (McCallum, 1996; Hasinoff, 2002; Shani & Brafman, 2004). In summary, the agent needs to learn which actions and environmental locations most effectively lead to the desired goal and associate unique sensory cues to these locations in order to make the appropriate decision on subsequent visits. Furthermore, the agent must make these complex associations without immediate feedback about the quality of its decisions.

Honey bees and other insects that forage over long distances face just these sorts of problems. Many decades of behavioral research on hymenoptera have illuminated a number of the strategies available to insects for addressing these issues. Much of this research has been dedicated to determining how insects use celestial cues and features of the terrain to guide their way through the environment (von Frisch, 1967; Wehner et al., 1996; Dyer, 1998; Collett & Collett, 2002 for reviews). These behavior studies, in conjunction with neurobiological evidence have also led to qualitative and quantitative hypothesis of how insects acquire (Cheng & Spetch, 1991; Lehrer, 1991, 1993; Zeil et al., 1996; Capaldi & Dyer, 1999) and employ (Anderson, 1977; Cartwright & Collett, 1983; Wehner & Rossel, 1983; Collett & Collett; 2004; Fry & Wehner 2005) this information during repeated trips through their environment. Unfortunately, these models are rarely organized into a common and cohesive framework. Instead, they focus on each strategy as if it operates independently rather than studying how the different strategies are integrated and coordinated by the organism (e.g. Giurfa & Menzel, 1997; Collett & Collett, 2002, 2004).

As biologists our goal is to unravel the mechanisms underlying complex behavior and to understand why organisms solve certain problems in their own particular way. In essence we attempt to reverse engineer well crafted biological machines that have been shaped by natural selection to solve specific problems. Computer scientists are faced with an inverse and perhaps more difficult problem. Their goal is to build, from the ground up, intelligent machines that resemble biological agents in being able to accomplish complex sequential tasks in real-world environments. Over the years engineers have looked to strategies employed by biological systems for inspiration when designing algorithms for complex tasks like navigation (reviewed by Franz, 2000). Biology benefits from these efforts because the formal models generated by computer scientists often provide explicit and testable hypotheses for how biological agents perform the complicated tasks they commonly face. In addition, an understanding of the problems faced by artificial systems may help sharpen questions about how organisms solve analogous problems.

The remainder of this chapter will keep with this tradition by applying a machine learning approach to what is currently known about insect navigation. In the following sections I will first define the components of the sequential decision-making framework developed by computer scientists to address complex serial tasks. This will include the nomenclature, algorithmic concepts and specific problems that arise therein. I will then discuss several strategies insects use for navigating through their environment and apply to them concepts from sequential decision-making. As I will show, it is not clear if this framework applies directly to the navigation behavior of biological agents. Nevertheless, its potential lies in how it encompasses a comprehensive set of problems that are very

similar to those experienced by navigating organisms. My goal is to explore this potential to see what can be gained by casting insect navigation in a sequential decisionmaking framework. This review will set the stage for later chapters which investigate particular components of the navigation problems faced by honey bees.

Sequential decision-making

Basic framework

Sequential decision-making addresses the general question of how autonomous agents that sense and act in their environments can learn to choose a sequence of optimal actions to achieve their goals. The sequential decision-making concepts I will describe and apply throughout this chapter come from several comprehensive reviews on the subject by Kaelbling et al. (1996), McCallum (1996), Mitchell (1997), and Sutton & Barto (1998). This approach starts by breaking the agent's world into three basic variables; (1) the set of perceptible *states* that define the disposition of the agent relative to its environment, (2) the set of possible behavioral *actions* that the agent can execute within each state and that take it to the next state, and (3) the *reward* values returned by the environment when the agent performs one of its actions at in some environmental state st. The dialog box below from Kaelbling et al. (1996) gives a reasonable idea of how these components interact.

Environment	You are in state 65.	You have 4	possible actions.
-------------	----------------------	------------	-------------------

Agent.....I'll take action 2.

Environment......You received a reinforcement of 7 units. You are now in state 36. You have 7 possible actions.

Agent.....I'll take action 5.

Environment......You received a reinforcement of 3 unit. You are now in state 4.

You have 5 possible actions.

Agent.....I'll take action 1.

Environment......You received a reinforcement of 10 units. You are now in state 87. You have 3 possible actions. etc.....

Figure 1-1. Agent/Environment interaction dialog

The goal of the agent is to learn a *policy* for controlling its behavior that in some way maximizes the amount of reward it receives.

The sequential decision-making framework allows the characteristics of states, actions and rewards to be flexible in a number of ways. This flexibility makes it possible to apply sequential decision-making learning algorithms to a wide variety of real problems and environments. For example, this framework does not require that all of the agent's actions to be eligible in every state of the environment. A navigating agent like a mobile robot should not turn left if there is wall obstructing its left side and it makes sense for the action -turn left to be ineligible in such circumstances. Also, many of the learning algorithms allow for actions to provide stochastic rewards and state changes. Probabilistic state change may be an issue for a mobile robot that experiences wheel slippage or an insect flying in heavy winds. It these cases the agent may attempt to execute the action "turn left" with considerable uncertainty of the actions outcome. Sequential decision-making algorithms are designed to accommodate the failure of actions and other probabilistic features like fluctuations in reward values associated with pairs of states and actions. Reward fluctuation is very real issue natural systems. For example, a foraging animal may discover a patch of resources that is initially very profitable. Over time this patch may be subject to depletion by competitors and the animal's own foraging visits. In such cases it might be beneficial for the agent to be able to reassess the value of familiar foraging locations over successive visits in order to track fluctuations in reward.

Another biologically relevant feature of the sequential decision-making framework is that the agent learns its policy by way of real-time experience within its environment. While acting in its environment the agent observes the reward and state changes associated with the different available actions in each environmental state. Over many episodes of experience in its environment (often thousands or more for robotic agents) the agent learns the payoffs for conducting each action in each state (Kaelbling et al., 1996; McCallum, 1996; Mitchell, 1997; Sutton & Barto, 1998). The optimal policy then selects for each environmental state the action with the highest associated reward. In navigational terms the actions with the highest values are usually those that bring the agent closer to some desirable physical location in its environment. In other sequential tasks like playing board games and optimizing manufacturing processes, the goal may be winning the game or completing the construction of some item (Tesauro, 1992; 1995).

Tasks like navigation and playing board games often require the agent to conduct a series of actions in a succession of states before receiving a reward. This common issue

of delayed rewards, which is often the norm in real-world tasks, is one of the specific problems that the sequential decision-making framework is designed to address. Using a variety of different methods, these algorithms reinforce the path to the goal by altering the reward values of previous actions based on their proximity to the final reward. In essence, each time the agent performs an action in a state, the reward it receives from this action is used to update reward values of the previously experienced pairs of states and actions that led it there. With enough experience, the agent eventually assigns the proper reward value to each pairing of state and action. One way that some sequential decision-making algorithms differ is in how much and how far back in time or space the current reward signal can reinforce previous states and actions. Some sequential decision-making algorithms update only the closest previous state-action pairs (for a more detailed review with explicit algorithmic examples see Kaelbling et al., 1996; Mitchell, 1997; Sutton & Barto, 1998).

Now that I have laid out the basic principles I would like to address some common difficulties agents encounter when confronted with sequential decision-making tasks and describe some of the solutions to these issues offered by computer scientists.

Exploration vs. exploitation

If the agent's ultimate goal, as mentioned above, is to learn a control policy that maximizes long-term reward. It should come to follow the path to the goal that its experience predicts is the most rewarding (e.g. shortest in time or space or requiring the least effort). The problem is that an agent may need to make decisions in the

environment before it has learned the optimal course of action. While learning about a particular environment, an agent faces a trade-off between acting on what it has experienced so far and exploring options yet to be discovered (Sutton & Barto, 1998). If the agent decides to exploit its environmental knowledge too soon it runs the risk of converging on a policy that is sub-optimal when compared to other unexplored alternatives. Conversely, if the agent chooses to continually select untried actions and visit unexplored states it stands to miss out on the rewards it knows are available. Reinforcement learning usually addresses this trade-off between exploration and exploitation by customizing the learning algorithm to the agent's environment (Wilson, 1996; Mitchell, 1997; Sutton & Barto, 1998; McMahon et al., 2005). In stable environments, a common strategy is to start the agent with a high probability of exploring that decreases with increasing environmental experience. This gives the agent a better chance to learn its environment before converging on an optimal exploitative policy. In environments where the reward and state structure are dynamic or stochastic it may be better to set a constant exploration probability to give the agent an opportunity to track environmental change (Mitchell, 1997; Sutton & Barto, 1998).

Defining state

In order for many sequential decision-making strategies to find an optimal policy, the states of the environment must possess three simplifying characteristics; the number of states is finite, the states are discrete, and that environment provides all of the information necessary for the agent to determine its current state and make a proper decision about what action to take next (Watkins & Dayan, 1992; Sutton & Barto, 1998). The previous

example of agent environment-interaction alludes to these properties (**Figure 1-1**). The dialog suggests that there is probably some <u>finite</u> number of <u>discrete</u> environmental states (designated by positive whole numbers in this case) and that the environmental signal tells the agent, without ambiguity, what state it is in. Notice that the environment in the dialog does not say "you receive a reward of 4 and you are *probably* in state 23 or maybe state 55" (Kaelbling et al., 1996).

Unfortunately, biological agents as well as autonomous mobile robots usually reside in environments that violate at least one, if not all of these assumptions. For instance, real-world environments are usually quite large as well as continuous which creates a number of complications. To begin with, large environments make storing a value for every state-action pairing impractical in terms of the agent's memory and processing power (Mitchell, 1997). Also, if the environment is large enough, the agent may not be able to visit every state. If the agent has not visited many of the states in its environment how can it really know what the optimal policy is? When the number of states large enough to exceed the agent's processing power or spatial accessibility the environment can be considered infinite for all practical purposes. Additionally, continuous state spaces cause problems because the agent can experience the perceptual features of known states from a nearly infinite number of different "views". This can cause even small environments to have huge state spaces if every view is treated as a separate state (Sutton & Barto, 1998).

Recognizing these problems, computer scientists have extended many sequential decision-making algorithms to handle these difficulties. One interesting method is for the agent to generalize its previous experience to novel states of the environment. Here, the

agent bases its action selection in new states on the optimal action for known states with similar perceptual characteristics (Tesauro, 1992, 1995; Mitchell, 1997). Such generalization techniques allow the agent to function when it encounters a new state in its current environment and sometimes to entirely new situations. Generalization algorithms can also be used to simplify continuous environments. Using a similar method the agent could learn the proper action for a single state and then apply this action to all of the perceptually similar neighboring states or views of states. Generalization algorithms are basically a way for the agent to convert an environment that is large or continuous, or both, into a manageable number of discrete states.

Transforming the agent's perceptual input into a simpler, more tractable data set is a common theme in sequential decision-making (Kaelbling et al., 1996; Mitchell, 1997; Sutton & Barto, 1998). In a similar way, honey bees appear to use something akin to a generalization algorithm when solving pattern discrimination tasks in y-maze environments. In these experiments bees are trained to find a nectar reward presented near one of two visual stimuli. When confronted with two novel visual patterns in later episodes, honey bees select the novel pattern that most closely resembles the rewarding pattern learned during training. A variety of experiments suggest that bees can generalize their actions to new stimuli based on perceptual similarities in color, shape, size and intensity (Ronacher & Duft, 1996; Stach et al., 2004). Another compelling example of generalization in honey bees comes from Zhang et al. (1996). In this experiment they trained bees through a maze via a fixed route to the food designated by a series colored cards. When the maze was reconfigured such that the color cards marked out a novel path, bees followed the cards rather than using the sequence of turns associated with the

reward in the original training configuration. This result suggests that during training bees learned to follow the color cards rather than a sequence of motor turns and that bees could apply a general card following strategy when dealing with novel trajectories through the maze environment.

So far I have discussed some strategies used by computer scientists and insects to deal with infinite and continuous states of the environment. Earlier in this chapter I also alluded to the third problem faced by navigating agents; that agents are not always able to effectively discern their current environmental state due to an overabundance of perceptual information or noisy and obscured sensory inputs. Under these circumstances the environment violates the third assumption known as the Markov property.

For an environment to be Markov the agent's current state must provide all the necessary perceptual information for it to choose the optimal action. More to the point, the correct action for an agent in some state $S_{(t)}$ cannot be dependent on the path through previous states leading to the current state $S_{(1...t)}$ (Kaelbling et al., 1996; Mitchell, 1997; Sutton & Barto, 1998). One can quickly get a feel for why this property is appealing. What if to take the correct action in its current state the agent has to remember the previous series of states or state-action pairings that led there? This is not a serious problem if one or a few environmental states require a memory of previous events to select the next action. On the other hand, if this is the situation for many or all states of the environment the memory and computational power necessary to retain and process this information would grow rapidly as the number of environmental states increases. It would be very costly to store in memory the complete sequence of events leading up to every state in even very restricted real-world environments.

Although computationally attractive, the Markov property, like the discrete and finite state assumptions discussed above, can be problematic in many real-world decisionmaking problems. Fortunately the environments in most real-world navigation tasks are inherently Markov. It is difficult to imagine a situation where an agent would need to modify its action selection based on the sequence of previously encountered states that led to its current position. Each decision made along the way to a goal state is usually based on the information provided in each intermediate state as it is encountered. For example, if I am at the corner of Oak Street and Sherman Boulevard the fact that I need to travel due west to reach the library is independent of whether I reached the intersection from the north or the south.

Even though most environments encountered during navigation tasks adequately fit the Markov assumption, the brief example above raises an important issue. Before an agent can choose the correct action in any given state, it must first be able to determine what state it is actually in. For instance, making a correct decision is more complicated in the example above if there are no street signs designating the corner of Oak St. and Sherman Blvd. I could use other landmarks such as houses or trees to define the intersection but if no useful cues are available, selecting the correct action becomes a difficult task. Usually if we look close enough we can find some information that uniquely defines each important decision point in our environment. Unfortunately this is not always an option for artificial agents utilizing noisy, low-resolution sensors such as infrared and sonar (McCallum, 1996; Shani & Brafman, 2004).

If an agent cannot resolve the sensory distinctions between similar places in the environment these states are said to be perceptually aliased. Perceptual aliasing occurs

when two states of the environment appear identical to an agent's perceptual system(s) but require different actions to optimally obtain the goal (Whitehead & Ballard, 1991). Aliased locations are non-Markov since the agent is unable make proper decisions using only the sensory information available at its current state. What is interesting is that in many cases the Markov violation results from the agent's restricted sensory and computational capacity rather than from ambiguities inherent in the environment.

When an agent encounters aliased states it can solve the ambiguity in at least two ways. One option, as described above, is to "look closer" with the perceptual system and try to uncover distinguishing characteristics of each aliased location. To return to the previous example, if no road signs are available I might commit to memory the color of one or more of the houses found at the corner of Oak and Sherman so as to distinguish it from other unmarked intersections. Another strategy is to use a memory of one or more unambiguous, previous states to aid decision-making in the aliased location. For example remembering that the intersection of Elm St. and Farragut Blvd. occurs immediately before Oak and Sherman and using this piece of information as part of my state definition at the aliased intersection.

The utile distinction algorithm is an example of how the sequential decision making approach uses a "look closer" strategy to resolve perceptual aliasing. In this strategy the agent learns in real-time to select the perceptual information that defines each state based on how useful the information is for predicting reward (McCallum, 1996). By learning which perceptual features of each state best distinguish it from other similar states the agent is able to convert perceptually aliased locations into tractable Markov states. There is some evidence that honey bees and wasps have special behaviors they perform to

selectively memorize important perceptual information in their environment (Lehrer, 1993; Zeil et al., 1996; Capaldi & Dyer, 1999). Perhaps like utile distinction algorithms, these strategies provide a means for insects to avoid storing unnecessary perceptual clutter.

Sometimes the critical distinctions necessary to resolve environmental ambiguities are below the resolution of the sensory system or grossly obfuscated by other sensory clutter or noise. In these cases searching for additional distinctions fails and it is usually not plausible to improve the sensory system of the agent because the sensory apparatus is developmentally and evolutionarily fixed (as is the case of natural agents) or because it is very expensive (as is often the case for artificial agents). Instead the agent may rely on a memory of previous states and use this information at the ambiguous choice point. Unfortunately this strategy is susceptible to the same issue that is avoided by the Markov assumption: as state space increases history increases and therefore memory capacity must increase.

A common strategy sequential decision-making algorithms use to resolve this problem is to collapse large amounts of information into a more manageable size (McCallum, 1996; Sutton & Barto, 1998). They try to reduce past experience into one or a few values that are updated after each transition from one state to another. This simplified parameter is part of the agents current state and summarizes information about past state-action pairs leading to the current choice point. Each state is therefore defined by a combination of the current sensory perception and an internal summary of previous experience. The addition of this information results in a Markov environment because the agent has all of

the information it needs to select the optimal action in each state. A nice example of this strategy found in nature is the path integration system of insects.

Path integration allows an insect to keep a running tally of its movements in relation to its starting point during its outbound and homeward paths. It is understood that insects continuously sum the distance and direction of their movements throughout their travel in a global path integration vector (reviewed in Collett & Collett, 2000). By updating its path integration vector after each action that generates motion, the insect maintains a compact summary of all Euclidean movement leading to its current position in the environment. If insects do have a notion of state, it is likely that the value of the path integration vector is a part of it. As is the case with path integration, state may include more than just the immediate perceptual information provided by the external environment. It may also include percepts internal to the agent such an internal measure of time or motivational states like hunger and sexual arousal.

Some difficulties for applying the sequential decision-making framework to animal behavior

The sequential decision-making framework provides a general formalization for carrying out a variety of sequential, goal-directed activities. Although this framework has been applied successfully to several machine learning tasks (see Sutton & Barto, 1998 for review), how well we can apply sequential decision-making ideas and algorithms to animal behavior and insect navigation in particular remains unclear. In this section I will address some of the complications of integrating the sequential decision-

making approach with animal behavior and discuss why it is worthwhile venture even with such difficulties.

As discussed previously, many computer science applications of sequential decisionmaking algorithms explicitly pre-define the perceptual components of state. Under these circumstances the agent's sensors and the environment are designed such that the agent always has access to the necessary information to make an appropriate decision at any location in its environment (Kaelbling et al., 1996; Mitchell, 1997; Sutton & Barto, 1998). The engineer must also pre-program the agent's set of possible actions. Predefining the state and action sets requires careful consideration and hard work by the programmer but in the end it greatly simplifies the learning task for the agent. Even though biological agents come equipped with complex sensory apparatuses and large behavioral repertoires they probably begin their lives with only a loose structure of what comprises state. The animal's state space is then altered and augmented through experience and learning. Since state definition in animals is a dynamic process rather fixed and predefined, it can be extremely difficult to determine all the information an animal uses to assess the state of its environment.

Also, because the computer scientist personally designs and implements the learning algorithms, he/she has explicit knowledge of how the agent generates its reward function. Determining an animal's reward function may seem an easy prospect when compared to a complete understanding of an animal's state or action set. After all, we can observe when a foraging animal receives a food reward. Several modeling studies based on primate neurobiology suggest that some animals may indeed distribute their experience of an observed reward among non-rewarding, intermediate environmental states as

predicted by a sequential decision-making approach (Daw & Touretzky, 2000; 2002). Still, gaps in our knowledge about the perception and internal information processing of animals will continue to place limits on how well we can apply the sequential decisionmaking framework to the behavior we observe in biological agents. Of course these constraints apply to any theoretical modeling of animal behavior where the exact neurological basis of the behavior under study is unknown. Often the reason for constructing mathematical models of biological systems is not only to explain our experimental observations but also to generate testable hypotheses for further investigations.

To determine if a sequential decision-making approach provides any additional explanatory power to what we already know about insect navigation it makes sense to first try to conceptually reduce the animal's problem into sequential decision-making components. Once we have a reasonable notion of what comprises the various states and actions for the animal, we may then be able to look at how the actions it takes in each state change with experience. We can then start to understand if insects, and perhaps other navigating animals, make use of learning and decision-making strategies similar to those proposed by the sequential decision-making framework. Even if during this integration we discover that these two fields cannot be directly applied to one another, we still stand to benefit from new questions and hypotheses that develop during the process.

In an effort to pave the way for further integration of these two fields, the remainder of this chapter will explore the results of existing behavioral experiments to gain insight into what honey bees might consider an environmental "state" for the purpose of making
navigational decisions. In addition I will discuss specialized behavioral actions used by honey bees to learn useful new perceptual information (and perhaps update their state).

Defining navigational states for honey bees

Many species of ants, wasps and bees are known to travel over long distances in search of provisions for their nest. The impressive navigational abilities achieved by the relatively simple nervous systems of insects have interested researchers for nearly a hundred years (see Frisch, 1967; Tinbergen, 1972; Collett & Collett, 2002 for review). Although studies investigating insect spatial abilities range over a variety of species, for a number of reasons the European honey bee *Apis mellifera* has become a particularly important model organism for the study of insect navigation.

One reason for the importance of honey bees is that they are excellent navigators able to travel as far as 10 kilometers in search of nectar and pollen and return to the nest (Visscher & Seeley 1982; Beekman & Ratnieks, 2000). Second, they are easy to care for and their behavior can be trained and observed in a range of natural and artificial environments (basic techniques for observation and training described in Frisch, 1967; a variety of environmental conditions reviewed in Collett & Collett, 2002). Finally, honey bees have a unique and observable communication system called the waggle dance which worker bees use to share the locations of valuable foraging patches with other hive-mates (Frisch, 1967). This is a valuable source of information for the study of animal navigation since it provides a direct window into the honey bees' internal representation of its environment, a window not available for any other animal (Dyer, 2002).

Using a variety of experimental techniques researchers have discovered a great deal about the role of visual information in honey bee navigation. I will focus mainly on what is known about visual navigation when hypothesizing about the inputs of a honey bee's environmental state since the bulk of existing data, including this dissertation, are primarily concerned with the insect visual system. This does not mean that odors and other sensory cues are irrelevant to foraging honey bees. In fact, olfactory cues exert enough control over honey bee foraging behavior that researchers take great care to control for the effects of odors during studies of visual navigation. Odors are commonly controlled by washing or replacing the surfaces of the experimental apparatus that bees come in contact with during training (see Srinivasan et al., 1997; Fry & Wehner; 2005 for examples) or by disabling pheromone glands in a technique called scent plugging (Towne & Gould, 1988). Both of these procedures are also described in the method sections of the following chapters.

Before discussing what a bee "knows" about the visual features of its environment, it is useful to describe what their visual system can perceive. First of all bees have color vision. Honey bees are able to see and distinguish among various shades of yellow, green and blue starting at ultraviolet wavelengths around 300 nm and extending all the way into the reddish around 670 nm (Frisch, 1967; Menzel & Backhaus, 1991). They are also able to detect patterns of polarized light (Rossel & Wehner, 1986). A bee can observe nearly the entire 360 degree panorama around its longitudinal axis and can learn spatial patterns with approximately 2 degrees of visual resolution in the frontal and lateral portions of its two compound eyes (Wehner, 1981; Srinivasan & Lehrer, 1988; Giurfa et al., 1996).

Path integration

With all of this visual information, honey bees have a lot to choose from to define the states of their environment. One well-studied source of environmental information available to honey bees is path integration (reviewed in Collett & Collett; 2000). As described previously, the path integration system keeps a running total of the bee's distance and direction from its nest. In honey bees, this computation is done entirely through visual perception. Bees compute the distance portion of their path integration vector by monitoring the flow of visual texture that passes over there visual field during flight (Esch & Burns, 1995; Esch et al., 2001). The direction of the vector is measured in relation to the sun and patterns of polarized light in the sky (Frisch, 1967; Rossel & Wehner, 1984; Dyer & Dickinson 1996). By consulting an updated path integration vector, a honey bee can quickly determine the distance and direction of its home from any state of the environment. Path integration is very attractive as state input because it reduces information about the bee's state history into a single value that can be used to disambiguate potentially aliased environmental states.

<u>Visual landmarks</u>

Insects rely on visual features to guide navigation in a number of ways (reviewed in Collett & Collett 2002). Although bees can rely on their path integration system to guide travel to and from the nest, they seem to prefer using visual landmarks when they are available (Frisch & Lindauer, 1954; Chittka, 1995; Chittka & Geiger, 1995; Chittka et al., 1995; Srinivasan et al., 1997; Collett & Collett, 2004; Vladusich, 2005; Chapter 4 of this

thesis). Previous work suggests honey bees have two strategies for using landmarks. The first strategy is to use salient en-route landmarks as a series of visual beacons. Under these circumstances bees links short path integration vectors to the visual properties (such as color shape and size) of intermediate landmarks encountered on the path to the goal (Collett et al., 1993; Chittka, 1995; Srinivasan et al., 1997; Zhang et al., 2000; Collett et al., 2002; Collett & Collett, 2004). On subsequent trips she recalls the movement vector associated with each beacon to determine her next action. This action then leads the bee to either the goal or the next landmark in the sequence. These shorter local vectors are believed to be stored separately from the constantly updated global path integration vector discussed above (Collett & Collett, 2004).

When available, bees use visual and olfactory features of the final target (such as floral odors and patterns or nest odor) as the terminal beacon once they gets close to the goal area (Frisch, 1967; Fry & Wehner; 2002). When the target does not supply salient perceptual information bees can learn to rely on aspects of nearby visual landmarks to locate the goal (Cartwright & Collett, 1983; Cheng, 1999a, b; see also Chapter 5 for a detailed analysis). It is known that bees learn the visual sizes (Cartwright & Collett, 1979; Cheng, 1999a, b, see Chapter 5), colors (Cheng et al., 1986) and angles between (Cartwright & Collett; 1982 see Chapter 5) nearby landmarks in order to pinpoint cryptic food locations over multiple visits. It is thought that bees store these properties as a 2dimensional retinotopic template (a kind of visual "snapshot") which they use in a sequential matching process to locate the goal on subsequent visits (Cartwright & Collett, 1983; Collett & Baron, 1994; Collett & Rees, 1997 see also Chapter 5). Visible landmarks appear to play an important role in honey bee navigation and we would be remiss to not include them as a part of state. Given the current evidence, we might speculate that landmarks play a role as boundaries, allowing bees to create discrete states within their normally continuous environments. As previously described, bees are able to generalize visual landmarks to new environmental stimuli and there is also weak evidence that bees can use memories of previously encountered visual landmarks to help select actions in later environmental locations (Zhang et al., 1999; Chapter 2 for detailed discussion).

Olfactory cues

It is well known that honey bees use environmental and self-generated olfactory cues both in the hive to transmit information about foraging sites to other workers, and at feeding sites to aid their nest-mates in pinpointing a profitable foraging location (see Frisch, 1967; Dyer, 2002 for review). Over long distances there will times when a bees is positioned far from both the nest and the intended foraging site, where odors are less concentrated or absent. How bees use olfactory cues to make navigational decisions during these portions of a trip is less clear due to methodological difficulties associated with detecting the odor concentrations bees experience during free flight.

More recently a few experiments have managed to identify roles that odors might play in a sequential decision-making context during navigation. Srinvasan et al. (1998) showed that a honey bee could learn to choose between two visual stimuli at the decision point of a y-maze based on an odor she experienced in a previous maze chamber. Bees learned to select a yellow color stimulus when the previous odor was a lemon scent and

blue stimulus when the odor was mango. A more current experiment by the Reinhard et al. (2004) showed honey bees could use a combination of path integration and odor stimuli to select among familiar foraging locations. In this study honey bees were trained to collect scented sugar water at two environmental locations; both 50 meters from the nest but 30 meters apart. Each feeder was provided a different scent. During tests they presented one of the familiar odors inside the nest and found that bees preferentially visited the feeding location where that odor was experienced during training. This shows that bees associate the outbound path integration vector leading to each feeder with the scent/taste experienced at that location

These experiments support a role for odors as a state-defining percept. Specifically they provide evidence that bees can use a memory of olfactory cues to help select actions in states that are non-Markov and ambiguous.

Motivational state

Dyer et al. (2002) provided evidence that honey bees base their flight heading on their state of hunger. In this experiment they captured bees in glass vials just as they were leaving for the nest after filling up on sugar water at a training feeder. Bees that were released shortly after capture flew along their normal trajectory toward the expected location of the hive. Bees that were held for an extended period of time (and hence digested much of their nectar supply) often flew in the opposite direction, which led to the location of the training feeder. That hunger appears to change what action is optimal in a bee's current environmental state strongly suggests that it plays a role as a state defining percept.

Exploratory behaviors and selective perception

As discussed in an earlier section, it unrealistic to assume that bees and other biological agents begin their lives with complete knowledge of their environment. This expectation is illustrated in **Figure 1-1** where the environment tells the agent explicitly what state it is. In reality biological agents probably begin with some basic information about the characteristics of the environment. This basic information is then supplemented and changed through learning to better define important states in the environment.

For many artificial tasks it may be tractable to predefine states for the agent in order to speed up the learning of an optimal policy. Kaelbling et al. (1996) describe a packaging task where programmers pre-defined as many as 200,000 discrete states. Defining the state space a priori is an even more difficult prospect for real-time, real-world navigational tasks. In these situations many computer scientists supply the agent with separate algorithms that allow it to learn on its own, the stable perceptual features associated with each state it experiences (reviewed in McCallum, 1996; Kaelbling et al., 1996). Behavior evidence suggests that honey bees and other flying hymenoptera also come equipped with a special strategies that aid the learning of useful environmental information. One such strategy is a specialized learning flight that honey bees perform at particular times and locations during navigation.

When the learning flight is conducted near the nest it is called an orientation flight. An orientation flight begins with the bee conducting a short bout of hovering flight while facing the entrance of the nest. She then widen her flight by performing one or more loops that start near the nest, range out as far as 18 meters and return to the starting point. An orientation flight can last as long as 8-10 minutes and bees usually perform several before setting out on the first actual foraging trips of their lives (Tinbergen & Kruyt, 1938; Iersel & Assem, 1964; Capaldi et al., 2000). Capaldi and Dyer (1999) performed a set of experiments showing that these flights improved honey bee's ability to find the nest. Their data suggest that this improved homing capacity is due to honey bees learning nest related landmarks during the orientation flight.

If the learning maneuver is conducted near a food source it is referred to as a turnback-and-look (TBL). The TBL is a highly stereotyped learning flight performed by honey bees and wasps upon leaving a newly rewarding foraging patch. Just like the beginning of orientation flights, the TBL begins with the insect flying in tight arcs while facing the goal. This arcing flight widens gradually into high circles and ends when the bee chooses a homeward flight path (Tinbergen & Kruyt, 1938; Lehrer, 1991; Zeil et al., 1996). The total length of a single TBL ranges from a few seconds to a couple minutes during which time bees are believed to use motion cues to select and learn the appearance of salient visual landmarks surrounding the goal (Zeil et al., 1996). This is where bees are thought to take their "snapshot" memory used to guide flight upon their return (Lehrer, 1993; Zeil et al., 1996). The flights become shorter after each subsequent visit to a foraging site until they eventually disappear. However, if a bee has difficulty finding the goal during a visit to a familiar foraging location, she will often begin a new round of TBLs after she eventually finds the reward. It appears that the TBL learning flight is initiated at new foraging sites and can be renewed at familiar sites when there is temporary uncertainty about the location of the goal (Wei et al., 2000). TBLs are thought

to play a critical role in landmark learning. Chapter 5 illustrates that the TBL is not necessary for the acquisition of certain spatial information although this does not exclude the possibility that the TBL allows for faster or more accurate learning.

These highly stereotyped flights appear to play an important role in the learning of visual features near important locations in the environment; namely the nest and profitable food locations. In addition to learning flights, many insects also perform spontaneous behaviors in response to particular visual stimuli. These behaviors, like learning flights, are suspected to aid the learning of depth information as well as visual features of the target stimuli and other nearby visual cues. Two such spontaneous behaviors are contour following and spontaneous flight toward large, salient objects.

When presented with a surface covered by a series of alternated black and white stripes, honey bees will spontaneously "scan" this pattern by flying back and forth along the contrast borders of the stripes (Lehrer et al., 1985). Generally bees prefer to perform scanning flights that run parallel rather than perpendicular to the contrast borders but, when two surfaces covered in the same bar pattern are presented at different heights, bees will modify their scanning flights to run perpendicular to the black and white stripes (Lehrer et al., 1988). The perpendicular scanning flight causes the higher, nearer surface to move faster over the retina than the lower more distant surface (Lehrer & Srinivasan, 1994). This motion parallax, which does not occur during scanning flights oriented parallel to the stripes, allows the bee to detect the differing heights of the patterned surfaces. For bees and wasps contour following is the general rule but they are able to modify their behavior to extract useful depth information from the visual scene.

Another behavior that might play a role in learning states is a phenomenon we call the spontaneous approach (Chapter 3). Honey bees and other insects are attracted to salient visual landmarks they encounter along a path to a goal. This attraction is strong and can cause the insect to deviate as far as 25 to 50 degrees from the flight heading of the straight goal path (Frisch, 1967; Chittka & Geiger, 1995; Chittka et al., 1995; Graham et al., 2003; Chapter 3). In ants, similar approach behavior appears to facilitate the learning of other visual features in the vicinity of the attractive landmark. This additional learning causes the ant to follow a path defined by the position of a salient off-route landmark even after the landmark is removed (Graham et al., 2003). Honey bees are found to revert to a straight goal path when the attractive intermediate landmark is absent (discussed in more detail in Chapter 3). At present the exact function of spontaneous approach behavior in honey bees under natural foraging conditions is unclear but I present it here as a possible learning behavior at least in ants.

These studies suggest that insects that are inexperienced with foraging or detect new or important perceptual cues are programmed to take specialized exploratory actions during which they learn important state information for use in later foraging episodes. In addition, some of these behaviors can re-emerge if a bee becomes uncertain about a familiar location. Since the turn-back-and-look takes place after feeding, perhaps it also serves as a means for bees to associate their recently experienced reward with the current state (during the initial arching flight near the goal) and distribute the value of this reward to the states leading up to the goal (during the much wider spiraling portion of the flight). It is not clear if we can hypothesize a similar reinforcement mechanism to the hivecentered orientation flights of new foragers since these bees do not receive an observable

food reward for their efforts. Perhaps the experience of re-entering their nest could serve as the terminal reward signal during these flights.

Summary

There is a great deal of evidence suggesting that honey bees base their navigational decisions on many sources of information. Bees appear to use many of these percepts in concert to select their actions. The path integration system can be used to associate visual landmarks, odor cues and hunger level to en-route flight vectors. Also, evidence suggests that bees can retain memories of previously experienced visual and olfactory percepts and use this information to aid action selection in later environmental states (this will be revisited in detail in Chapter 3). Having multiple exploratory strategies to learn new information and using multiple perceptual features to make decisions also provides the agent with back-up systems that it can apply when it is uncertain about its current state of the environment. The Back-up strategies are a common theme in insect navigation (Dyer, 1981; Collett & Collett, 2004).

Below I present a sequential decision-making dialog for honey bees similar to **Figure** 1-1 with more explicit notions of state and actions.

Environment							
External perceptual information							
There are two large black objects separated by 30 retinal units.							
L eff object is 10 horizontal by 20 vertical retinal units							
Dight object is 5 horizontal by 20 vertical retinal units.							
Right Object is 5 horizontal by 20 vertical fetinal units.							
Slight Horal odor present.							
Internal information							
Distance from the nest is 40 optical flow units.							
Direction from the nest is 30 degrees left of the sun.							
Honey crop 1/10 th full							
Bee							
Best match to current information is state 25.							
Actions for state 25.							
1.) move toward largest black object.							
2.) continue on current flight trajectory.							
3.) turn toward nest.							
4) begin to land							
Action 1 selected							
Environment							
Reward signal							
Reinforcement of 10 units received							
External percentual information							
External perceptual information							
I wo black objects on rear retina.							
One blue object occupying 20 vertical by 10 horizontal units							
on frontal retina.							
Internal information							
Distance from nest is 50 optical flow units							
Direction from the nest is 25 degrees left of the sun							
Honey crop 1/11 th full							
No odors							
Bee							
Update value of state-action pairing 1-25							
Current information matches state 26							
Actions for state 26							
1) maintain harding toward blue londmost							
1.) maintain nearing toward blue failuthark							
2.) turn back toward the nest $2 \times 1 \times 1 \times 1$							
3.) begin landing							
Action 1 selected etc							

Figure 1-2. Hypothetical state assessment and action selection of a honey bee

Prospects for an integrative research program

The general goal of this chapter was to introduce the overarching structure and ideas of the sequential decision-making framework and relate these concepts to some of what is known about insect navigation. Although it remains unclear how easily we can use notions from sequential decision-making to inform research questions in animal navigation (and vice versa), it is clear that such a broad conceptual overlap justifies continued efforts to integrate these two fields.

The rest of this thesis contains a series of experiments dealing with visually guided honey bee navigation. In the tradition of earlier research on this subject, the studies presented herein address questions about what visual information honey bees pay attention to as well as the strategies they employ to use this information to make navigational decisions. Some of these projects were directly inspired by concepts from the computer science literature. Others were generated through logical progression from the insect navigation literature and my own observations of honey bee behavior. Regardless of how these questions developed, they are all quite applicable to the ideas laid out in this introduction.

Chapter 2 presents work that was directly inspired by the sequential decision-making framework. This chapter deals with the issue of state and if honey bees use visual memory to aid decisions at perceptually ambiguous environmental locations.

Chapter 3 identifies and investigates the spontaneous approach behavior described previously in this chapter as a candidate behavior for gathering information. These experiments take a step towards understanding the properties of visual stimuli that drive this behavior and how this response affects subsequent navigation. Chapter 4 is concerned with the how honey bees might use en-route visual landmarks to define state and prompt actions when they are distant from the goal. The data presented in this chapter address the use of local vectors by honey bees over long distances and characterize a previously unreported strategy associated with en-route landmarks. Specifically, bees will repeatedly return to the most recently experienced enroute landmark when they are unsuccessful in finding the goal state.

Finally, chapter 5 looks closely at the snapshot model of visual navigation (Cartwright & Collett, 1983), which can be regarded as a generalization mechanism for simplifying the task of generating flight control actions based on visual landmarks in continuous environments. In these experiments we are concerned with the relative importance of template matching and visual beacons for determining states and actions when honey bees are near the goal.

To conclude, I hope that this chapter and the rest of this thesis illustrate the potential for developing new and useful ideas and research questions through the integration of ideas derived from both animal navigation and sequential decision-making.

Chapter 2

The use of visual memory by honey bees to solve a perceptually aliased maze task.

Introduction

The relatively simple nervous systems and impressive wayfinding skills of central place foraging insects such as the desert ant *Cataglyphis* and the European honey bee *Apis mellifera* have made them ideal model organisms for the study of animal navigation. Over the last ten decades researchers have learned much about the sensory inputs and behavioral strategies insects use to collect provisions from distant locations in their environment and return to the nest (reviewed in Frisch, 1967; Dyer, 1998; Collett, 2002; Wehner, 2003). In recent years, computer scientists and engineers have begun to look at the navigational strategies of insects and other animals for inspiration when constructing goal-directed, mobile robots (reviewed in Franz et al., 1998; Franz & Mallott, 2000).

Both animals and artificial agents face the problem of relying upon imperfect sensors to determine where they are in the environment. One version of this problem, called *perceptual aliasing*, (Whitehead & Ballard, 1991) describes situations where two different locations of the environment produce the same pattern of perceivable stimuli but require different responses from the agent to attain the goal.

Computer scientists have offered two general solutions to this problem. First the agent may disambiguate the two locations by consulting a record, stored in the form of short term memory, of the distinctive sequence of experienced stimuli leading up to the

aliased locations (reviewed in McCallum, 1996; Hasinoff, 2002; Shani & Brafman, 2004). Alternatively, the agent may take special actions designed to extract additional, disambiguating sensory information from the aliased scene. In situations where additional information is not available or is below the resolution of the agent's perceptual system the memory approach provides a more tractable solution to the problem. The difficulty with using previous experience to solve aliased choice points is determining just how much memory is necessary to disambiguate each state. After all, agents need to disambiguate their environment while remaining within the limits of their memory and information processing capabilities. Studies of how animals deal with similar problems may help guide further research on artificial agents.

It is not unreasonable to suggest that animals such as honey bees, which possess multiple sensory inputs that are higher in resolution than those typical of robots, will at times encounter perceptual aliasing in their natural environment. For example, a foraging honey bee should try to avoid revisiting a depleted flower patch but would not want to miss the opportunity to visit a similar looking nectar-rich patch in an adjacent location. If the two patches are close together they would likely share the same visual and olfactory cues and thus appear quite similar from a bee's point of view. Under these circumstances she might need some mechanism to internally "separate" the two patches in order to make an informed decision about where to land next.

There are several lines of evidence suggesting that bees can disambiguate perceptually aliased visual scenes. In one of the first experiments addressing this issue Collett & Kelber (1988) trained bees to find a sugar water feeder in relation to a landmark array within two geographically separate enclosures. The interiors of the two enclosures were

identical except for one feature: the goal was located in a different position relative to the landmarks in each hut. This setup presumably produced ambiguity since the information provided by the local stimuli did not adequately predict the position of the goal once a bee was inside an enclosure.

The authors found that with training, foraging bees searched at the correct feeder location in each enclosure. Given the arrangement of the apparatus, particularly that the enclosures were separated by 33 meters, bees could have employed a variety of strategies to accomplish the task.

Before describing in detail the sources of information available to honey bees to solve the hut task I would first like to mention how these strategies fit within the context of the computer science strategies of remembering and collapsing perceptual information. The first three mechanisms described below, namely path integration, hunger state and time of day, are strategies that combine multiple stimuli perceivable at the choice point into a single manageable source of information. These information signals do not retain information specific to any previously experienced environmental locations. However, the last three strategies presented, visual memory, olfactory memory and trial order do require honey bees to store memories of information specific to previous environmental states. Finally, I would like to stress that the following strategies are not mutually exclusive. It is possible that bees use any number of these sources of information simultaneously.

Path integration

It is well documented that honey bees and many other insects use a process called path integration to continuously monitor their distance and direction from the nest on outbound foraging trips (see Collett & Collett, 2000 for review). Honey bees measure direction in relation to celestial cues (Rossel & Wehner, 1986; Dyer & Dickinson, 1996) while the distance component is measured by integrating the amount of texture passing over her visual field during flight (Esch & Burns, 1995; Srinivasan et al., 2000). Since the two enclosures in the Collett & Kelber (1988) experiments were at different locations, bees could have distinguished between the two visually identical landmark arrays by consulting their current path integration value. Srinivasan et al. (1999) presented some evidence that bees could use the distance signal of their path integration system in this capacity.

<u>Time of day</u>

Another strategy is to consult an internal clock to associate time of day to reward locations. Gould (1987) showed that honey bees could be trained to approach the same artificial flower from different directions when the correct approach was predicted by the time of day, with one landing direction rewarded in the morning and the other direction rewarded in the afternoon. It is unlikely that time of day was a useful cue in the Collett & Kelber (1988) experiment since bees were trained to both enclosures throughout the day with only a few minutes separating visits to each. However, Gould's experiment suggests that bees can use time of day to reduce environmental ambiguity.

Hunger state

Since the feeder in each hut only provided half of a full nectar load, a bee could have solved the task by associating the different levels of nectar in her honey crop with the different locations of the feeder in each enclosure. It is important to note that this strategy could only work if bees visited the enclosures in a fixed sequence during each bout of foraging. Dyer et al. (2002) provided evidence that bees may use hunger level to choose between an outbound or homeward flight trajectory. Whether or not such information can inform short-range local decisions is an open question.

<u>Trip Order</u>

The sequence in which bees visited Collett and Kelber's (1988) huts provided another possible source of useful information. As mentioned above, the reward in each enclosure dispensed about half the volume of a full nectar load. This required bees to visit both huts on each foray to achieve a full honey-crop before returning to the nest. If bees visited the huts in a consistent order and learned this sequence of decisions they could use this information to guide search in each enclosure. Collett et al. (1993) showed that bees could successfully navigate through a small scale maze-like arena by learning to string together a sequence of motor commands. Whether or not bees can use this ability over the much larger foraging scale of the Collett and Kelber (1988) experiment is an unexplored question.

Visual and olfactory memory

Finally, bees may learn the visual or olfactory features of the environment just prior to entering each hut and once inside associate this memory with the correct search location. Zhang et al. (1999) provided direct evidence that bees are able to disambiguate a choice point using short-term visual memory. Using a v-maze design, they presented bees with a large yellow or blue color card in the first maze chamber. After passing through a hole in the center of the color card, the bee entered a second chamber where it was presented with two different patterned cards, a horizontal black and white grating on one arm of the y-maze and a vertical black and white grating on the other. In order to choose correctly, bees needed to associate the color of the card (the indicator stimulus) in the first chamber with the correct grating orientation at the choice point. Specifically, bees had to associate the memory of the blue color card with a horizontal grating and the yellow memory with a vertical grating. After identifying and passing through the correct pattern the bee entered a third chamber. This chamber presented the bee with two more stimuli. Here the bee needed to choose between a concentric ring pattern and a pie-like radial grating. Again bees had to associate the patterns in the third chamber with the patterns and colors in the previously visited chambers in order to predict reward position. The authors showed that bees were able to make this series of memory associations. In addition they found that once trained, bees could make informed decisions regardless of which order they encountered the stimuli in the maze.

Srinivasan et al. (1998) performed another similar maze experiment where olfactory cues served as the indicator stimulus. Again bees were proficient at solving the task,

showing that honey bees were able to associate olfactory memories with spatial decisions.

Inspired by the robotics problem of perceptual aliasing, we were initially interested in the relative priority with which honey bees used these sources of information for solving environmental ambiguities. We decided to focus on the interaction of path integration and visual memory since the methodologies used in previous work on these strategies could be integrated into a single experiment with relative ease (Srinivasan et al., 1999; Zhang et al., 1999). The idea was to provide informative path integration information as well as a salient visual indicator stimulus immediately before an ambiguous choice point. Once bees solved this task we intended to make either the visual memory association or the path integration vector uninformative to decision-making in order to determine how heavily each of these strategies contributed to their performance. Unfortunately we could not consistently reproduce the 75-80% correct choice performance reported in either of these previous works on visual memory associations or path integration (Zhang et al., 1999; Srinivasan et al., 1999). Our inability to reproduce visual memory associations in a maze task prompted a series of experiments to determine the critical methodological difference(s) responsible for our inability to obtain a result. The goal of this chapter is to examine under a range of conditions, the strength and significance of the memory association phenomena relative to other honey bee learning abilities.

Methods

General Training and Testing Procedures

All experiments conducted at Michigan State University (MSU) took place at a field site on the southern end of campus from June through September of the specified year. To limit natural forage and ensure a high level of motivation the bees and apparatus were contained in a mesh flight cage. Experiments conducted at the Australian National University (ANU) took place inside a greenhouse throughout the month of January 2001. The bees used in ANU experiments could forage either at our apparatus inside the greenhouse or outside under natural conditions. Once we trained bees to the apparatus inside the greenhouse they continued to faithfully visit the maze and rarely switched to foraging outdoors. Subjects used within each experiment came from the same nest.

Although we systematically varied different training and testing methods across experiments the general premise was the same. In order to receive the sugar water reward a bee had to fly first into the choice chamber via a large hole in the center of the indicator stimulus. Once in the choice chamber she would make a choice by flying through a hole in one of the two separate decision stimuli. After making a correct decision, bees were permitted to fill completely from a sugar-water feeder in the end chamber. In order to choose correctly, subjects needed to associate each of two indicator stimuli with its correct decision stimulus. For example, if the indicator stimulus was a yellow card, bees would only receive reward by passing through the vertical grating in the choice chamber. Conversely, when the indicator stimulus was blue, bees needed to select the horizontal grating. At no time was either the food or the indicator stimulus viewable from the choice chamber. This ensured that only a memory of the indicator stimulus could serve to disambiguate the choice point. The y-maze designs used for the MSU and ANU experiments are detailed in **Figures 2-1, 2-2** and **2-3**.

During a training visit, bees flew from the hive, entered the maze, received a reward and returned home. Bees were allowed up to four visits to each presentation of a single stimulus configuration. When the majority of bees had accumulated three to four visits to a stimulus configuration we alternated to the next configuration. We pseudo-randomly alternated the four possible stimulus configurations so that no one decision stimulus or arm of the maze was rewarded for more than eight consecutive visits. This helped prevent bees from continuously selecting one particular arm of the maze or decision stimulus.

Training lasted for seven to twelve days at the end of which even the least experienced bee had received a minimum of 60 rewards in the maze. 150 or more visits to the maze were usually necessary to receive 60 rewards since bees were only rewarded for correct decisions. Prior to training, we marked all bees with enamel paint for individual identification. Odors in the maze were controlled either by the presence of exhaust fans attached to the ends of the reward chambers (Zhang et al., 1999) or by the use of scent plugs to block recruitment pheromones (Towne & Gould, 1988). We applied a scent plug by spreading a heating mixture of bees wax and bow rosin over the two terminal body scales of a bee's abdomen. When the mixture cooled it formed a hard seal which prevented the bee from exposing her Nasanov gland in order to emit pheromones.



Figure 2-1. *MSU maze apparatus with tunnel*: The approach tunnel was 100cm long by 22cm wide and covered in random Julesz patter pixel size of 1 cm². The height of the apparatus was 22 cm throughout. Entrance holes, 4cm in diameter, were cut through the center of the stimuli. The tops of the maze chambers were covered with custom fitted Plexiglas® lids. Both horizontal and vertical grating patterns had a period of 20 degrees of visual angle when viewed from the center of the choice point.

ANU Apparatus



Figure 2-2. ANU cylinder maze apparatus: The tops of the maze chambers were covered with custom fitted Plexiglas® lids. Entrance and exit holes were 3 cm in diameter. All stimuli were 18 cm x 18 cm. Both horizontal and vertical grating patterns had a period of 18 degrees of visual angle when viewed from the center of the choice point.



Figure 2-3. *MSU maze apparatus without a tunnel*: The height of the apparatus was 22 cm throughout. Entrance holes 4cm in diameter were cut through the center of the stimuli. The tops of the maze chambers were covered with custom fitted Plexiglas® lids. Both horizontal and vertical grating patterns had a period of 20 degrees of visual angle when viewed from the center of the choice point.

At the end of training we tested the bees to determine if they made the proper associations between the indicator and decision stimuli. Before testing began we replaced all maze chambers with clean ones to remove any last remnants of non-volatile olfactory cues. We tested bees in 20 to 30 minute episodes interspersed with 20 to 60 minute bouts of training. During testing each bee entered individually to make her choice. To receive a score the bee had to pass through a choice stimulus and physically enter one of the final chambers. Only the first of such choices was recorded on each testing visit since some bees swiftly turned around and re-entered the choice point when they did not immediately find a reward feeder. The above protocol was repeated for seven different experiments. We varied specific details of the procedure for each subsequent experiment in an attempt to make the memory associations easier for honey bees to learn. The rest of the method section will enumerate the differences between each experimental condition.

Experimental variations

Condition 1

We conducted the first five experiments (conditions 1-5) at MSU over the summer of 2000. In the first condition bees experienced the colored indicator stimulus at the entrance of a 1m tunnel leading to the choice point (**Figure 2-1**). During training, bees were required to find their way back out of the maze after performing a correct decision and receiving a reward. Incorrect decisions led to a non-rewarding sham feeder filled with water. After bees tasted the water, we allowed them to back-track to the choice point and choose again. Feeders were absent during testing and after a bee made a test decision we released it from the maze rather than requiring the bee to find her own way out.

Condition 2

Training and testing for the second condition was identical to method 1 except the color card was presented at the end of the tunnel just prior to the choice point. In this case we hoped that having the indicator stimulus closer in time/distance to the choice point might facilitate learning.

Condition 3

For the third method we also trained and tested bees with the indicator stimulus at the end of the tunnel immediately before the choice point. In addition we removed the sham feeder from the incorrect arm of the maze during training. Bees tended to linger at the unrewarding water feeder in condition 2. We thought that the presence of the sham feeder might have caused an erroneous reward signal for the bees and therefore undermined the association of the indicator stimuli with the correct decision stimuli.

Condition 4

The fourth experiment preserved all of the changes in conditions 2 and 3. In addition, during training we immediately released bees from the apparatus after an incorrect decision. We also released correct bees after feeding rather than having them fly back out of the entrance. We started waving bees out after incorrect choices because we thought that letting bees "correct" their decisions by backing out of the wrong end chamber and flying into the correct one might be causing some kind interference via the formation of multiple associations. We released correct bees after feeding in order to speed up the visitation rate to our apparatus.

Condition 5

Training and testing for condition 5 maintained all of the changes as condition 4. In addition we blocked the incorrect path throughout training in an attempt to facilitate learning of the correct associations.

Condition 6

Experiments at the ANU campus were conducted over the month of January. The cylinder y-maze design is detailed in **Figure 2-2**.

We did not use scent plug to control for odors. Instead, small fans were mounted on the back of the end chambers. The fans were supposed to draw air, along with any recruitment pheromone, out of the apparatus.

Over the first two to three days of training, the hole in the incorrect choice stimulus was blocked with a sheet of acetate thus forcing the bees through the correct pattern. This forced choice training was not reported in the methods of the original paper (Zhang et al., 1999). During this time a 6cm x 6cm paper baffle was gradually inserted over the entrance hole to each chamber. The purpose of the baffles was to prevent bees in the maze from seeing the features of subsequent and previous maze chambers. After the second or third day we began proper training. We removed the acetate sheet at this point so individuals could make decisions freely. After a mistake, bees were released from the incorrect chamber and permitted to re-enter the beginning of the maze. Those that made correct choices received a reward and were released to return to the hive. Any bee that made two consecutive incorrect decisions was corrected using the acetate sheet during its third attempt with that stimulus configuration. Again the acetate correction procedure was not reported in the original paper (Zhang et al., 1999). The procedure of reward and correction continued throughout proper training and testing. Proper training lasted for six to eight days giving a minimum of 60 round trips per individual. The training and testing procedures were identical except for the first 1 or 2 days of constant forced choices.

47

÷

During tests, if a bee made two consecutive incorrect choices without returning to the nest it received a forced choice with acetate covering the incorrect maze chamber on its third visit. This sequence of events counted as a single incorrect decision. If the bee's first choice was incorrect and she made a correct second choice before returning to the nest, she was allowed to feed and was released to the nest. This was also scored as a single incorrect decision. During testing each bee was permitted to make as many as four such decisions during any given presentation of a single stimulus configuration.

Condition 7

The seventh experiment was conducted at MSU during the summer of 2001. The training and testing procedures for condition seven were the same as the ANU experiment with the following exceptions: 1.) The maze chambers were cubes (Figure 2-3) rather than cylinders 2.) Baffles were not installed between maze chambers. 3.) We conducted the experiment outdoors in a flight cage, not inside a greenhouse. 4.) We used scent plugs for odor control rather than exhaust fans. Table 2-1 lists the differences between all seven experimental procedures.

Simple discriminations

It was important for us to be sure our bees could sense the difference between the blue and yellow stimuli as well as the horizontal and vertical gratings. In the first of two experiments the indicator stimulus was removed and either horizontal or vertical grating was made permanently rewarding. We alternated the left and right positions of the gratings to ensure only the pattern and not the position predicted reward. Test decisions

48

÷

were not rewarded and after a choice we released the bee from the end chamber by lifting the Plexiglas® lid. After testing, the same group of bees was trained to the previously unrewarding pattern (reversal learning) and tested in the same manner. We then repeated this assay with a separate group of bees using the color stimuli at the choice point rather than the gratings. Each control experiment was completed in less than a day with a group of ten to twenty bees.

	Condition 1	Condition 2	Condition 3	Condition 4	Condition 5	Condition 6	Condition 7
Location	MSU	MSU	MSU	MSU	MSU	ANU	MSU
Odor Control	Scent Plugs	Scent Plugs	Scent Plugs	Scent Plugs	Scent Plugs	Ventilation Fans	Scent Plugs
Maze Layout	Y-maze With Tunnel	Y-maze With Tunnel	Y-maze With Tunnel	Y-maze With Tunnel	Y-maze With Tunnel	Cylinder Maze	Y-maze
Distance of Indicator From Choice Point	1 M	1 mm	1 mm	l mm	1 mm	1 mm	1 mm
Apparatus Indoors						+	
Baffles Present						+	
Sham Feeder In Incorrect Choice Box	+	+					
Incorrect Bees Permitted to Turn Around	+						
Shaped Choices During Training					+	+	+
Released Bees After Decision				+	+	+	+
Rewarded Test Trials						+	+
Bees let in one at a time during training						+	+

 Table 1-2. Condition variations across maze experiments.

Results

<u>Overall</u>

Figure 2-4 shows the control data for simple color and pattern discriminations as well as the complementary reversal learning data. Bees performed well on the simple discrimination tasks in a short training period. Learning to choose the opposite color or pattern (reversal learning condition) required slightly more training but also resulted in 80-90 % correct decision.



Simple Discrimination Choice Performance

Figure 2-4. Simple discrimination over time: Percent correct choices based on the number of visits to the apparatus for color (solid line) and grating (dashed line) discrimination tasks. First bees we trained bees to one of two stimuli followed by a reversal stage where the other alternate stimulus predicted reward.

Figure 2-5 provides a summary of choice performance for each condition of the visual memory task. This analysis used only the first choices of bees for each presentation of the four stimuli configurations. By analyzing only the first decisions for each stimuli configuration we control for effects of local, non-associative strategies such as win-stay-loose-shift (Demas & Brown, 1995). For each condition we tested choice performance against an expected random decision frequency (50/50) using a chi square goodness of fit test. Only the ANU (condition 6) experiment resulted in a choice frequency statistically better than 50% ($\chi^2(1,n=124) = 5.45$, p=.02) while MSU bees in condition 4 performed significantly worse than random ($\chi^2(1,n=163) = 9.331$, p=.002). Since condition 6 was the only experiment to produce the desired associative learning, we broke the data down to more closely inspect the performance of individual bees in this condition (**Figure 2-6**). Binomial tests revealed that only one bee performed statistically better than random (n=17, p=0.0472), although seven of the ten bees chose the correct option more often than the incorrect one.



Figure 2-5. Complex association performance across experimental conditions: Percent correct choices across bees for each condition. Numbers inside the bars are the sample sizes for each condition. Data for condition 6 was collecting at ANU. All other experimental took place at MSU.



Condition 6: Individual First Choice Performance

Figure 2-6. Complex association performance of individual bees trained and tested at ANU: Percent correct decisions for each individual bee. Numbers inside the bars are the total number of decisions performed by each individual.

Figure 2-7 shows the first choice performances throughout training for condition 4, 6 and 7. Since we did not record training decisions in conditions 1-3 and all training choices were forced in condition 5 these experiments are not represented in the figure. The figure presents the percent of correct choices in ten-visit blocks in each experimental condition not including after the initial forced choice training. Conditions 6 and 7 plots are longer than condition 4 because testing visits were also rewarded and could therefore be included. In all three conditions, choice performance throughout training and testing fluctuated around 50% with the ANU data remaining above random most consistently.



Figure 2-7. Complex association over time: Percent correct decisions made by all bees in 10 trial blocks for condition 4 (circles), 6 (squares) and 7 (triangles).

Dependence of second choices

Table 2-2 shows the occurrence of correct and incorrect second choices given correct or incorrect first choices for the last day of training/testing. Again we looked at conditions 4, 6 and 7 since 2^{nd} decisions were permitted in these experiments. In the case of condition 4 we looked at the last day of training since multiple choices were not permitted during tests in this experiment. We analyzed each contingency table using a chi square test of independence. The second decision honey bees made in each stimuli configuration was dependent on the first for condition 4 ($\chi^2(1,n=255)=19.35$, p<0.0001)
and condition 6 ($\chi^2(1,n=90)$ =3.731, p=.050). By looking at the table we can see that honey bees that made a correct first decision tended to also make a correct second choice in these two experiments. When the first decision was incorrect bees made a more or less random second decision. Bees in condition 7 appeared to choose randomly regardless of previous experience. This analysis was conducted to determine if bees were using a local strategy such as win-stay-loose shift rather than making the desired associations between the presented stimuli.

		Cond	ition 4	
2nd choice				
		Correct	Incorrect	
lst choice	Correct	103	46	69% stayed correct
	Incorrect	44	62	58% stayed incorrect
		Cond	ition 6	
		2nd ch	noice	
]	Correct	Incorrect	1
lst choice	Correct	41	18	69% stayed correct
	Incorrect	15	16	52% stayed incorrect
		Cond	ition 7	
		2 e = = = =	hoice	
	ſ	Compart		1
	Comment	<u> </u>	incorrect	520/
si cnoice	Correct	40	41	53% stayed correct
	Incorrect	49	55	153% stayed incorrect

Table 2-2. Choice Dependencies: The tables show the relationship of honey bee decisions on two consecutive trips to the same stimulus configuration in the three training conditions where such behavior was permitted.

Discussion

Any animal or robot that travels significant distances to acquire resources must learn

features of its environment that facilitate safe and reliable travel from one destination to

another. This can be difficult considering the wealth of information the environment provides to even simple sensory systems. In order to reduce computational load the navigator must selectively attend to useful information and disregard the remainder as noise. Even when the agent's sensory system and behavioral repertoire are finely tuned to its environment, as is the case for animals, sometimes two or more locations in the environment share very similar sensory information. If these locations require different behavioral responses the agent is faced with an ambiguity (Whitehead & Ballard, 1991; McCallum, 1996; Hasinoff, 2002; Shani & Brafman, 2004). When an agent finds itself in an ambiguous environmental state it might resolve its uncertainty by using some form of stored information about the previous experience that led to its current location. In a series of experiments, we investigated whether or not honey bees could use short-term, visual memory to solve a perceptually ambiguous maze task.

In the summer of 1999 we set out to determine if honey bees could use visual memory to solve a perceptually ambiguous decision task in a maze apparatus modeled after Zhang et al. (1996). After several attempts by our laboratory which showed no evidence of learning, Zhang et al (1999) published a study showing that bees could indeed learn just this sort of strategy in a y-maze similar to ours. As a follow up we designed several experiments intended to extend their results.

The results of the first follow up study provided no evidence that bees could learn the association between the indicator and decision stimuli. At this point we decided to replicate the Zhang et al. (1999) result before we proceeded with any further studies. In a series of experiments we systematically changed training and testing conditions in order to isolate the factor(s) responsible for our negative results (**Table 2-1** *conditions 1*

through 5). After several unsuccessful attempts during the summer of 2000 (**Figure 2-5** *conditions 1 through 5*) we communicated our difficulties with the ANU laboratory.

Over the winter of 2001 we collaborated on an experiment in their laboratory that replicated their result with statistical significance (**Figure 2-5** condition 6). The summer of that year we again conducted the experiment at our MSU field site, modifying our apparatus and procedures to closely match those used in Australia while retaining key elements of our design (**Table 2-1** condition 7). Once again we found no evidence that bees could learn the desired association. In the following sections we will enumerate the differences in our training methods and discuss how each of them may have hindered learning in our MSU experimental conditions.

Training methods

We will focus on the methodological differences between the ANU experiment (condition 6) where we found learning, and MSU condition 7 where our methods most closely resembled condition 6 but did not produce a significant result. There were four fundamental differences in the training methods between these two experiments; 1.) At MSU we used scent plugs instead of ventilation fans to control for odors, 2.) the MSU experiment was conducted in full view of the sky instead of indoors, 3.) the MSU maze was constructed from square wooden boxes rather than plastic cylinders, 4.) and the MSU maze did not have baffles obstructing the entrance holes between maze chambers. We will address each of these differences in turn in the following sections.

Odor control

It seems unlikely that the application of scent plugs impeded learning in condition 7. Since these experiments our laboratory has consistently used scent plugs on honey bees to control for odor in a variety of behavioral assays and have yet to notice even subtle changes in flight or navigation ability. In addition, the bees trained for the simple discrimination tasks were also scent plugged. If scent plugging substantially affecting learning we might expect that choice performance in the simple discrimination would also suffer. **Figure 2-4** shows that bees in these experiments displayed near-perfect decision-making, which is consistent with previous results from similar studies (Frisch, 1967; Giurfa, 2004).

Indoors vs. outdoors

It is well documented that honey bees rely heavily on celestial cues during natural foraging tasks (Frisch, 1967; Rossel & Wehner, 1986; Dyer & Dickinson, 1996). Perhaps there is a worry that the availability of the sun and patterns of polarized light in the sky overshadowed the local color and pattern stimuli provided in the outdoor MSU maze apparatus. In a recent investigation of the use of visual landmarks and odometry by honey bees Vladusich et al. (2005) reported that the precision of honey bee search flights in their outdoor apparatus was poorer than that found by Srinivasan et al. (1997) in indoor experiments using similar methodologies. If celestial information does eliminate learning in these complex tasks this would certainly cast doubt on the importance of such strategies during natural honey bee navigation where a full view of the sky is practically the norm.

Cylinder maze chambers vs. boxes

The only major difference between the two apparatuses was the layout of the choice chamber. The construction of the MSU y-maze was such that from the entrance of the choice chamber there was an angular distance of approximately 80 degrees between the centers of the two decision stimuli (**Figure 2-3**). The angle between decision stimuli in the ANU apparatus was approximately 45 degrees (**Figure 2-2**). This means that the ANU decision stimuli filled more of a bee's frontal retina during her entrance into the choice chamber than did the decision stimuli in the MSU maze. This design difference should have had little affect on honey bee performance since Giger & Srinivasan (1997) showed that bees learn to discriminate similar grating patterns equally well with frontal and lateral portions of their retina. Also, the fact that bees performed well in the simple discrimination suggests that the shape of the choice chamber and angular distances between our decision stimuli are probably not the cause of the negative result. The simple discriminations were conducted at MSU in the same cubical y-maze chambers with the same decision stimuli as used in the indicator task.

Baffles

During training and testing in the MSU y-mazes bees were allowed swift and unobstructed passage between all of the chambers in our maze apparatus. Without the presence of baffles bees had the potential to view both the indicator stimulus and the decision stimuli simultaneously just before passing from the indicator chamber into the choice chamber. The baffles present between maze chambers in the ANU apparatus prevented bees from seeing the decision patterns (or colors) from anywhere inside the

indicator chamber. It is possible that a view of the decision stimuli from the indicator chamber caused bees to shift their attention away from the indicator stimulus and therefore hindered learning of the desired association. The baffles may have played another less obvious role. Dyer & Chittka (2004) provided evidence that bumble bees perform better on difficult visual discrimination tasks when they take more time to make decisions. The baffles in the ANU experiment blocked a straight flight from chamber to chamber causing bees to "slow down" and spend more time in each chamber before moving on to the next. This means that bees in the ANU maze received more exposure to the indicator stimulus (as well as the decision stimuli) than MSU bees during each trip through the maze. This could have helped reinforce the complex association over successive visits as well as ensuring bees had enough time to recognize each stimulus before moving on to the next maze chamber. Whether or not the speed/accuracy trade off found in bumble bees is present for honey bees in these types of maze tasks is unknown.

A local strategy for solving the maze task

Although bees did not learn the desired stimuli association in our MSU experiments there is some evidence that they may have improved their ability to find the reward using a different strategy. The data presented in **Table 2-2** show that during training in some experiments, bees that chose correctly on their first visit to the stimuli configuration had a higher probability of choosing correctly again on their next visit if the stimulus configuration remained unchanged. Such an advantage on the second visit was not seen for bees that were incorrect on the first trial. One could argue that the pattern of high

correct choice frequency reported in **Table 2-2** resulted from bees learning the relationship among the maze stimuli. This is certainly not the case for bees in condition 4 whose poor performance during unrewarded test trials rules out such an interpretation. Here we suggest that bees may have implemented a win-stay strategy where they were biased to repeat the same sequence of actions so long as those actions remain rewarding. Win-stay and win-shift strategies for solving discrimination and spatial tasks are well known for rats, (Olton et al., 1981) primates, (Levine, 1959) and birds, (Kamil et al., 1977) as well as for honey bees (Demas & Brown, 1995). Still, it is difficult to explain the choice results of the ANU experiment with a win-stay strategy since bees performed the task relatively well even when we counted only their first choices to each presentation of the four stimulus configurations (**Figure 2-5**).

The impressive learning results reported by Zhang et al. (1999) for the even more complicated dual-decision maze task were generated by including all honey bee choices (1st, 2nd, 3rd, etc.) made to all stimulus configurations. The authors assumed each visit by a bee to the maze was statistically independent based on a previous maze study which found no dependence of choices across multiple visits (Srinivasan & Lehrer, 1988). The problem with this assumption is that in the 1988 experiments bees were trained to solve simple discrimination tasks. Our results as well as Srinivasan & Lehrer (1998) show that bees choose with remarkable accuracy during simple discriminations within twenty or fewer visits (**Figure 2-4**). Srinivasan and Lehrer (1988) only trained subjects for approximately two hours by which time bees had reached their peak performance of 90-100% correct choices. Perhaps the many hours of variable performance experienced by bees in the complex association task allowed a win stay strategy to take root. If this is the

case, analyzing multiple choices per bee during a single presentation of a stimulus configuration could artificially inflate the smaller effect of bees actually learning the proper associations between stimuli. However, this objection does not apply to analyses based on first choices.

Puzzling results

There are two results that we find especially hard to interpret. First, why is it that bees in condition 4 chose significantly below chance level during tests? **Table 2-2** suggests that during training these bees used some sort of win-stay strategy since performance disappeared during unrewarded test visits. Regardless of whether test visits were unrewarded, we took great care in all of our experiments to alternate the four possible stimulus configurations so as not to bias choice frequencies above or below random.

It is also curious that the bees in condition 7 learned neither the desired association nor a local win-stay strategy. Instead their choices were distributed at random throughout the experiment. Thus far we are unable to account for these results.

Conclusions

The long training periods necessary for bees to learn the complex association in the ANU maze often consumed the entire foraging life span of the honey bees, and yet yielded modest performance at best (**Figure 2-7**). This stands in contrast with the quick and accurate learning of bees during simple discrimination tasks (**Figures 2-4**). This

contrast raises questions about the importance of complex memory associations in the foraging strategies of bees. Given that there is some evidence for learning in the ANU experiment and the absence of clear evidence of experimental artifacts, we would not suggest that the ability does not exist. Perhaps under natural foraging conditions bees make these associations over a shorter time span which would make it of some value to bees. At present however it seems unlikely that this strategy plays a very important role in the bees' lives. Hopefully future work will be able to isolate the critical element(s) necessary to quickly and reliably produce this behavior in a maze environment.

Can honey bees use visual memory to resolve ambiguous decision points? The answer appears to be yes but high sensitivity to slight methodological changes and the long training periods necessary for learning to occur suggest that making these complex associations is quite difficult for honey bees. Such associations probably play a minor role when compared to other foraging strategies such as simple discriminations for which learning is fast and accurate.

Chapter 3

How honey bees use a novel off-route landmark near the goal.

Introduction

In order to return efficiently to a location in the environment, a mobile agent must select stable sensory cues that reliably mark the path to its intended goal. Within a short distance from the goal location however, cues emanating from the goal itself often provide enough information for successful orientation. When the target is too distant or difficult to perceive from its current location, the agent must rely on other information to guide its path. This distinction between proximate and distant navigation was first recognized by Watson & Lashley (1915). Subsequent research has shown that over extended journeys to and from their homes, biological agents such as insects and mammals use distinct landmarks encountered en-route to their destination to guide travel (reviews: *insects*- Collett, 1996; Wehner, 1996; Dyer, 1998; Collett & Collett, 2002. *mammals*- Benhamou et al, 1995; Etienne et al 1996). For an agent to successfully utilize en-route landmarks to reach a destination, it should commit to memory only those landmarks that are easily recognized from a direct path to the goal and are unlikely to change over time.

Often, mobile agents must deal with a wide variety of changes to familiar landscapes in order to successfully navigate through their environment. What sorts of changes require attention? This seems a simple question at first. If a change in perceptual information at a critical decision point makes an agent unsure of the correct next action, it should make an effort to resolve this uncertainty. The question becomes more difficult when considering subtler changes in the appearance of known landmarks or perhaps changes in the size and configuration of landmarks that are viewable but relatively distant from a given route. Which of these changes should be incorporated into the agent's knowledge of the environment? In humans, addressing these sorts of questions experimentally can be difficult. Erecting, demolishing or otherwise changing large landmarks in natural human environments is rarely tractable (Gollege, 1999; Maguire et al., 1999). Fortunately, ethologists and psychologists study way-finding and other spatial behavior in a variety of organisms whose environments are better suited for controlling and manipulating the cues available for navigation. In our experiments we examine wayfinding behavior in the European honey bee *Apis melifera* which is known to use visual cues for navigation over a range of spatial scales (reviewed in Collett, 1996; Dyer, 1998; Collett & Collett, 2002)

During her lifetime, a foraging honey bee makes dozens if not hundreds of trips to nectar locations as distant as ten kilometers from her home (Frisch, 1967; Visscher & Seeley, 1982; Beekman & Ratniek, 2000). Although a bee spends only about ten days foraging outside the nest, (Neukirch 1982) the sheer number and length of her forays provides sufficient opportunity for her to experience environmental changes analogous to those described above. Dating back to the classic investigations of digger wasp homing conducted by Tinbergen (1972), knowledge about insect visual navigation has come from studies where an experimenter changes a feature of the animal's familiar environment (e.g. Anderson, 1977; Cartwright & Collett 1983; Chittka, 1995; Srinivasan et al., 1997; Collett et al., 2002; Chittka & Gieger, 1995; Cheng 1999a, b; Fry & Wehner 2005;

Vladusich et al., 2005). In these experiments, the subject's response to the change provides the observer insight into which stimuli or parameters of the stimuli are of value for goal directed movement. These sorts of experiments have uncovered a great deal about the sensory information honey bees use to revisit a familiar foraging location and subsequently find their way back to the nest.

However, such studies generally do not examine how the navigating animal updates her knowledge of the environment upon experiencing a change to the environment that subsequently endures. We focus on a simple version of this problem by studying how bees respond to the presentation of an additional landmark introduced in a familiar location. The work in this chapter continues in the spirit of earlier experiments in that we trained bees to find a reward in regard to landmarks within a carefully controlled environment. After we trained bees thoroughly with the initial landmark array, we added a landmark intermediate (en-route) to the goal to understand how bees respond to and subsequently deal with the new element along their familiar route. In addition, after we trained bees for an extended period of time in the presence of an introduced landmark we observed their flight behavior when this element was removed. In a second set of experiments we varied the distance of the novel landmark from the straight goal path or its visual size as seen from the starting point to determine how these two landmark characteristics modulated affect honey bee flight paths.

Based on our results and previous work done with ants and bees (Lehrer, 1993; Lehrer & Srinivasan, 1994; Graham et al., 2003) we suggest that insects and perhaps other organisms might use specialized behaviors to pay particular attention to salient landmarks in their environment. These landmarks may then be used to facilitate the learning of

other stable but more subtle visual cues along their travel routes. In addition, we propose that honey bees may divide an otherwise complex and continuous path into a series of discrete segments designated by the visual cues they identify during these specialized behaviors. In a sequential decision-making framework such strategies could serve to reduce large, continuous environments into a discrete and more tractable number of states as well as help the honey bee select sensory cues that reliably define important states of the environment

Methods

<u>Apparatus</u>

Our apparatus, modeled after a design developed by Fry et al. (2000), consisted of an indoor arena with a camera positioned above to track bee flight paths (**Figure 3-1a, c**). The arena floor was 200 cm in diameter with a 22 cm high wall running continuously around the perimeter. A white, nylon tent extended up another 178 cm from the edge of the wall with the camera positioned into a hole sewn in the center of the ceiling. The floor and wall of the arena were painted flat white and the wall had a 2 cm in diameter entrance hole centered 11 cm from the arena floor. A 1.2 meter, clear, plastic entrance-tube led from an outdoor flight cage to the entrance hole. The end of the tube was fitted with an elbow that terminated in a small 4 x 4 cm square platform. This platform encouraged bees to fly rather than walk from the entrance into the arena. To provide approximately homogenous lighting inside the arena, we spaced four 500-watt halogen work lamps evenly around the outside of the tent.

We tracked bee flight paths in the arena using a pan-tilt Sony surveillance camera and BIObserve Trackit software (Fry et al., 2000). Using the bee's contrast against the white arena background, Trackit recorded the bee's position and body axis in the arena at 60 Hz. Bee track paths were plotted and analyzed using custom programs written in Matlab.



Figure 3-1. Arena **a.**) Overall dimensions **b.**) Pre-training feeder box **c.**) Detailed side-view of arena. The arena is a one way system. Bees walk into the arena via the entrance tube, fly from the entrance platform to the feeder-hole and crawl into the feeder box below. After feeding they walk out to the flight cage via the exit tube and return to the nest.

Training and Testing Procedures

Before we trained a group of bees to the inside of the arena they first underwent a preliminary training process. During preliminary training bees learned, step by step, how to crawl into a feeder hole, through a one-way door and how to leave through an exit tube after feeding. We conducted the pre-training in the flight cage for approximately two days using a replica of the one-way feeding system found below the arena floor (**Figure 3-1b**). By the end of pre-training, bees were proficient at maneuvering through the apparatus and could be let into the arena.

Inside the arena, bees needed to fly from the entrance platform to a 1.5 cm in diameter hole in the arena floor using the landmarks provided. During the first few hours of training bees often got lost in the arena, requiring us to release confused bees from the tent periodically. When we had a group of at least 40 bees cycling through the arena we captured each bee and paint marked it for individual identification. Training continued for and additional three to five days to ensure that each bee had a minimum of 60 visits to the landmark array inside the arena. We started each experiment with such a large group of subjects because many bees stopped returning during the often protracted training and testing procedures. A starting group that ranged between 40-50 bees usually dwindled to around 10-20 by the end of an experiment.

At the end of training we allowed bees into the arena one at a time for testing using a small door in the entrance tube. During tests the feeder-hole was covered with a small sheet of white paper and each bee was tracked during the first 1 to 3 minutes of search flight. After the tracking period she was placed in the feeder box to receive reward and exit to the nest. Testing in all experiments began by tracking bee flights with the

landmarks in the original training configuration. We then conducted follow-up tests where we manipulated some aspect of the original landmarks. Test trials tended to slow down visitation to the arena so we interspersed testing sessions with bouts of en masse training to keep motivation high.

To control olfactory cues we used a mixture of beeswax and bow rosin to cover the Nasanov gland on the abdomen of each bee (Towne & Gould, 1988). Plugging this gland prevented bees from emitting pheromone that might otherwise be used to aid navigation. The application of the scent plug occurred during the paint marking of individual bees. To further control odors during testing we placed a fresh, clean feeding apparatus under the arena and thoroughly washed the entire arena floor with a solution of odorless soap and warm water. In addition we replaced training landmarks with clean, identical testing landmarks.

<u>Data Analysis</u>

To quantify honey bees' responses to an introduced landmark we used two basic measurements; the flight direction flown from the entrance of the arena and the distance traveled from the entrance to the feeder. Henceforth, these measurements will be referred to as heading and path length. Flight headings were calculated using the Euclidean coordinates of the arena entrance and the point at which each bee had traveled a net distance of 50 cm from the entrance. At this distance bees had an opportunity to commit to a flight direction but had not yet passed by the novel landmark in any of our experiments. We measured path length for each bee by summing all of her x and y displacements starting at 10 cm from the entrance and ending 20 cm from the feeder hole.

Truncating the beginning and ends of the flight path was necessary to exclude walking behavior which could artificially inflate the path length values.

Results

Experiment 1

The purpose of conducting the first two arena experiments was three-fold. First we wanted to quantify the magnitude of bees' initial, spontaneous response to a novel landmark. The second goal was to track how this response changed over prolonged experience with the novel element left in place. Finally, we wanted to know if differences in the visual properties of the initial training array affected bees' overall response to a novel landmark.

To address these questions we trained bees to find the feeder hole using either a distant, extended panoramic landmark (**Figure 3-2a** for the floor plan) or two nearby cylinder landmarks (**Figure 3-3a**). After training, we first tested bees in the training configuration (**Figure 3-2a, 3-3a**). Next we positioned an additional landmark approximately half way to and 39⁰ (panorama) or 43⁰ (two cylinders) to the left of the feeder hole (**Figure 3-2b, 3-3b**). After tracking bees' initial response to the novel landmark we continued training with the cylinder present. We conducted additional testing trials on each bee's tenth visit (**Figure 3-2c, 3-3c**) and again at twenty (**Figure 3-2d, 3-3d**) or more visits. We performed an additional manipulation in the panorama experiment where we removed the added landmark after twenty or more visits and tracked bees' responses in the original training configuration (**Figure 3-2e**).

Headings

The data in Experiment 1 were collected over three years and the placement of the novel landmark was not exactly the same during replications. In order to provide a common measure for the panoramic and 2-cylinder training groups we transformed the angular headings into the ratio of the angular distance from feeder to the novel landmark. Statistical comparisons were done using these values rather than the raw angular data. We performed t-tests on a series of eight comparisons; 1-2.) training condition vs. first trip with the novel landmark present within both panorama and cylinder trained groups, 3-4.) training condition vs. 20+ trips with the novel landmark present within both panorama and cylinder trained groups, 5.) training condition vs. novel landmark removed within the panorama trained group, 6-8.) panorama vs. cylinder trained bees for the training, first trip and 20+ trip conditions. The Bonferroni method for correcting for multiple comparisons was applied to all the p-values reported throughout the results section.

Flight headings to both the 2-cylinder and panoramic training arrays, although variable, clustered in the direction of the feeder hole. Figures 3-2a and 3-3a show the actual flight paths from the entrance to the 50 cm cut-off of bees in the two training conditions. A summary of the data is presented to the left of the flight plots. Each black dot around the semi-circles represents a measurement of a single bee. The angle and strength of the mean heading vector is represented by direction and length of the arrow in the semi circle. When the novel landmark was introduced (Figures 3-2b & 3-3b), flight headings of bees in both treatments shifted dramatically toward it. This shift in heading was statistically significant for the panoramic (t(42) = 4.26, p = 0.0008) and the two-

cylinder landmark conditions (t(27) = 5.46, p = 0.0008) when we compared the mean headings in the training condition to those observed during the first visit with the new landmark. Measurements of flight headings after 10 and 20 or more subsequent trips revealed no decline in approach to the added element in both training arrays. After an additional 20 or more visits with the novel landmark present, the mean heading remained significantly different from the training condition for both the panoramic (t(70) = 2.95, p = 0.035) and 2-cylinder training arrays (t(36) = 4.18, p = 0.0016) (Figures 3-2d and 3-3d).

Within the panorama treatment we did not find a significant difference when we compared the headings of the training condition with the headings of bee flights after the novel landmark was removed. This suggests that, at least in the panorama landmark group, bees shift back to a more goal-directed flight path when no off-route landmark is available. **Figure 3-4** provides a summary of heading changes for both experiments over each condition.



Figure 3-2. Experiment 1 panorama headings plan view. The arrows within the semicircle represent the mean heading vector at 50 cm from the entrance a.) Training tracks and headings. The square is the location of the feeder hole and the arc represents the panoramic landmark. The panorama was a piece of curved Plexiglass® 11 cm in height and covered in black and white Julesz pattern (1 cm squares). b) First trip with new landmark. The added landmark was a black cylinder measuring 10 cm in diameter and 24 cm in height c.) ~ Tenth visit with new landmark present. d.) Twenty or more visits with new landmark present e.) Landmark removed after ten visits.



Figure 3-3. Experiment 1 2-cylinder headings. The arrows within the semicircle represent the mean heading vector at 50 cm from the entrance **a**.) A plan view of training array where the square is the location of the feeder hole and the two filled circles represent the training array. Training cylinders were flat black with a diameter of 10 cm and a height of 24 cm. **b**.) First trip with new landmark. The added landmark was a black cylinder measuring 10 cm in diameter and 24 cm in height **c**.) ~ Tenth visit with new landmark present. **d**.) Twenty or more visits with new landmark present.



Figure 3-4. Experiment 1 heading summary. A summary of heading changes when the training array was the panorama (open squares) and 2 cylinders (filled circles). Angles were transformed into a ratio of the angular distance from feeder to the novel landmark in order to provide a common angular measure for the panoramic and 2-cylinder training groups. Error-bars show the standard error of each mean. The mean angular distance of the training condition was statistically different from the 1st trip and 20+ trip means in both the 2-cylinder and panoramic treatments. The 9-11 trips condition was not compared to the training mean for either 2-cylinder or panoramic treatments. The mean angular headings of the panoramic training and removed conditions were compared and found insignificant. There were also no differences between the panorama and 2-cylinder bees for any experimental condition.

Path length

The length of the straight goal path and the dog-leg path from the entrance to the added landmark and then to the goal, was a different length in the panorama than it was in the cylinder training conditions. In order account for these differences and provide a common measure for comparing between training conditions we applied the following data transformations. First, for each path length value we divided the total path length

value by the length of the straight goal path in each training condition. This accounted for the longer distance between the entrance and the feeder hole during the 2-cylinder conditions than in the panorama conditions. An additional adjustment was necessary to account for the position of the added landmark since the added landmark was placed farther from the straight goal path in the panoramic condition than in the 2-cylinder condition. Here we calculated the ratio of the dog-leg goal path to the straight goal path for each condition. We then reduced the path length values for panorama trained bees by the percent difference in these two ratios. Statistical comparisons were performed using these values rather than the raw path lengths of each bee. We used t-tests to perform a series of eight comparisons; 1-2.) training path lengths vs. first trip with the novel landmark within both panorama and cylinder trained groups, 3-4.) training path lengths vs. 20+ trips with the novel landmark within both the panorama and cylinder trained groups, 5.) training path lengths vs. landmark removed within the panorama trained group, 6-8.) path lengths of the panorama vs. cylinder trained bees for the training, first trip and 20+ trip conditions. The Bonferroni method for correcting for multiple comparisons was applied to all the p-values reported throughout the results section.

The total distance traveled from the goal to the feeder varied little throughout our manipulations for both treatment groups (**Figure 3-5**). In the 2-cylinder treatment path length did not vary with the initial addition or continued presentation of the off-route landmarks. The total distance bees traveled to reach the goal was approximately 1.5 times the length of the straight path whether the novel landmark was present or absent.

Panorama trained bees had a statistically longer path length than 2-cylinder bees during the first trip to the novel, off-route landmark (t(89), p = 0.034). We found no

other significant differences in mean path length in any comparisons either within or between the panoramic and 2-cylinder training treatments. Although it appears that the path length of panorama bees went down after subsequent training with the new landmark, the mean path lengths for the 1st, 9-11 trips, or 20+ trips conditions were not statistically different from the training condition. **Figure 3-6** shows the actual flights of panorama and 2-cylinder bees to the training array and on their first trip with the novel landmark.



Figure 3-5. Experiment 1 path length summary. A summary of path length changes when the training array was the panorama (open squares) and 2 cylinders (filled circles). Distances were transformed into multiples of the calculated straight path to the feeder hole to provide a common measure for the panoramic and 2-cylinder training groups. Error-bars show the standard error of each mean. The panorama and 2-cylinder treatments were found to be statistically different in only the 1st trip condition after conducting a Bonferoni correction for multiple comparisons.



Figure 3-6. Experiment 1 flight paths. a.) Panorama- training. b.) Panorama- first trial with added landmark c.) 2-cylinder- training d.) 2-cylinder- first trial with added landmark. In the panorama conditions the distance of the added landmark from the straight path was larger than others. Path lengths were normalized to prevent this from confounding our results.

Experiment 2

We conducted a second experiment to see how the nearness of a novel landmark to the goal path affected the approach response. For this study we changed the initial training landmark configuration by shifting the feeder position and the landmark array approximately 25⁰ laterally. This allowed us to position a novel landmark farther from the training array than possible in previous experiments. We trained bees with two black cylinders near the feeder hole and tested in four conditions; a.) with only the training array, b.) with a novel landmark, identical to the elements of the training array positioned beyond the training array and 36 degrees to the left of the straight path, c.) with an

identical novel landmark added 56 degrees to the left of the goal path but before the training array. d.) with a half-sized novel landmark positioned 38 degrees to the left of the straight path, half way to the feeder hole (**Figure 3-7**). We reduced the size of the novel landmark in manipulation d because we wanted to be confident that the approaches in condition c were due to the angular proximity of the novel element to the straight path and not its visual size as seen from the entrance. Since no normalization was necessary between training and testing conditions we used the raw angle and path length data for statistical comparisons. We applied a Bonferroni correction for three comparisons to all reported p-values.

Bee flight directions in the training condition clustered in the direction of the feeder while headings in the four addition manipulations tended to be more scattered. The landmarks added nearest and farthest from the straight goal-path failed to elicit enough approaches to achieve a mean heading significantly different from training (**Figure 3-7** conditions **b** & **d**). However, the mean heading for the large and off-route novel landmark condition (**Figure 3-7c**) was found to be statistically different from training values by a Watson/Williams two sample test for angular data ($F_{1, 14} = 7.82$, p = 0.043). Path lengths remained around 1.5 times the straight path in all the added landmark manipulations and were not found to differ from the training condition (data not shown). These results suggest that size is the dominating factor for eliciting an approach response from honey bees.



Figure 3-7. Experiment 2 headings. The arrows within the semicircle represent the mean heading vector at 50 cm from the entrance **a**.) Training array **b**.) The added landmark was the same size as the training elements and placed far off route beyond the training array. **c**.) The added landmark was the same size as the training elements and placed far off route before the training array. **d**.) The added landmark was one half of the size of the training elements, placed nearer the goal path and before the training array

Experiment 3

When we removed the off-route landmark in the panorama condition at the end of experiment 1 we noticed that bees immediately shifted back to a flight heading centered on the feeder (**Figures 3-4 & 3-5**). This result suggested that the approach behavior did not result in bees learning a fixed dog-leg flight path which we might have expected based on other experiments (Graham et al., 2003; Collett et al., 1993). We conducted experiment 3 as a follow to determine if training with 2-cylinder array rather than a panorama produced similar results. In this experiment we presented the en-route landmark *throughout* training giving bees in excess of 60 visits in the presence of the off-

route landmark. This time we trained bees in the presence of the off-route landmark from the start to see if greater experience with their dog-leg flight paths would result in a learned motor pattern.

The off-route landmark was identical in size to those near the feeder and located 29⁰ to the left of and approximately half way along the straight path to the feeder hole (**Figure 3-8a**). Testing was done in the training configuration and during bees' first trip with the off-route landmark removed.



Figure 3-8. Experiment 3 angle summary a.) Training array which included an en-route landmark. b.) Flight headings of bees when the off-route training landmark was removed.

To determine if well trained bees preferred a heading directed at the off-route landmark rather than the feeder we compared the training headings of experiment 3 (Figure 3-8a) with the training headings of experiment 1 (2-cylinder training Figure 3-3a). This comparison is justified since the positions and sizes of the near goal landmarks were identical in both experiments. We found the mean headings in these training conditions to be statistically different by a Watson\Williams two sample test for angular data ($F_{1, 43} = 5.23$, p = 0.027). When we removed the landmark approximately 66% of the bees oriented their flight to the feeder hole. This implyies that most subjects did not maintain a fixed, dog-leg motor pattern once the en-route landmark was eliminated. However this shift was not found to be significant. We attribute the lack of statistical significance to the low sample size in the removal condition.

Discussion

Numerous experiments have been conducted to identify which visual features of the environment are important for honey bee navigation both nearby and en-route to a goal. A variety of visual cues including celestial information, motion cues, the color and shape of the goal as well as the sizes, shapes, colors and spatial relationships of nearby landmarks have all been shown to contribute in the orientation and execution of both outward and homeward journeys of honey bees (Collett & Collett 2002 for review). And although several authors have proposed mechanisms for how honey bees might deploy their knowledge of visual information to attain the goal (i.e. Wehner & Rossel, 1983; Cartwright & Collett, 1983; Anderson, 1977; Fry & Wehner, 2005) there have been only a few models of how honey bees might acquire such information in the first place (i.e. Zeil et al., 1996; Lehrer, 1996; Lehrer & Bianco, 2000). To this end, this set of experiments was designed to identify some of the factors that cause bees to spontaneously attend to and subsequently incorporate a new landmark along a familiar route.

In experiments where two distinct cylinder landmarks indicated the goal position, the data suggest that the amount of attention a landmark received was a function of its visual size as seen from the starting point. An off-route landmark reliably elicited approaches when it appeared larger in the visual field than the landmark elements near the feeder.

This was the case even when the off-route landmark was quite eccentric to the straight goal path. The eccentricity of the off-route landmark ranged from $29-40^{\circ}$ in the 2-cylinder condition of experiment 1 and was positioned 56° from the straight goal path in experiment 2c. In all of our experiments the landmarks were designed such that bees had a simultaneous view of the training array and the off-route landmark from the entrance platform of the arena. What is note-worthy is that bees approached the off-route landmark even though they had an opportunity to observe a rather large angular disparity between the off-route landmark and the landmarks located near the feeder hole prior to beginning their flights (**Figure 3-7c**).

In contrast to the two cylinder results, a novel off-route landmark also elicited approaches when the training array was an extended landmark that subtended a very large visual angle (**Figure 3-2**). This result appears to break with the "size dictates approach" phenomenon observed in all of the two-cylinder training experiment. From a bee's point of view the panorama would have appeared as a thin gray band due to its relatively short stature (11 cm tall) and random julesz coloring. Conversely, the 24 cm tall, solid-black off-route landmark would have had a very different appearance against the white arena background. Perhaps the added landmark's visual distinctness from the training array was responsible for drawing attention during the panorama treatment. Another possibility is that there is an upper limitation on how much size can influence honey bee flight trajectories where very large landmarks such as the panorama are to big to elicit directed approach responses. The path length results in **Figure 3-5** show that panorama bees spent significantly more time exploring the added landmark than bees trained with a two-cylinder array. Although this difference in path length could be construed as

evidence that bees treated the cylinder as something new and different from the panoramic training array, additional experiments are necessary to support a visual distinctness explanation of the approach behavior.

Although our findings suggest a role for landmark size in modulating approach behavior, it is likely that the distance of a landmark from the straight path also plays a role in determining what bees adopt as an en-route landmark in larger, more natural foraging conditions. Early observations by Frisch (1967) suggested that even over long distances bees will deviate from the straight goal-path to incorporate conspicuous landmarks into their route. Later studies showed that bees will ignore these familiar enroute landmarks when they are experimentally shifted to a position more than 25-30⁰ lateral to the previously learned goal route (Chittka & Geiger, 1995; Chittka et al., 1995). Within the confines of our experimental apparatus, bees always approached a salient landmark no matter how far it led them off of their training heading. At least two methodological differences could explain the discrepancy between our data and previous findings.

In the previous studies, bees were trained over hundreds of meters of countryside in the presence of celestial cues that provided bees access to high quality path integration information. During tests the heading of the path integration vector directly conflicted with the path dictated by the shifted landmarks. Path integration during training in our spatially restricted and visually sparse arena may have provided little assistance in selecting a goal directed heading. Even if bees had poor or no path integration information in our arena, they still had a view of the near-goal landmarks no matter their position in the arena. The accuracy of mean flight headings toward the feeder hole in all

of our training conditions as well as previous studies done in a similar arena environment suggest that under these circumstances a view of goal related landmarks is sufficient to produce properly directed approach to and search behavior at the goal area (Fry & Wehner, 2002).

Differences in the spatial scale of ours and previous experiments might also account for the behavioral differences in another way. The lack of differences in path lengths found in nearly all of our experiments, suggests that honey bees experience little or no increase in flight distance when approaching our off-route landmarks. In the single case where path lengths initially increased with the introduction of the landmark, (panorama condition **Figure 3-5**) over subsequent visits, bees quickly settled on a path that was no longer than they took to the training array. This is striking because their initial heading deviated from the straight path and yet their overall path length was no longer. Deviating to eccentric beacons over longer foraging distances might yield much larger time and energy costs. Under these conditions one might expect a trade off between the navigational value of incorporating an en-route landmark and the costs accrued during large deviations from the straight path to the goal.

Whether an off-route landmark was newly introduced to a familiar scene or present from the onset, it always received attention based size and possibly visual novelty. This shift in heading toward the more salient off-route landmark persisted for as long as we continued to train and test bees in its presence (**Figures 3-4, 3-8**). However, when the landmark was removed, bees trained with a panoramic landmark retained no tendency to head in the direction of or visit the location that the off-route landmark previously

occupied. The same trend probably holds for 2-cylinder trained bees although the low sample size prevents us from confidently drawing this conclusion (Figure 3-7, 3-8b).

These results contrast with the findings in a similar study performed with wood ants. Using an arena setup akin to ours, Graham et al. (2003) found that wood ants were also attracted to an off-route cylinder landmark positioned between the starting location and the goal. In a manipulation where the landmark was removed, subjects continued to choose a path that included an approach to the missing landmark. The authors suggest that the intermediate landmark provided a cue which ants used to learn other visual information in the scene, such as light and shade gradients created by the various folds and boundaries of the curtain that enclosed their rectangular arena. They hypothesize that ants used the subtle remaining cues to guide movement in the absence of the original landmark.

This might be an attractive strategy from a machine learning perspective since an agent could use particularly salient environmental stimuli to initiate a "closer look" at other useful but more subtle perceptual information to define important states in their environment. Such a strategy may be especially useful for insects and robots that possess relatively low resolution perception. Large landmarks like trees and rocks that occur in different locations along a route, although very salient, may appear similar to the low resolution visual system of insects as well as the sonar and infrared of robot navigators. A strategy of using large salient landmarks as cues to learn less obvious, but more defining state information would allow an agent to reduce the amount of information it needs to store while retaining reliable information at critical choice points. The navigating agent could initially divide its environment into a series of states each defined

by a salient but perhaps ambiguous landmark. When the agent detects the presence of each of these landmarks along its path, it could then do a more detailed matching of other subtle cues in the vicinity to determine its current location more reliably.

In our arena we did not find evidence that bees used a learning strategy similar to that reported by Graham et al (2003). This is likely may be due to differences in our methodology. In our experiments, the round arena and homogenous lighting conditions provided very little in the way of large-scale spatial patterns for bees to associate to the location of the off-route landmark. In addition, Graham et al (2003) did not provide the ants with any landmarks near the feeder. With the panorama or two large cylinder landmarks consistently designating the position of the reward, bees may not have found it necessary to encode the location of the off-route landmark at all.

Cheng & Spetch (1999) reported a similar phenomenon in a different arena study. In their experiment they trained bees to first find reward to the left of a single, blue, cylinder-landmark. This was followed a subsequent period of training where they placed an additional yellow, cylinder-landmark to the left of the reward such that the food was positioned equidistant from both cylinders. During tests they switched the relative positions of the landmarks, placing the blue cylinder to the left and the yellow cylinder to the right of the food. Under these conditions bees searched more often at the position designated by the original blue landmark even though it lead them outside of the 2element landmark array. The authors suggest that the initial training with the single blue cylinder blocked the acquisition of additional spatial information given by the location of the yellow landmark during the second round of training. Although our bees approached the en-route landmark while it was present, perhaps a similar blocking mechanism

prevented them from acquiring a permanent dog-leg path when the en-route landmark was removed.

Another possibility is that the approach behavior we observe is not under the direct control of a honey bee decision-making process. Perhaps it is an innate and spontaneous attraction elicited in the presence of any salient landmark bearing certain visual properties. Other spontaneous behaviors have been observed in honey bees such a propensity to fly parallel to contrast borders (Lehrer et al., 1985) and an innate preference for particular colors (Giurfa et al., 1995) and patterns (Lehrer et al., 1995) while foraging. Under the proper conditions, honey bees are able to alter all of these spontaneous tendencies in order to achieve the goal. If further inspection reveals that this behavior is instinctive as we expect, it will be interesting to see if it is also subject to alteration under the right set of circumstances.

i

Chapter 4

Honey bees use an en-route landmark as a resetting point for search flights.

Introduction

It is known that insects use path integration as well as conspicuous intermediate landmarks to guide navigation to distant foraging sites and back to the nest (reviewed by Frisch, 1967; Dyer, 1998; Collett & Collett, 2002, 2004). Path integration (dead reckoning) is the process of continually computing ones current net distance and direction from a point of origin (reviewed by Collett & Collett, 2000). In the case of central-place foragers the point of origin is usually the nest. A honey bee calculates the distance portion of a path integration vector by summing the amount of texture passing across her visual field during flight (Esch & Burns, 1995; Srinivasan et al., 2000). The direction component is guided by patterns of polarized light in the sky and the position of the sun while compensating for its movement throughout the day (reviewed by Rossel & Wehner, 1986; Dyer & Dickenson, 1996). The advantage of path integration is that it provides a means of wayfinding when the topography of a landscape is unfamiliar. Using dead reckoning, an insect in search of new foraging patches can set an immediate homeward vector from a new and perhaps unfamiliar feeding location. Honey bees are in the unique position to use path integration to visit new and unfamiliar outbound foraging locations because experienced foragers communicate path integration coordinates of food via the dance language (Dyer, 2002). By observing a waggle dance inside the hive a

90

÷
forager honey bee can retrieve the distance and direction of a rewarding patch from a fellow forager and set an outbound vector for this site even if it is in an area she has never experienced during previous foraging bouts (Frisch, 1967; Dyer, 2002).

Path integration is certainly an efficient and ubiquitous form of insect navigation (Collett & Collett, 2000; Collett & Collett, 2002) but it is susceptible to the accumulation of errors in the distance measure (Srinivasan et al., 1997) and is difficult to implement in overcast weather when celestial cues are not available (Dyer & Gould, 1981). Although bees that are unfamiliar with a feeding site have only path integration initially, they quickly learn visual landmarks that demarcate the location of their nest (Tinbergen, 1972; Capaldi & Dyer, 1999) and familiar foraging routes (reviewed by Dyer, 1998; Collett & Collett, 2002). Several previous experiments show that honey bees learn to visit a series of intermediate landmarks on their way to a familiar location. Evidence suggests that during learning, bees generate a series of path-segment memories called local vectors where the appearance of each conspicuous intermediate landmark is linked with the distance and to the next landmark in the sequence or to the goal when she is near her destination. (Chittka, 1995a, b, c; Collett et al., 1993, 2002; Collett & Baron, 1995; Srinivasan et al., 1997; Zhang, 2000). In these studies bees preferred to execute the local vector associated with the most recently experienced landmark on their trip even when the landmark appeared in the wrong compass direction (Chittka, 1995a, b; Frisch 1967), at the wrong distance from the nest or previous landmark (Chittka, 1995; Srinivasan et al, 1997) or out of place in the learned sequence (Collett et al, 1993). Although these results suggest bees rely heavily on conspicuous visual landmarks to guide their way in familiar terrain, they also provide data showing that bees will use other, more global information

such as celestial cues, distance information provided by the path integrator or panoramic cues such as distant landmarks or changes in scenery when a local landmark appears grossly out of context.

Like these previous experiments, our own studies (see chapter 3) illustrate a strong tendency of bees to deviate from a goal directed path to approach conspicuous landmarks found along the way to their destination. Furthermore, results suggest that bees are inclined to approach off course, en-route landmarks regardless of whether they were available from the beginning of learning, or newly introduced to a previously learnt goal path (see chapter 3). Though the off-route landmarks in our experiments did not appear to impose a residual local vector upon their removal, the fact that bees have an irresistible attraction to very eccentric, albeit salient, landmarks led us to wonder at the functional value of such a response in the context of navigation

Previous work suggests that insects use en-route landmarks in at least three ways: 1.) as intermediate "checkpoints" for correcting errors accumulated by their global path integration system (Srinivasan et al, 2000) 2.) as cues for segmenting their path into a series of shorter, more precise trajectories called local vectors (Collett et al., 1993; Srinivasan et al., 1997; Collett et al., 2002; Collett & Collett, 2004) 3.) as cues to facilitate the learning of additional visual information along the path (Graham et al., 2003).

From a sequential decision-making perspective the first strategy suggests a means for the agent to cross check perceptual information within a state in order to reduce uncertainty. The second two strategies are systems for associating states with actions

where the actions are the execution of a local vector or obtaining more perceptual information to help determine their current location.

Unfortunately, the results of two of the experiments supporting the local vector hypothesis allowed honey bees a view of the en-route landmark from the goal (Srinivasan et al., 1997; Collett et al., 2002). In such cases, bees may have sharpened their search by using the provided landmark in a template matching strategy (Cartwright & Collett, 1982) rather than executing a short and accurate local path integration vector associated with the landmark. To better understand how the presence of conspicuous, en-route landmarks affected honey bee search behavior we trained bees to locate a sugar water feeder in a long tunnel with or without an en-route landmark in place. We subsequently removed the feeder and observed how bees searched for the missing goal. To avoid the difficulties associated with interpreting the results of previous experiments we designed our apparatus such that bees were unable to perceive the landmark from the goal region.

Although our initial intent was to examine whether the presence of an en-route landmark improved search and if so how, the bees' behavior revealed that the landmark served a function that had not previously been identified. Specifically, bees used the landmark as an anchor point to launch repeated searches when they were unsuccessful in finding the goal. Therefore, the focus of our experiment and analysis became to observe and characterize this undocumented search strategy of honey bees.

93

÷

Methods

<u>Apparatus</u>

The tunnel measured 7.0 m in length and 11 cm in width with walls extending 21 cm above the tunnel floor (**Figure 4-1a**). We covered the walls and floor with panels of black and white Julesz pattern with a pixel size of 1 cm² to provide visual stability and a distance measure via optic flow for the flying bees (Srinivasan et al., 1997, Srinivasan et al. 2000). The top of the tunnel was covered with a thin layer of nylon mesh to contain our subjects. The whole apparatus was positioned on a stand in an open field which provided no view of exterior landmarks from inside the tunnel. During data collection we recorded search flights with a Sony model TRV140 8 mm/digital camera positioned 3.2 m above the tunnel floor. This height ensured that the camera would not be visible to the subjects as they flew through the tunnel. We used the camera zoom to get the largest frame of view while still allowing us to keep track of a bee's position during her search flight. This resulted in a viewable area of about 1.6 meters of tunnel centered approximately on the feeder position.

Tunnel Top View



Figure 4-1. *Tunnel* **a.**) Top and side views of tunnel with bin labels around the feeder and landmark. Each bin was 20 cm in length. The field of view of the camera spanned from the beginning of bin -4 to the end of bin 3 (about 1.6 m around the feeder). F and L mark the locations of the feeder and landmark during training and testing. **b.**) The landmark was constructed from a piece of PVC tubing cut lengthwise. The landmark formed a hump on the tunnel floor 5 cm in height with yellow facing the entrance and Julesz pattern facing the feeder.

Training and testing

For the tunnel experiment we used two separate groups of 20 bees from the same hive. Each group was trained, as per Frisch (1967), to a cryptic sugar-water feeder that was then moved down to a distance of 5.4 m from the entrance (**Figure 4-1a**). Bees visited the feeder for three days, giving each bee a minimum experience of 60 visits to the food. Over the first day of training we marked all of the visiting bees with a small dot of enamel paint. We marked each group of bees with a different color so as not to re-use the same individuals across experimental conditions. Marking bees on the first day of each experiment also ensured that test bees all had similar training experience. During training observers avoided standing near enough to the apparatus to become landmarks except during brief periods when refilling the feeder.

To control for odors during tests, we removed all Julesz pattern in the tunnel within 2 m of the feeder position and scrubbed the floor and walls with warm water. After applying new pattern to the tunnel walls and floor we let bees into the tunnel individually to video record their search behavior in the absence of the feeder. A test trial commenced when the bee began forward flight from the tunnel entrance toward the feeder and terminated when she returned to entrance. During tests an observer sat below the tunnel and watched the trial by peering into the entrance. This allowed us to observe, in a limited way, what the bees were doing when they were beyond the view of our camera. Each trial lasted from 1 to 3 minutes depending on how intensively the bee searched for the feeder. At the end of test trials we captured each bee and held it until the end of the experiment to avoid recording multiple trips by any one individual. Testing lasted one day at the end of which all captured bees were released near the hive entrance.

Data analysis

We transferred video from 8 mm tape to DVD format using Adobe Premier. The video analysis was done using Cyberlink PowerDvd. For scoring purposes we divided the tunnel into 20 cm bins centered on the feeder position (**Figure 4-1a** top view). The angle of the camera lens allowed us to view from bin -4 to bin 3 with the position of the feeder centered in bin zero. To score each bee we played the video back frame-by-frame and recorded the sequence of visits and time spent in each bin for the entire length of her trial.

Results

Our arena experiments in chapter 3 suggested that bees have a spontaneous and apparently irresitable attraction to salient landmarks placed intermediate to the goal location. We conducted the tunnel experiment to determine what advantages an intermediate landmark might provide over longer travel distances to a food source. To address this question we trained bees in two tunnel treatments. In the first, we trained and tested bees in the presence of a conspicuous yellow landmark placed at the half-way point of the tunnel 1.7 m from the feeder location (**Figure 4-1a** top view). We painted the entrance-side of the landmark half yellow and the feeder-side of the landmark half Julesz so that the yellow coloring disappeared from view shortly after bees flew over it (**Figure 4-1a**). This landmark, unlike the Srinivasan et al. (1997) and Collett et al. (2002) landmarks, was not visible from the location of the food. For comparison purposes we trained a separate group of bees to find the feeder in at the same position in the tunnel without a landmark present.

Search distribution

Our first analysis was to compare searching behavior between landmark and nolandmark bees. For this measure we used the initial 30 seconds of search starting from when each bee entered into bin -1 for the first time. We used only the first 30 seconds of flight in order to capture as much search behavior as possible while minimizing other behaviors. During tests, bees tended to land, crawl and attempt upward with longer unrewarded experience in the tunnel. Each three seconds of time spent in a bin was considered a visit and we totaled every visit from all bees to each bin for each condition.

97

ł

The totals in the unobservable bins -4/-8 and 3/9 were generated by dividing all offcamera search time by the number of off camera bins on either side of the camera. We chose bin -8 for the entrance-side limit because bees rarely visited bins -9 through -24 whether or not the landmark was present. This technique of analyzing honey bee search distributions differed from Srinivasan et al. (1997) and Collett et al. (2002) who analyzed the first two to three turns each bee made while walking along the side of the tunnel. In our experiment we used an over head camera to prevent the experimenter from affecting honey bee flight paths. This prevented us from conducting a similar turn-analysis since we could not accurately observe any turns bees made while outside the field of view of the camera.

Figure 2-2 shows that bees in the no-landmark condition have a nearly bell shaped search distribution centered over the feeder position while landmark bees searched almost uniformly across bins. This was surprising because previous studies showed a sharper search distribution when a landmark was present (Srinivasan et al., 1997; Collett et al., 2002). In addition, where bees spent their time when out of our camera's view was very different between treatments. Landmark bees spent 42% their off-camera time on the landmark side while no-landmark bees only 9% in the same tunnel region. Conversely, landmark bees spent 52% of their off-camera time at the end of the tunnel while no-landmark bees spent 52%. The difference between bee off-camera time was found to be statistically different by a G-test of independence ($\chi 2(1, n= 284)=72, p < 0.0001$). The overall search distributions for the landmark and no-landmark conditions were also statistically different (Chi Square goodness of fit analysis ($\chi 2(7, n=205)= 25.2$, p=.0007)).



Figure 4-2. *Tunnel search distributions*. A summary of where bees spent their time during the first 30 s of search after crossing into bin -1. In both the landmark and no-landmark conditions the feeder was located in the center of bin 0 throughout training.

Survivorship analysis

Figure 4-2 shows that having a landmark produced a more evenly distributed search near the feeder than if there was no landmark during training and testing. During tests we noticed that landmark bees performed very short bouts of back-and-forth search around the feeder location followed by successive trips back to the landmark. The distribution in Fig 4-2 shows all of this flight behavior and hence looks flat. In contrast, no-landmark bees flew to the feeder location and searched intently. Eventually they broke from a focused search and flew straight to the entrance or very end of the tunnel. We performed a survivorship analysis to quantify these differences Figure 4-3. Using a log-likelihood analysis, we found that bees in the no-landmark condition were quicker to visit the

÷

entrance or end of the tunnel than bees trained and tested with a landmark

 $(\chi^2(df=1)=7.183, p=0.007).$



Figure 4-3. Ends-of-tunnel visitation rates of bees trained and tested with a landmark (dotted line) and without a landmark (solid line) present in the tunnel.

Flight direction

The survivorship analysis, coupled with our observations during tests, suggested that bees trained with a landmark were more likely to begin a new attempt to find the goal from the landmark rather than from the ends of the tunnel. To further quantify this phenomenon we analyzed flight paths of bees in the goal region. We wanted to

1 : 4

determine how often bees flying toward the entrance continued on a straight flight off camera (toward the landmark) rather than to turning back toward the goal. We looked at how often bees that flew from bin -1 to bin -2 continued into bin -4 in both the landmark and no-landmark treatment groups (**Figure 4-4**). We found that landmark bees were statistically more likely to continue on a negative, landmark-directed heading (traveling all the way to the landmark in most cases) than bees trained without a landmark (G-test of independence: $\chi^2(1, n=75) = 12.89$, p=.00033).



Flight Direction Near The Goal

Figure 4-4. Proportion of bees that flew into bin -4 after previously flying from bin -1 to bin -2. To count as entering bin -4 bees had to fly directly from bin -1 to -4 and off camera without making a turn along the way. In addition, after entering bin -4 bees had to remain off camera for a minimum of three seconds.

Discussion

Bees trained to find food in a randomly textured channel with no view of landmarks searched with reasonable precision at the feeding location during unrewarded test trials suggesting that bees use their updated path integration vector with fair accuracy. This result is consistent with previous findings of Srinivasan et al. (1997), Collett et al. (2002), and Vladusich et al. (2005) who trained honey bees in similar tunnel environments. Bees trained with a landmark positioned en-route to the feeder seemed to rely heavily on the landmark during tests. This finding is also consistent with previous results to the extent that when a visual landmark is available bees use it extensively. However, our observations of search behavior in the landmark condition suggest that bees used the enroute landmark in a way that was different from any previously proposed strategy.

When Srinivasan et al. (1997) trained bees to find food in a randomly textured channel they found that bees searched more precisely for the food when trained and tested with an en-route landmark than without. The authors suggested that the sharper search behavior was the result of bees using the en-route landmark to correct the distance estimate of their global path integration vector. A different interpretation of the Srinivasan et al. (1997) result is that rather than using the landmark to correct their global path integration vector, bees split the tunnel route into two path segments bounded by the entrance of the tunnel, the position of the landmark, and the location of the feeder. In this case the bees used the to execute separate and shorter local vectors resulting in the tight, bell-shaped search patterns over the goal region. This interpretation is supported by other work dealing with the use of en-route landmarks by insects (see Collett & Collett, 2004 for review).

In the similar tunnel experiments of Srinivasan et al (1997) and Collett et al. (2002) the presence of an intermediate landmark produced search pattern centered neatly on the training position of the feeder. In contrast, it is clear that in our experiments the en-route landmark did not sharpen search near the feeder position as expected (**Figure 4-2**). Rather, bees trained and tested with a landmark intermediate to the goal searched in a nearly uniform distribution throughout the recordable range of our camera. What we

noticed was that bees used the landmark as waypoint to launch successive flights toward the position of the feeder. After a brief pass through the feeder location, bees returned to the landmark as if to begin anew. This behavior is characterized in Figure 4-4 which shows that landmark-bees tended to continue off camera toward the landmark when their flight was headed in that direction. No-landmark bees were much more likely to turn back to the feeder location under the same circumstance. In addition **Figure 4-3** shows that bees trained and tested with an intermediate landmark were slower to leave the middle and visit the end or beginning of the tunnel than no-landmark bees. The tendency of no-landmark bees to "time-out" and travel to the ends of the tunnel more quickly than landmark bees may be evidence that the intermediate landmark gave bees a greater certainty that they were in the correct general location of the goal so long as they could continually re-check the position of the landmark. This sort of resetting behavior and overall dependency on an en-route landmark was not reported in any of the previous tunnel experiments designed to determine the priority or interplay of path integration and visual landmarks during honey bee navigation. In our experiments bees seemed to ignore their global path integration vector and engage in repeated executions of a local search flights from the position of the en-route landmark.

There is a key difference between our experiments and previous work that may have led us to observe this resetting behavior where others did not. As stated in the methods our en-route landmark was virtually invisible from the feeding location. From the feeder bees could only see the Julesz back-side of the landmark which stood only 5 cm high, approximately 2^0 degrees of visual resolution. The small visual angle, barely within a bee's visual limit, and blending pattern should have rendered the landmark undetectable

103

i

from near the feeder location in our experiment. In contrast, the landmarks in earlier studies were not camouflaged and were well within the bees' visual resolution from the training position of the feeder. We suggest that in earlier studies bees may have had an opportunity to use a visual memory of the landmark's appearance from the goal to sharpen their search. This is a well known strategy of bees and other insects for returning to familiar foraging locations (Cartwright & Collett, 1983; Collett & Reese, 1997; Cheng 1999a, b). Here we suggest another strategy for using visual landmarks en-route to a goal site. Our data suggest that this "re-setting" strategy provides an online means for bees to ensure they are on the correct path. Having increased certainty about their current location seems to allow bees to persist longer in the goal location before returning to the nest. The resetting strategy does not necessarily undermine the local vector hypothesis. In fact, a bee returning to the landmark, must at a minimum, have some knowledge of where to find the landmark in the tunnel from its current position near the food. Bees could use very well use global path integration information or a local vector to keep track of the position of the landmark.

Local vectors are potentially attractive from a sequential decision-making perspective as a strategy for breaking down complex and continuous environments into a discrete set of states and actions. However, the resetting behavior we report herein seems more akin to information gathering strategies such as learning flights (Lehrer, 1991, 1993; Capaldi & Dyer, 1999) and the spontaneous attraction to salient landmarks (Graham et al., 2004) that insects use to collect new information about important environmental locations. The difference between resetting and these other strategies is that bees appear to use the enroute landmark to gain a more certain assessment of their current state rather than to learn new information about their environment.

÷

Chapter 5

How honey bees use of landmarks near the goal: an examination of the role of snapshot guidance.

Introduction

Many species of central place foraging insects use visual features of the terrain to return to important locations in their environment. Since Tinbergen's (1972) classic experiments showing that the digger wasp *Philanthus triangulum* learns visual cues to locate its burrow, there has been a great deal of work to better understand what visual properties of landmarks insect rely on for navigation as well as how these properties might be used to revisit various locations in the environment (Frisch, 1967; Collett, 1996; Wehner et al., 1996; Dyer, 1998; Collett & Collett, 2002).

Based on early behavioral evidence in ants, (Wehner & Raber, 1979; Wehner, 1981) hover flies, (Collett & Land, 1975) and honey bees, (Anderson, 1977; Cartwright & Collett, 1979, 1982), Collett & Cartwright (1983) proposed a mechanism of navigation whereby insects guide their return to a location by sequentially matching their current view of the environment to a two dimensional eidetic memory (snapshot) of landmarks learned at the goal. Since it was proposed, this model has inspired an extraordinary amount of theoretical and experimental work about the contents and use of snapshots for wayfinding in insects as well as being cited as a potential strategy for navigation in a myriad of other mobile organisms (*insects:* Collett & Baron, 1994; Collett & Rees, 1997; Cheng, 1998, 1999a, b; Graham & Collett, 2002; Durier et al., 2003. *mammals:* Biro &

Dawkins, 2003; Manser & Bell, 2004. *birds:* Dawkins & Woodington, 2004. *fish:* Schuster, 2002.). The model aroused such interest because it could explain the apparent ability of animals to orient to a goal in a flexible way, e.g. along novel approach paths.

Thus far when discussing insect navigation in a sequential decision-making framework I have talked rather loosely about how honey bees and other insects might couple actions to environmental states. What makes the snapshot model so attractive is that it explains in detail how visual landmarks might generate actions in the form of steering commands. In addition, the snapshot hypothesis presents a tractable way for an agent to generalize a single memory, namely their snapshot, to the continuously changing view of landmarks they experience within any given environmental state. For this reason the snapshot model has also been implemented in a number of simulations and robotic applications by computer scientists (Trullier et al, 1997; Franz et al., 1998; Bianco & Cassinis, 2000; Lambrinos et al., 2000; Moller, 2000; Lambrinos, 2003).

Over the years several lines of evidence have emerged suggesting that honey bees and other hymenoptera do indeed use some sort of template matching strategy to direct their navigation when near a familiar goal location. The first support for the snapshot model stems from the initial studies of search distribution data by Cartwright & Collett (1982). In these experiments the authors demonstrated that when confronted with a single, cylindrical landmark, honey bees seem to use the retinal image size of the landmark to locate the goal position. They provided evidence that honey bees can rely on interlandmark angles rather than individual landmark sizes to pinpoint the goal when the training array is composed of multiple landmarks. In a follow up paper, Cartwright & Collett (1983) formulated a computational model of snapshot navigation. The algorithm, which was refined by testing the performance in a simulation task, produced behavior that was consistent with earlier results and also made predictions about how bees should guide their approach to the goal from more distant locations.

In addition to providing predictions of movement patterns in the vicinity of the goal, the model also raised the question of how bees might anchor their snapshots to the external environment. If an organism is to match its memory to a given visual scene it must somehow align its internal snapshot in a compass direction consistent with the expected orientation of landmarks at the goal. The authors suggested two possible solutions to this problem. The first was a retinotopically based memory where bees approach the goal from a consistent compass direction and adjust their body axis such that the current view of the landmarks falls upon the same retinal positions experienced during the "taking" of the snapshot. The alternative hypothesis was that bees associate a compass heading to the snapshot which is abstracted away from retinotopic coordinates. In this scenario the snapshot representation maintains its compass heading the way a dashboard compass stays in line with magnetic north. This sort of "free-floating" template memory should allow more flexibility since bees could use it to recognize and match their memory to a current visual scene from any body axis orientation.

In later experiments Collett & Baron (1994) and Collett & Rees (1997) provided evidence that bees and wasps use a retinotopic rather than a free-floating snapshot memory to return to a learned foraging location. The main support for retinotopy came from Collett & Baron (1994) who showed that bees prefer to align their body axis in a consistent direction during repeated visits to a foraging location. In general, bees preferred to approach a goal position while facing south rather than each bee settling on a

different but personally consistent compass heading. This set of experiments also showed that when near the goal bees can use the earth's magnetic field to maintain this consistent body axis orientation. Although previous results do not provide direct evidence of eidetic memory, they do suggest that honey bees (Collett & Baron, 1994) and wasps (Collett & Rees, 1997) possess a flight strategy that could potentially simplify the use of visual landmarks during template matching.

Another indirect line of evidence for a snapshot strategy comes from studies of learning flights called turn-back-and-looks (TBLs) that both honey bees and wasps perform after receiving a reward at a new foraging location. The TBL is a highly stereotyped series of arcing and widening lateral flights during which the insect acquires position and depth information of surrounding landmarks (Lehrer, 1991, 1993; Zeil et al., 1996). TBLs are thought to play a significant role in the selection and memorization of landmarks near the goal and are the putative mechanism for taking a snapshot. Collett & Baron (1994), Zeil et al. (1996) and Collett & Rees (1997) showed that bees and wasps, upon return to a learned food location, perform lateral flight maneuvers similar to those seen during a departing TBL. These lateral flights are believed to help establish a body alignment and view of landmarks consistent with that seen during the learning flight and therefore aid in obtaining a match between the world and the memorized snapshot. Again these findings corroborate but do not explicitly test the notion of snapshot memory use in honey bees.

Despite these later experiments and the enthusiasm generated from outside the insect navigation community, it is clear that the original version of the snapshot algorithm does not adequately explain how insects guide their final flight to a familiar goal location. The

only hard evidence for the use of an image matching mechanism in honey bees comes from the early search distribution data of Cartwright & Collett (1982) and a similar set of experiments performed by Cheng (1999a, b). All of these data are consistent with the use of snapshots when insects are close to the goal. However, the original model proposed by Cartwright & Collett (1983) also makes clear predictions about how insects should steer their flight during their approach from *outside* a landmark array. Unfortunately neither the original nor subsequent evidence suggest a role for template matching other than for a tightly restricted area very near the goal (Collett & Baron, 1994; Collett & Rees, 1997; Cheng, 1999a, b; Fry & Wehner, 2005). To date no one has directly compared the predictions of the Cartwright & Collett (1983) model with approach flights of honey bees trained with landmark arrays equivalent to those used for the original simulations and over the spatial scale that these simulations tested.

The goal of our experiments was to evaluate honey bee flight paths under similar landmark conditions as those of Cartwright and Collett (1979, 1982). To this end we used detailed tracking data to compare the direction of honey bee approach flights with predictions generated by the snapshot model. In order to provide a fair evaluation we used landmark arrays that were virtually identical to those used in the early experiments and simulations that generated the original algorithm. We also inspected bees' body axes during these flights to check for consistent body alignment and lateral, scanning flight near the goal as observed by Collett & Baron (1994) and Collett & Rees (1997). This allowed us to evaluate the snapshot model in a more comprehensive way than has been previously done.

A further feature of our approach is that bees were denied the opportunity to perform a learning flight, or TBL. This allows us to evaluate the necessity of this learning procedure for acquiring spatial information near the goal.

Method

Apparatus

Our apparatus, modeled after a design developed by Fry et al. (2000), consisted of an indoor arena with a camera positioned above to track bee flight paths (**Figure 3-1a, c**). The arena floor was 200 cm in diameter with a 22 cm high wall running continuously around the perimeter. A white, nylon tent extended up another 178 cm from the edge of the wall with the camera positioned into a hole sewn in the center of the ceiling. The floor and wall of the arena were painted flat white and the wall had a 2 cm in diameter entrance hole centered 11 cm from the arena floor. A 1.2 meter, clear, plastic entrance-tube extended from the arena wall into an outdoor flight cage which housed the nest. The arena-end of the entrance tube was fitted with an elbow that terminated in a small 4 x 4 cm square platform. This platform encouraged bees to fly rather than walk from the entrance into the arena. To provide approximately homogenous lighting inside the arena we spaced four 500-watt halogen work lamps evenly around the outside of the tent.

We tracked bee flight paths in the arena using a pan-tilt Sony surveillance camera and BIObserve Trackit software (Fry et al., 2000). By using a bee's contrast against the white arena background, Trackit recorded the bee's position and body axis in the arena at 60 Hz. Bee track paths were plotted and analyzed using custom programs written in Matlab.

111

ł



Figure 5-1. Arena **a.**) Overall dimensions **b.**) Pre-training feeder box **c.**) Detailed side-view of arena. The arena is a one way system. Bees walk into the arena via the entrance tube, fly from the entrance platform to the feeder-hole and crawl into the feeder box below. After feeding they walk out to the flight cage via the exit tube and return to the nest.

Training and Testing

Before we let a group of bees into the inside of the arena they first underwent a preliminary training process. During preliminary training bees learned, step by step, how to crawl into a feeder hole, through a one-way door and how to leave through an exit tube after feeding. We conducted the pre-training in the flight cage for approximately two

÷

days using a replica of the one-way feeding system found below the arena floor (**Figure 3-1b**). By the end of pre-training, bees were proficient at maneuvering through the apparatus and could be let into the arena.

Inside the arena, bees needed to fly from the entrance platform to a 1.5 cm in diameter hole in the arena floor using the landmarks provided. During the first few hours of training bees often got lost in the arena, requiring us to release disoriented bees from the tent periodically. When we had a group of at least 40 bees cycling through the arena we captured each bee and paint marked it for individual identification. At this time we also scent plugged each bee by applying a mixture of beeswax and bow rosin over the Nasanov gland on its abdomen (Towne & Gould, 1988). Plugging this gland prevented bees from emitting pheromone that might otherwise be used to aid navigation.

Training continued for and additional three to five days to ensure that each bee had a minimum of 60 visits to the landmark array inside the arena. We started each experiment with such a large group of subjects because many bees stopped returning over the long periods of training often required. A starting group that ranged between 40-50 bees usually dwindled to around 10-20 by the end of an experiment.

At the end of training we allowed bees into the arena one at a time for testing using a door in the entrance tube. During tests the feeder-hole was covered with a small sheet of white paper and each bee was tracked during the first 1 to 3 min of search flight. After the tracking period she was placed in the feeder box to receive reward and exit to the nest. Testing in all experiments began by tracking bee flights with the landmarks in their original training configuration. We then conducted follow up tests where we manipulated some aspect of the original landmarks. Delays created by test trials tended

113

÷

to slow down visitation to the arena so we interspersed testing sessions with bouts of group training to keep motivation high.

To further control odors during testing we placed a fresh, clean feeding apparatus under the arena and thoroughly washed the entire arena floor with warm water. In addition we replaced training landmarks with clean, identical testing landmarks.

Results

<u>Experiment 1</u>

The goal of these experiments was to ensure that bees used similar landmark information as those in Cartwright & Collett (1982) studies. This was especially important to test since bees in our arena were not provided the opportunity to perform turn-back-and-look learning flights (Lehrer, 1991, 1993). When given a single nearby landmark, do bees locate the goal using the landmark's image size? When given a symmetrical array of three landmarks, do bees locate the goal using inter-landmark angles rather than the image size of individual elements of the array?

During training, a single matt-black, cylindrical landmark 12 cm high and 5 cm in diameter was placed 60 cm behind and offset slightly to the right of the feeder-hole location. We trained 20 to 40 bees for 2 full days after which we tracked their search behavior. We then exchanged the training landmark with one twice its size, located in the same position and tracked the bees during their first trip with the new landmark. **Figure 5-2a** shows that doubling the landmark's size results in bees searching at a location more distant from landmark than during training. This result is consistent with

114

i

that of previous work suggesting that bees will use retinal image size to judge distance in the presence of a single conspicuous landmark (Cartwright & Collett, 1982).

To determine if bees could use inter-landmark angles to guide their search for the goal we placed three identical cylindrical landmarks behind the feeder. Each landmark was positioned 60 cm from the feeder hole and 60 degrees from its neighboring landmark(s). We trained a new group of bees for two full days in the presence of the 3 landmark array. On day 3 we tracked bees' flight to the training array. We then doubled the size of the training landmarks and tracked the bees during their first trip with the new landmarks. The change in size of the array had little effect on search behavior. Although bees demonstrated a slight tendency to search somewhat farther away when we doubled the size of the array elements, this effect was much less pronounced when compared with the results of the single landmark condition (**Figure 5-2b**). In general these results are also consistent with previous findings showing that bees can use inter-landmark angles to return to a feeding location when provided with a multi-element landmark array (Cartwright & Collett, 1982). These experiments also show that the bees can learn such patterns without the benefit of learning flights.



Figure 5-2. Search distributions Plan view of arena. a.) Search distributions of bees trained with a single landmark and tested with the training landmark (solid contour, n=24) and a landmark twice the training size (dotted contour, n=12). b.) Search distributions of bees trained with three landmarks and tested with the training landmarks (solid contour, n=15) and landmarks twice the training size (dotted contour, n=13). Contours were generated by first dividing the arena into 20 cm² bins. We then calculated where among the bins each bee spent her first 10 seconds of search flight. The contour line is drawn around the top one third most visited bins for all bees in each condition. The small black circle (FH) designates the location of the feeder hole and the large black circles are landmarks (LM).

Experiment 2

The next set of experiments was conducted to look at the flight trajectory of bees during their approach to the landmark array before they began search for the feeder. Do the predictions of the snapshot model closely match the flight trajectories of bees approaching from outside the landmark array?

The training landmarks were set up exactly as in experiment 1b but using the large cylinder landmarks (24 x 10 cm) and with the array shifted 70 degrees clockwise (**Figure 5-3a**) to more closely match the array setup of the Cartwright & Collett (1982) experiments. Bee flights were first tracked with the landmarks presented in the training condition. We followed this by removing each of the training landmarks in turn and tracking bee approaches to the remaining two landmarks. Finally we added a landmark en-route. The landmark added for this manipulation was a 16 tall x 15.5 cm in diameter

coffee can painted flat black. Half of its surface was covered with white poster board so that it blended into the white arena floor and walls. This landmark was positioned so the black surface could be seen from the entrance but not from within the landmark array so it could not serve as an additional landmark during their search close to the goal. Earlier experiments show that bees have a spontaneous tendency to approach a newly added landmark (see chapter 3). We used this phenomenon to alter honey bee's normal flight path. We then observed their subsequent flight into the goal region in order to compare honey bee flight trajectories from a novel position with the predictions of the snapshot model.

Approaching the goal

To determine if bees were using sequential image matching to guide their approach to a landmark array, we implemented the snapshot algorithm (Cartwright & Collett, 1984) in MatLab (see caption **Figure 5-3** for details). The output of the algorithm is a vector field where each arrow gives the predicted flight direction if the bee is matching a snapshot taken at the goal to her current view of the environment (**Figure 5-3a**). The bottom left quadrant of **Figure 5-3a** exhibits a band of arrows that arc away from the nearest landmark toward the feeder. According to the model's prediction, bees entering this area should fly upward and away from the landmark toward the goal. Of the 20 bees tested with this configuration of landmarks, over 50% flew from the entrance toward landmark 1 before turning left to search at the proper location (**Figure 5-3b**). This initial trajectory toward this landmark required bees to fly perpendicular, and in some cases directly against, the flight path predicted by the model.



Figure 5-3 Model predictions and flight paths to a 3 landmark array. **a**.) Vector field predicted by the snapshot model (implemented in Matlab). This plot predicts the direction a bee should travel from each 10 cm² bin if she is using the template matching algorithm proposed by Cartwright & Collett (1984). **b**.) Actual flight tracks of the approaches of honey bees to the landmark array. Excessive looping behavior was removed in order to provide a clearer picture of approach trajectories. ET designates the location of the entrance tube and FH the feeder hole. *Our snapshot model was implemented in Matlab as follows:*

First, a snapshot of the landmarks was taken at the feeder location. This snapshot was represented numerically by the angular orientation and width of each landmark as they would fall on the retinal. The predictions of the model were determined in the following manner. First an image of the landmarks, containing the same information as the feeder snapshot, was taken at each of 400 locations separated by 10 cm increments. Next, the center of each landmark in the feeder snapshot was paired with the center of the closest landmark in the current view based on angular distance. From each pairing we summed the tangential vectors that rotated the landmark on the retinal snapshot to match the feeder snapshot. The resultant unit vector determined the direction the bee should move from each arena location.

Landmark removals

Several lines of evidence suggest that rather than using a snapshot, bees select a single landmark as a beacon to guide flight to the goal area from more distant locations (Collett & Baron, 1994; Collett & Rees, 1997; Fry & Wehner, 2005). Figure 5-3b appears to be consistent with these observations since a high proportion of our bees began their flights by heading first to the landmark nearest the entrance.

In a previous experiment Collett et al. (1993) showed that bees could be trained to find

a feeder by flying through a series of dividers. Bees had to fly through a series of holes

in these dividers which resulted in a zig-zag flight path through their apparatus. After a period of training, the authors removed the dividers and found that bees continued on their fixed zig-zag flight path even though they could have taken a straight flight directly to the goal. This suggested that bees can learn to navigate using a series of turning rules; a sort of fixed motor program. Similarly, bees in our arena flew an indirect path to the goal, first heading directly toward a single landmark followed by flying to the area within the three landmark array. To determine which array element(s) bees preferred to use as their initial beacon and to see if our bees also learned a turning rule associated with their indirect routes to the goal, we removed each of the three landmarks in turn and tested five different bees in each of the three removal conditions.

The first removal is the most informative (**Figure 5-4**). When the landmark 1 was removed none of the five bees approached the vacant area that the landmark occupied during training. Bees' straight, goal-directed paths in the absence of landmark 1 are not consistent with the use of a motor program learned through extended experience along the original dog-leg path (**Figure 5-3b**). It is also interesting to note that several bees flew to the center of the remaining two element array. This is consistent with the results of Fry & Wehner (2002, 2005) who found that bees had a tendency to fly to the center of mass of a landmark array when the individual array elements were identical in size, shape and color. In the other two removal conditions (**Figure 5-5b**, **d**) all of the bees again approached landmark 1. Under these conditions landmark 1 was visually larger than elements 2 and 3 when viewed from the entrance. These results are consistent with the data presented in chapter 3 confirming that bees have a strong attraction to those landmarks that appear largest in their visual field.



Figure 5-4 Model prediction and flight paths of bees after the removal of the nearest array element. **a**.) Vector field predicted by the snapshot model. This plot predicts the direction **a** bee should travel from each 10 cm² bin if she is matching the 3-element template learned in training to the 2-element array. **b**.) Actual flight tracks of the approaches of honey bees to the landmark array. Excessive looping behavior was removed in order to provide a clearer picture of approach trajectories.

i



Figure 5-5 Model prediction and flight paths of bees after the removal of the farthest array elements. $\mathbf{a} \& \mathbf{c}$.) Vector field predicted by the snapshot model. These plots predict the direction a bee should travel from each 10 cm² bin if she is matching the 3-element template learned in training to the 2-element arrays. $\mathbf{b} \& \mathbf{d}$.) Actual flight tracks of the approaches of honey bees to the remaining landmarks. Excessive looping behavior was removed in order to provide a clearer picture of approach trajectories.

Pushing bees off course

Chapter 3 showed that bees will spontaneously investigate a novel landmark in a familiar environment. We used this phenomenon to push bees off their normally traveled route in order to determine if they would subsequently employ an image matching strategy to guide the remainder of their flight to the reward location. Once off their path

i

but beyond the introduced landmark bees again tended to fly directly to the next, nearest landmark before turning toward the feeding location (**Figure 5-6**). Since the half of the introduced landmark facing the training array was covered in white paper and was not visible to the bees, it should not have altered flight trajectories once it was out of view.



Figure 5-6 Pushed off course. a) Vector field predicted by the snapshot model. Arrows left of the added landmark have been removed since the approaches to this landmark rendered these headings unrepresentative. The novel landmark disappeared from view once bees were amongst the 3-element array. b.) Actual flight tracks of the approaches of honey bees to the added landmark and training array. Excessive looping behavior was removed in order to provide a clearer picture of approach trajectories.

Array configuration

Bees in the above experiments did not appear to randomly select which element of the array to use as a beacon. Instead, they seemed to be selecting the element that was nearest their starting location, whether that was the entrance or some position beyond a novel en-route landmark. In the first removal condition, when there was no element that was considerably closer to their starting position, bees centered their approached between the two remaining landmarks (**Figure 5-4**). This led to an investigation of the approaches of bees to the landmark array in the replication experiment where the landmarks were

positioned symmetrically around the arena entrance (Figure 5-2b). In this configuration there was no single landmark that was substantially closer to the entrance: the two closest landmarks were at approximately the same distance. Given the results above, one might expect that individuals would have selected one of the two close, lateral landmarks and used it as a beacon under these circumstances. Instead, bees headed toward the center of mass of the array in a way similar to the flight behavior shown the removal condition, and other previous experiments (Figures 5-4, 5-7; Fry & Wehner, 2005).



Figure 5-7 Model predictions and flight paths to a 3 landmark array. **a**.) Vector field predicted by the snapshot model. This plot predicts the direction a bee should travel from each 10 cm² bin if she is using the template matching algorithm proposed by Cartwright and Collett (1984). **b**.) Actual flight tracks of the approaches of honey bees to the landmark array. Excessive looping behavior was removed in order to provide a clearer picture of approach trajectories.

Body axis analysis

Two lines of evidence in support of an image matching hypothesis deal with body axis orientation near the goal. The first comes from Collett & Baron (1994). In a set of experiments honey bees were trained to find food in relation to a single landmark positioned either to the east, west, north or south of a feeder. They found that bees tended to consistently face south when approaching the feeder, except when the landmark was north of the goal, in which case bees learned to approach the goal facing north. By facing north bees had a better view of the landmark.

If bees use an image matching strategy to find the food in our arena we might expect them to align their bodies in a manner that maximizes their use of the landmarks provided. To investigate this hypothesis we looked at the mean body axis heading for bees near the goal under three conditions: 1.) training with a single landmark (from experiment 1) 2.) training with 3 symmetrical landmarks (from experiment 1) 3.) training with 3 symmetrical landmarks positioned 70 degrees to the right of the entrance (from experiment 2). Figures 5-8a, b illustrates the strength and direction of the mean body axis orientation of individual bees for each landmark condition. Table 5-1 provides the mean values as well as the 95% angular confidence interval around the mean. Flight headings of bees trained and tested with a single landmark were significantly clustered with a mean angle of 206° (Rayleigh test: *1 landmark*: n=22, r=0.387, p < 0.05) (Figure 5-8a). Flight headings of bees trained and tested with three landmarks centered on the entrance were also significantly clustered but with a mean angle of 183⁰(Rayleigh test n=15, r=0.59, p<0.05) (Figure 8-5b). However, flight headings of bees trained and tested with three landmarks that were not centered on the entrance were randomly distributed with a mean angle of 155° (Rayleigh test n=20, r=0.28, p > 0.05) (Figure 8-5c). In all cases the confidence interval of the mean angle included 180° which was the angle corresponding to due south in our arena. Although subjects in all of these experiments generally faced south during their initial pass over the goal area (as found by Collett & Baron (1994)), our results are much more scattered, especially in the single and offset 3-landmark conditions.



Figure 5-8. Body axis headings near the goal. Large black circles represent the landmark array set at there angles relative to the feeder hole (center of the plot circle). The open square represents the angular direction of the arena entrance. Each thin arrow is the average body axis angle of a single bee during her first pass through the goal area defined by a 20 cm circle centered on the feeder hole. The dark black line within the unit circle represents the direction (angle) and strength (length) of the mean vector of all bees in that condition (values given in Table 5-1). a.) bees trained and tested with one landmark 60 cm from the feeder hole b.) bees trained and tested with three equidistant/angular centered on the arena entrance c.) bees trained and tested with three right of the entrance.

	Mean Angle	95% Confidence Interval
a.) Single Landmark	± 206°	± 50°
b.) Three Landmarks centered on entrance	± 183°	± 35°
c.) Three Landmarks offset from entrance	± 155°	±>100°
Collett & Baron (1994)	± 180°	± 22°

Table 5-1. Body axis headings near the goal. Angular means and confidence intervals of headings shown in Figure 5-8.

The second line of evidence comes from Collett & Rees (1997). In this experiment they trained bees to find a feeder indicated by a single cylindrical landmark. They found

that when bees neared the goal they would perform bouts of lateralized flight while facing the landmark. It was argued that this behavior provides a view of the landmark consistent with that learned during earlier turn-back-and-look orientation flights where bees and wasps are thought to memorize portions of there snapshots. The next experiment examined this by closely inspecting how honeybees oriented their bodies while searching near the goal.

Experiment 3

In order to better mimic the conditions of Collett & Rees (1997) we trained bees to find the feeder hole in relation to a large cylinder landmark positioned 10 cm from the goal. In one condition the landmark was placed directly behind the feeder hole such that the arena entrance, feeder hole and cylinder were arranged in a straight line (**Figure 5-9a**). In a second condition the landmark was placed to the right of the feeder hole orthogonal to the straight path from the entrance to the feeder (**Figure 5-9b**).

In our arena, bees generally keep their body axis in line with their flight heading even near the goal. Although we saw a great deal of circling behavior, we rarely observed bouts of sideways flight near the goal (**Figure 5-9**). Of the 30 bees trained to these configurations the dashed circle in **Figure 5-9b** highlights one of the very few brief instances in which a bee took up a lateral flight pattern near the landmark and feeder hole.


Figure 5-9 Body axis headings near the goal. Single bee tracks with body axis. Asterisks represent the bees' head. The direction (but not distance) of the arena entrance is represented by the letter E and the feeder hole by F. The filled black circle designates the landmark. The feeder-hole and landmark are positioned correctly and to scale. To conserve space the entrance is located at the correct direction from the feeder but at a closer distance than during actual training and testing. Numbers help designate the sequence of the path.

Discussion

Recent work by Fry and Wehner (2005) as well as our own observations led us to reevaluate the use of a template matching strategy by honey bees to guide flight near a familiar goal location. We began by implementing in Matlab the snapshot algorithm proposed by Cartwright & Collett (1987) to generate predictions of honey bee flight paths for similar landmark arrays at the same spatial scale used in early experiments. Although the actual approach flights of bees to these arrays only occasionally matched those predicted by the model, we found that two aspects of honey bee flight behavior within our arena were consistent with previous data in support of snapshot guidance. First, honey bees' search behavior was based on the angular size of, and angles among, visual landmarks (Cartwright & Collett, 1979, 1982). Second, bees tended to maintain a fixed, south-facing body axis orientation when near the goal location. In other respects the snapshot models was not upheld.

Support for the snapshot model

As with Cartwright and Collett (1979), doubling the training size of a single landmark in our arena caused honey bees to search farther from the landmark than during training (**Figure 5-2a**). Also consistent with Cartwright and Collett (1982), bees trained with a constellation of three landmarks did not alter their search behavior when we double the size of the array elements (**Figure 5-2b**). Our results indicate that bees judge the distance of a landmark from the goal using the retinal size of landmarks or, in more complex arrays, the angular distance between landmarks. In addition, it is worth noting that we also reproduced a peculiarity found in the earlier experiments: our results as well as those of Cartwright and Collett (1979, 1982) show that when tested with the training landmark(s) bees tend to search at a location nearer the landmarks than was the location of the feeder during training.

The second line of support is our corroboration of the earlier findings of Collett & Baron (1994) showing that bees take a consistent body axis orientation when very close to the goal. As stated above this is not direct evidence for the snapshot hypothesis but it would certainly be a means for bees to simplify template matching. Our evidence suggests that bees do indeed prefer a southern facing body orientation when they are near the goal (**Figure 5-8, Table 5-1**).

Though our evidence on body axis orientation confirms earlier work, our results need to be qualified in two ways. First, Collett & Baron's (1994) analysis of body axis orientation to a single landmark excluded the portions of flight where bees frontally faced the landmark. The authors did this to ensure that portions of flight where bees treated the landmark as a beacon did not conceal the otherwise robust pattern of southern orientation

near the goal. In our analyses of the 3-landmark arrays, it was not clear which landmark might serve as a beacon. Rather than lose a considerable amount of body axis data, we did not omit any portion of bee flights near the goal. For the sake of consistency we treated our single landmark data in the same manner. It is very likely that the greater scatter in our body axis data can be attributed to our inclusion of landmark-oriented flight.

The second qualification of our body axis result has to do with the orientation of the apparatus. Because our arena is confined to a relatively small room with a single north-facing window through which bees entered, our landmark arrays were always positioned to the south of the entrance and feeder-hole. This limitation prevented us from knowing which environmental cue(s) bees used to maintain a given body axis near the goal. When available, bees can use the earth's magnetic field (Collett & Baron, 1994) and the position(s) of the presented landmark(s) (Collett & Baron, 1994; Cheng, 1999b; Fry & Wehner, 2005). We are unable to determine whether bees oriented to magnetic cues or the landmarks or both since both of these cues were available during our experiments. However, bees could not have oriented their body axis based on celestial cues since the experiment was performed indoors without a view of the sky.

The snapshot model did a generally poor job of predicting honey bee flight paths to our landmarks arrays. Only in rare instances did the vector fields generated by the model correspond with bee approach trajectories. In addition, we observed very little lateral flight near the goal as described by Collett & Baron (1994) and Collett & Reese (1997). I discuss each of these incongruencies in turn in the following sections.

<u>Beacons</u>

The approach flights of bees from outside the landmark array conformed to the snapshot model best in instances where no one landmark in the array appeared larger from the arena entrance (Figure 5-4, 7). In cases where one of the elements was visually larger (Figure 5-3, 5) bees set a course to this landmark and approached it until they were within the array. This was especially noticeable in the condition where we added a novel landmark to pull bees off course (Figure 5-6). As soon as bees no longer had a view of the new landmark they usually selected the nearest (and therefore largest) next landmark as a beacon to continue their flight into the landmark array.

It is unclear whether bees used a snapshot mechanism like that of Cartwright & Collett (1983) to guide flight to our symmetrical landmark (**Figures 5-4**, 7) arrays since several individuals in each case appeared to orient to a single landmark for at least a portion of their approach flight. The look-and-turn strategy of flight guidance recently proposed by Fry & Wehner (2005) may better accounts for our results. This model predicts that when landmarks are identical in size, shape and color bees will fly to the center of mass of the array but if one landmark of the array is more salient it gets adopted as a beacon. Depending on where a bee finds herself inside our arena any one landmark may appear larger than the other elements present and subsequently get adopted as a beacon. This certainly seems the case for bees that were pushed off course by a novel landmark. It is important to note that Collett & Baron (1994) and Collett & Rees (1997) also suggested a critical a role for nearby landmarks as beacons during the approach to a goal. Other research has suggested that bees will select a single element as a beacon out of a multiple

landmark array based visual properties such as color, shape and size (Collett et al., 1996; Cheng, 1999a; Fry & Wehner, 2005 see also chapter3).

Once bees selected a landmark as a beacon and presumably approached it throughout the 60 + training trips they displayed no tendency to visit that location when we removed the focal landmark (**Figure 5-3, 4**). This suggests that if bees associated goal directed local vectors with the beacon landmark (i.e. fly to the largest black stripe, make a hard left and fly for 60 cm, look for the feeder hole) these commands went unused when the landmark was absent. This result contrasts with previous data showing that bees and ants maintain fixed routes in the absence of the original cues that marked the path (Collett et al., 1993; Graham et al., 2003).

Scanning flight and TBLs

Figure 5-9b highlights one of the few times a bee conducted a lateral flight over the goal location in our arena. As discussed previously, this sort of "scanning" flight is believed to mimic the lateral arcing flights performed during the turn-back-and-look learning flights bees and wasps perform before departing a newly visited foraging site. The TBL is reported to aid the taking of the snapshot by giving information about the sizes and distances of nearby landmarks (Lehrer, 1991, 1993; Zeil et al., 1996). Lateral flight on return to the goal is believed to help align a bee's snapshot to the current view of local landmarks in a way that is similar with their view during previous learning flights. Bees in our apparatus were denied TBLs due to the one-way nature of the arena. The absence of TBLs may account for the lack of subsequent lateral flight during their search near the goal. We find it interesting that bees still learned relevant landmark properties

such as inter-landmark angles and the retinal size of a single landmark without the aid of their natural learning flights. Whether or not our data cast doubt on the putative role of TBLs in landmark learning is worth further investigation.

Conclusions

In general it appears that image matching per se is a strategy bees use only when very close to the goal. As in other studies we found that successive use of beacons accounts for much of the flight behavior we observed within our arena (Collett et al., 1993; Collett & Rees, 1997; Cheng, 1999a, b; Fry & Wehner, 2002). We also provide evidence that when the elements of a landmark array are of the same shape and color bees select a beacon based its visual size as seen from their current location, with a preference for the largest element. It would be interesting to determine how the visual properties of landmarks such as size, color and shape interact during beacon selection in honey bees.

That insects also use visual landmarks as beacons during navigation over long distances is well documented (Frisch, 1967; Chittka et al., 1995; Chittka, 1995; Chittka & Gieger, 1995; Graham et al., 2003). Beacons seem to play an important role in both capacities and a more comprehensive investigation into how insects and other animals select and use them might reveal that very similar navigational strategies are at work on both large and small spatial scales.

From a sequential decision making perspective, these results suggest that bees prefer to simplify their large and complex foraging environment by segmenting their routes into a series of states each defined by a prominent visual beacon. We provide little evidence that bees use a snapshot-like generalization strategy while they are en-route to a familiar

feeding location. Although template matching is an attractive computational strategy for simplifying large continuous state spaces, bees probably use snapshots over a relatively small number of states very near the final goal.

REFERENCES

- Anderson, A. M. 1977. Model for Landmark Learning in Honeybee. Journal of Comparative Physiology 114 (3):335-355.
- Beekman, M., and F. L. W. Ratnieks. 2000. Long-range foraging by the honey-bee, Apis mellifera L. *Functional Ecology* 14 (4):490-496.
- Benhamou, S., P. Bovet, and B. Poucet. 1995. A Model for Place Navigation in Mammals. *Journal of Theoretical Biology* 173 (2):163-178.
- Bianco, G., and R. Cassinis. 2000. Biologically-inspired visual landmark learning for mobile robots. In *Advances in Robot Learning, Proceedings*.
- Biro, D., T. Guilford, and M. S. Dawkins. 2003. Mechanisms of visually mediated site recognition by the homing pigeon. *Animal Behaviour* 65:115-122.
- Capaldi, E. A., and F. C. Dyer. 1999. The role of orientation flights on homing performance in honeybees. *Journal of Experimental Biology* 202 (12):1655-1666.
- Capaldi, E. A., A. D. Smith, J. L. Osborne, S. E. Fahrbach, S. M. Farris, D. R. Reynolds,
 A. S. Edwards, A. Martin, G. E. Robinson, G. M. Poppy, and J. R. Riley. 2000.
 Ontogeny of orientation flight in the honeybee revealed by harmonic radar.
 Nature 403 (6769):537-540.
- Cartwright, B. A., and T. S. Collett. 1979. How Honeybees Know Their Distance from a near-by Visual Landmark. *Journal of Experimental Biology* 82 (OCT):367-372.
- Cartwright, B. A., and T. S. Collett. 1982. How Honey Bees Use Landmarks to Guide Their Return to a Food Source. *Nature* 295 (5850):560-564.
- Cartwright, B. A., and T. S. Collett. 1983. Landmark Learning in Bees Experiments and Models. *Journal of Comparative Physiology* 151 (4):521-543.
- Cheng, K., T. S. Collett, and R. Wehner. 1986. Honeybees Learn the Colors of Landmarks. Journal of Comparative Physiology a-Sensory Neural and Behavioral Physiology 159 (1):69-73.
- Cheng, K. 1999a. Landmark-based spatial search in honeybees. I. Animal Cognition 2 (2):73-78.
- Cheng, K. 1999b. Landmark-based spatial search in honeybees. II. Animal Cognition 2 (2):79-90.

- Cheng, K., and M. L. Spetch. 2001. Blocking in landmark-based search in honeybees. Animal Learning & Behavior 29 (1):1-9.
- Chittka, L. 1995. The influences of landmarks and distance estimation of honey bees. Animal Behaviour 50:1709-1710.
- Chittka, L., and K. Geiger. 1995. Honeybee Long-Distance Orientation in a Controlled Environment. *Ethology* 99 (2):117-126.
- Chittka, L., J. Kunze, C. Shipman, and S. L. Buchmann. 1995. The Significance of Landmarks for Path Integration in Homing Honeybee Foragers. *Naturwissenschaften* 82 (7):341-343.
- Collett, T. S., and M. F. Land. 1975. Visual Control of Flight Behavior in Hoverfly, Syritta-Pipiens L. Journal of Comparative Physiology 99 (1):1-66.
- Collett, T. S., and A. Kelber. 1988. The Retrieval of Visuo-Spatial Memories by Honeybees. Journal of Comparative Physiology a-Sensory Neural and Behavioral Physiology 163 (1):145-150.
- Collett, T. S., S. N. Fry, and R. Wehner. 1993. Sequence Learning by Honeybees. Journal of Comparative Physiology a-Sensory Neural and Behavioral Physiology 172 (6):693-706.
- Collett, T. S., and J. Baron. 1994. Biological Compasses and the Coordinate Frame of Landmark Memories in Honeybees. *Nature* 368 (6467):137-140.
- Collett, T. S., and J. Baron. 1995. Learnt Sensorimotor Mappings in Honeybees -Interpolation and Its Possible Relevance to Navigation. *Journal of Comparative Physiology a-Sensory Neural and Behavioral Physiology* 177 (3):287-298.
- Collett, T. S. 1996. Insect navigation en route to the goal: Multiple strategies for the use of landmarks. *Journal of Experimental Biology* 199 (1):227-235.
- Collett, T. S., and J. A. Rees. 1997. View-based navigation in Hymenoptera: Multiple strategies of landmark guidance in the approach to a feeder. Journal of Comparative Physiology a-Sensory Neural and Behavioral Physiology 181 (1):47-58.
- Collett, M., and T. S. Collett. 2000. How do insects use path integration for their navigation? *Biological Cybernetics* 83 (3):245-259.
- Collett, T. S., and M. Collett. 2000. Path integration in insects. Current Opinion in Neurobiology 10 (6):757-762.

- Collett, M., D. Harland, and T. S. Collett. 2002. The use of landmarks and panoramic context in the performance of local vectors by navigating honeybees. *Journal of Experimental Biology* 205 (6):807-814.
- Collett, T. S., and M. Collett. 2002. Memory use in insect visual navigation. *Nature Reviews Neuroscience* 3 (7):542-552.
- Collett, M., T. S. Collett, S. Chameron, and R. Wehner. 2003. Do familiar landmarks reset the global path integration system of desert ants? *Journal of Experimental Biology* 206 (5):877-882.
- Collett, T. S., and M. Collett. 2004. How do insects represent familiar terrain? Journal of *Physiology-Paris* 98 (1-3):259-264.
- Daw, N. D., and D. S. Touretzky. 2000. Behavioral considerations suggest an average reward TD model of the dopamine system. *Neurocomputing* 32:679-684.
- Daw, N. D., and D. S. Touretzky. 2002. Long-term reward prediction in TD models of the dopamine system. *Neural Computation* 14 (11):2567-2583.
- Dawkins, M. S., and A. Woodington. 2000. Pattern recognition and active vision in chickens. *Nature* 403 (6770):652-655.
- Demas, G. E., and M. F. Brown. 1995. Honey-Bees Are Predisposed to Win-Shift but Can Learn to Win-Stay. *Animal Behaviour* 50:1041-1045.
- Dickinson, J. A. 1994. Bees Link Local Landmarks with Celestial Compass Cues. Naturwissenschaften 81 (10):465-467.
- Durier, V., P. Graham, and T. S. Collett. 2003. Snapshot memories and landmark guidance in wood ants. *Current Biology* 13 (18):1614-1618.
- Dyer, F. C., and J. L. Gould. 1981. Honey Bee Orientation a Backup System for Cloudy Days. Science 214 (4524):1041-1042.
- Dyer, F. C. 1987. Memory and Sun Compensation by Honey-Bees. Journal of Comparative Physiology a-Sensory Neural and Behavioral Physiology 160 (5):621-633.
- Dyer, F. C., and J. A. Dickinson. 1994. Development of Sun Compensation by Honeybees - How Partially Experienced Bees Estimate the Suns Course. Proceedings of the National Academy of Sciences of the United States of America 91 (10):4471-4474.

- Dyer, F. C., and J. A. Dickinson. 1996. Sun-compass learning in insects: Representation in a simple mind. *Current Directions in Psychological Science* 5 (3):67-72.
- Dyer, F. C. 1998. Spatial cognition: lessons from central-place foraging insects. In Animal cognition in nature : the convergence of psychology and biology in laboratory and field, edited by R. P. Balda, I. M. Pepperberg and A. C. Kamil. San Diego: Academic Press.
- Dyer, F. C., M. Gill, and J. Sharbowski. 2002. Motivation and vector navigation in honey bees. *Naturwissenschaften* 89 (6):262-264.
- Dyer, F. C. 2002. The biology of the dance language. *Annual Review of Entomology* 47:917-949.
- Dyer, A. G., and L. Chittka. 2004. Bumblebees (Bombus terrestris) sacrifice foraging speed to solve difficult colour discrimination tasks. Journal of Comparative Physiology a-Neuroethology Sensory Neural and Behavioral Physiology 190 (9):759-763.
- Esch, H. E., and J. E. Burns. 1995. Honeybees Use Optic Flow to Measure the Distance of a Food Source. *Naturwissenschaften* 82 (1):38-40.
- Esch, H. E., S. W. Zhang, M. V. Srinivasan, and J. Tautz. 2001. Honeybee dances communicate distances measured by optic flow. *Nature* 411 (6837):581-583.
- Etienne, A. S., R. Maurer, and V. Seguinot. 1996. Path integration ln mammals and its interaction with visual landmarks. *Journal of Experimental Biology* 199 (1):201-209.
- Franz, M. O., B. Scholkopf, H. A. Mallot, and H. Bulthoff. 1998. Where did I take that snapshot? Scene-based homing by image matching. *Biological Cybernetics* 79 (3):191-202.
- Franz, M. O., and H. A. Mallot. 2000. Biomimetic robot navigation. *Robotics and Autonomous Systems* 30 (1-2):133-153.
- Frisch, K. v., and M. Lindauer. 1954. Himmel Und Erde in Konkurrenz Bei Der Orientierung Der Bienen. *Naturwissenschaften* 41 (11):245-253.
- Frisch, K. v. 1967. *The dance language and orientation of bees*. 1st Harvard University Press pbk. ed. Cambridge, Mass.: Harvard University Press.
- Fry, S. N., M. Bichsel, P. Muller, and D. Robert. 2000. Tracking of flying insects using pan-tilt cameras. *Journal of Neuroscience Methods* 101 (1):59-67.

- Fry, S. N., and R. Wehner. 2002. Honey bees store landmarks in an egocentric frame of reference. Journal of Comparative Physiology a-Neuroethology Sensory Neural and Behavioral Physiology 187 (12):1009-1016.
- Fry, S. N., and R. Wehner. 2005. Look and turn: landmark-based goal navigation in honey bees. *Journal of Experimental Biology* 208 (20):3945-3955.
- Giger, A., and M. Srinivasan. 1997. Honeybee vision: analysis of orientation and colour in the lateral, dorsal and ventral fields of view. *J Exp Biol* 200 (8):1271-1280.
- Giurfa, M., J. Núñez, L. Chittka, and R. Menzel. 1995. Colour preferences of flowernaive honeybees. Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology 177 (3):247-259.
- Giurfa, M., M. Vorobyev, P. Kevan, and R. Menzel. 1996. Detection of coloured stimuli by honeybees: Minimum visual angles and receptor specific contrasts. Journal of Comparative Physiology a-Sensory Neural and Behavioral Physiology 178 (5):699-709.
- Giurfa, M., and R. Menzel. 1997. Insect visual perception: complex abilities of simple nervous systems. *Current Opinion in Neurobiology* 7 (4):505-513.
- Giurfa, M. 2004. Conditioning procedure and color discrimination in the honeybee Apis mellifera. *Naturwissenschaften* 91 (5):228-231.
- Golledge, R. G. 1999. Wayfinding behavior : cognitive mapping and other spatial processes. Baltimore: Johns Hopkins University Press.
- Gould, J. L. 1987. Honey-Bees Store Learned Flower-Landing Behavior According to Time of Day. *Animal Behaviour* 35:1579-1581.
- Graham, P., K. Fauria, and T. S. Collett. 2003. The influence of beacon-aiming on the routes of wood ants. *Journal of Experimental Biology* 206 (3):535-541.
- Hasinoff, S. W. 2002. Reinforcement Learning for Problems with Hidden State. Technical Report: University of Toronto, Department of Computer Science.
- Iersel, J. v., and J. v. d. Assem. 1964. Aspects of Orientation in Diggerwasp Bembix Rostrata. Animal Behaviour S:145-162.
- Kaelbling, L. P., M. L. Littman, and A. W. Moore. 1996. Reinforcement learning: A survey. Journal of Artificial Intelligence Research 4:237-285.
- Kamil, A. C., T. B. Jones, A. Pietrewicz, and J. E. Mauldin. 1977. Positive Transfer from Successive Reversal Training to Learning Set in Blue Jays (Cyanocitta-Cristata). *Journal of Comparative and Physiological Psychology* 91 (1):79-86.

- Lambrinos, D., R. Moller, T. Labhart, R. Pfeifer, and R. Wehner. 2000. A mobile robot employing insect strategies for navigation. *Robotics and Autonomous Systems* 30 (1-2):39-64.
- Lambrinos, D. 2003. Navigation in desert ants: the robotic solution. *Robotica* 21:407-426.
- Lehrer, M., R. Wehner, and M. Srinivasan. 1985. Visual Scanning Behavior in Honeybees. Journal of Comparative Physiology a-Sensory Neural and Behavioral Physiology 157 (4):405-415.
- Lehrer, M., M. V. Srinivasan, S. W. Zhang, and G. A. Horridge. 1988. Motion Cues Provide the Bees Visual World with a 3rd Dimension. *Nature* 332 (6162):356-357.
- Lehrer, M. 1991. Bees Which Turn Back and Look. Naturwissenschaften 78 (6):274-276.
- Lehrer, M. 1993. Why Do Bees Turn Back and Look. Journal of Comparative Physiology a-Sensory Neural and Behavioral Physiology 172 (5):549-563.
- Lehrer, M., and M. V. Srinivasan. 1994. Active Vision in Honeybees Task-Oriented Suppression of an Innate Behavior. *Vision Research* 34 (4):511-516.
- Lehrer, M., G. A. Horridge, S. W. Zhang, and R. Gadagkar. 1995. Shape Vision in Bees -Innate Preference for Flower-Like Patterns. *Philosophical Transactions of the Royal Society of London Series B-Biological Sciences* 347 (1320):123-137.
- Lehrer, M. 1996. Small-scale navigation in the honeybee: Active acquisition of visual information about the goal. *Journal of Experimental Biology* 199 (1):253-261.
- Lehrer, M., and G. Bianco. 2000. The turn-back-and-look behaviour: bee versus robot. Biological Cybernetics 83 (3):211-229.
- Levine, M. 1959. A Model of Hypothesis Behavior in Discrimination-Learning Set. *Psychological Review* 66 (6):353-366.
- Levine, M. 1965. Edited by A. M. Schrier, H. F. Harlow and F. Stollnitz, *Behavior of nonhuman primates: modern research trends*. New York,: Academic Press.
- Maguire, E. A., N. Burgess, and J. O'Keefe. 1999. Human spatial navigation: cognitive maps, sexual dimorphism, and neural substrates. *Current Opinion in Neurobiology* 9 (2):171-177.
- Manser, M. B., and M. B. Bell. 2004. Spatial representation of shelter locations in meerkats, Suricata suricatta. *Animal Behaviour* 68:151-157.

- McCallum, A. K. 1996. Reinforcement Learning with Selective Perception and Hidden State. PhD, University of Rochester.
- McMahon, A., D. Scott, and W. Browne. 2005. An autonomous explore/exploit strategy. Paper read at Proceedings of the 2005 workshops on Genetic and evolutionary computation, at Washington, D.C.
- Menzel, R., and W. Backhaus. 1991. Colour vision in insects. In Vision and visual dysfunction, edited by J. Cronly-Dillon. Boca Raton: CRC Press.
- Mitchell, T. M. 1997. Machine Learning. New York: McGraw-Hill.
- Moller, R. 2000. Insect visual homing strategies in a robot with analog processing. Biological Cybernetics 83 (3):231-243.
- Neukirch, A. 1982. Dependence of the Life-Span of the Honeybee (Apis-Mellifica) Upon Flight Performance and Energy-Consumption. *Journal of Comparative Physiology* 146 (1):35-40.
- Olton, D. S., G. E. Handelmann, and J. A. Walker. 1981. Spatial memory and food searching strategies. In Foraging Behavior: Ecological, Ethological, and Psychological Approaches, edited by A. C. Kamil and T. D. Sargents. New York:: Garland Press.
- Reinhard, J., M. V. Srinivasan, and S. W. Zhang. 2004. Olfaction: Scent-triggered navigation in honeybees. *Nature* 427 (6973):411-411.
- Ronacher, B., and U. Duft. 1996. An image-matching mechanism describes a generalization task in honeybees. Journal of Comparative Physiology a-Sensory Neural and Behavioral Physiology 178 (6):803-812.
- Rossel, S., and R. Wehner. 1984. Celestial Orientation in Bees the Use of Spectral Cues. Journal of Comparative Physiology 155 (5):605-613.
- Rossel, S., and R. Wehner. 1986. Polarization Vision in Bees. Nature 323 (6084):128-131.
- Schuster, S., and S. Amtsfeld. 2002. Template-matching describes visual patternrecognition tasks in the weakly electric fish Gnathonemus petersii. *Journal of Experimental Biology* 205 (4):549-557.
- Shani, G., and R. I. Brafman. 2004. Resolving perceptual aliasing in the presence of noisy sensors. In *Proceedings of the Eighteenth Annual Conference on Neural Information Processing Systems*: MIT Press.

- Srinivasan, M. V., and M. Lehrer. 1988. Spatial Acuity of Honeybee Vision and Its Spectral Properties. Journal of Comparative Physiology a-Sensory Neural and Behavioral Physiology 162 (2):159-172.
- Srinivasan, M. V., S. W. Zhang, and N. J. Bidwell. 1997. Visually mediated odometry in honeybees. *Journal of Experimental Biology* 200 (19):2513-2522.
- Srinivasan, M. V., S. W. Zhang, and H. Zhu. 1998. Honeybees link sights to smells. *Nature* 396 (6712):637-638.
- Srinivasan, M. V., S. W. Zhang, J. Berry, K. Cheng, and H. Zhu. 1999. Honeybee navigation: linear perception of short distances travelled. *Journal of Comparative Physiology a-Sensory Neural and Behavioral Physiology* 185 (3):239-245.
- Srinivasan, M. V., S. W. Zhang, M. Altwein, and J. Tautz. 2000. Honeybee navigation: Nature and calibration of the "odometer". *Science* 287 (5454):851-853.
- Stach, S., J. Benard, and M. Giurfa. 2004. Local-feature assembling in visual pattern recognition and generalization in honeybees. *Nature* 429 (6993):758-761.
- Sutton, R. S., A. G. Barto, and NetLibrary Inc. 1998. Reinforcement learning an introduction. Cambridge, Mass.: MIT Press,.
- Tesauro, G. 1992. Practical Issues in Temporal Difference Learning. *Machine Learning* 8 (3-4):257-277.
- Tinbergen, N., and W. Kruyt. 1938. Über die Orientierung des Bienenwolfes (*Philanthus triangulum Fabr.*) III. Die Bevorzugung bestimmter Wegmarken. Z Vergl Physiol 25:292–334.
- Tinbergen, N. 1972. The animal in its world; explorations of an ethologist, 1932-1972. Cambridge, Mass.,: Harvard University Press.
- Towne, W. F., and J. L. Gould. 1988. The spatial precision of the honey bees' dance communication. *Journal of Insect Behavior* 1 (2):129-155.
- Trullier, O., S. I. Wiener, A. Berthoz, and J. A. Meyer. 1997. Biologically based artificial navigation systems: Review and prospects. *Progress in Neurobiology* 51 (5):483-544.
- Visscher, P. K., and T. D. Seeley. 1982. Foraging Strategy of Honeybee Colonies in a Temperate Deciduous Forest. *Ecology* 63 (6):1790-1801.
- Vladusich, T., J. M. Hemmi, M. V. Srinivasan, and J. Zeil. 2005. Interactions of visual odometry and landmark guidance during food search in honeybees. *Journal of Experimental Biology* 208 (21):4123-4135.

- Watson, J. B., and K. S. Lashley. 1915. *Homing and related activities of birds*. Washington,: Carnegie Institution of Washington.
- Wehner, R., and F. Raber. 1979. Visual Spatial Memory in Desert Ants, Cataglyphis-Bicolor (Hymenoptera, Formicidae). *Experientia* 35 (12):1569-1571.
- Wehner, R. 1981. Spatial vision in arthropods. In Handbook of sensory physiology. [Berlin, New York,: Springer-Verlag.
- Wehner, R., and S. Rossel. 1983. Polarized-Light Navigation in Bees Use of Zenith and Off-Zenith E-Vectors. *Experientia* 39 (6):642-642.
- Wehner, R. 1996. Middle-scale navigation: The insect case. Journal of Experimental Biology 199 (1):125-127.
- Wehner, R. 2003. Desert ant navigation: how miniature brains solve complex tasks. Journal of Comparative Physiology a-Neuroethology Sensory Neural and Behavioral Physiology 189 (8):579-588.
- Wei, C. A., S. L. Rafalko, and F. C. Dyer. 2002. Deciding to learn: modulation of learning flights in honeybees, Apis mellifera. Journal of Comparative Physiology a-Neuroethology Sensory Neural and Behavioral Physiology 188 (9):725-737.
- Whitehead, S. D., and D. H. Ballard. 1991. Learning to Perceive and Act by Trial and Error. *Machine Learning* 7 (1):45-83.
- Wilson, S. 1996. Explore/Exploit Strategies in Autonomy. In From animals to animats 4 : proceedings of the Fourth International Conference on Simulation of Adaptive Behavior, edited by P. Maes, J. Pollack, J. A. Meyer and S. W. Wilson. Cambridge, Mass.: MIT Press.
- Zeil, J., A. Kelber, and R. Voss. 1996. Structure and function of learning flights in bees and wasps. *Journal of Experimental Biology* 199 (1):245-252.
- Zhang, S. W., K. Bartsch, and M. V. Srinivasan. 1996. Maze learning by honeybees. Neurobiology of Learning and Memory 66 (3):267-282.
- Zhang, S. W., M. Lehrer, and M. V. Srinivasan. 1999. Honeybee memory: Navigation by associative grouping and recall of visual stimuli. *Neurobiology of Learning and Memory* 72 (3):180-201.
- Zhang, S. W., A. Mizutani, and M. V. Srinivasan. 2000. Maze navigation by honeybees: Learning path regularity. *Learning & Memory* 7 (6):363-374.

