Decision-Making and Action Selection in Honeybees: a Theoretical and Experimental Study

Lianne F. S. Meah

Supervisors: Prof James Marshall, Prof Andrew Barron, Prof Eleni Vasilaki

Department of Computer Science / Department of Biological Sciences
Faculty of Engineering / Faculty of Science and Engineering

A thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy

March 2018

I dedicate this thesis to Steven, Kris and Peter.

Declaration

I declare that the research presented in this document is original work, except where references are made to the work of others. This work has been completed at the University of Sheffield and Macquarie University as part of a joint PhD program, and some of the experimental work was done at Paul Sabatier University, Toulouse. Any illustrations which are not the author's own have been used with the explicit permission of the originator and are specifically acknowledged.

Acknowledgements

I wish to thank all of my supervisors for their ideas, support and guidance, without which this work could not have been completed. My thanks go to James Marshall for his advice, encouragement and ongoing support over the years, and for his general assistance as my supervisor. Thank you also to Eleni Vasilaki for her continuous guidance and patience (and also for all the coffees!). Thanks also go to Andrew Barron in Sydney for many insightful discussions. His expertise on honeybees contributed significantly to the maturation of this project. Many thanks also to Jean-Marc Devaud at Paul Sabatier University in Toulouse for kindly making room for me in his lab and allowing me to carry out my biological experiments, and to Kevin Gurney and Ken Cheng for their valuable input.

Thanks go to all of my colleagues, from all the labs I have worked in, for their interest in my research, their valuable suggestions and for all the good times. Thanks also go to my friends here in the UK and to the new ones I made over in Australia. Thanks especially to Peter, for his unconditional support, guidance and wisdom, and for being beside me throughout all these years. Thank you to Tim for all his love and support in Sydney, and to my second family over there who welcomed me into their lives. Thanks also go to my mother for supporting my academic choices and for always believing in my success.

This research was supported by an Engineering and Physical Sciences Research Council grant, as well as an International Macquarie Research Excellence Scholarship (iMQRES).

Abstract

Decision-making is an integral part of everyday life for animals of all species. Some decisions are rapid and based on sensory input alone, others rely on factors such as context and internal motivation. The possibilities for the experimental investigation of choice behaviour in mammals, especially in humans, are seemingly endless. However, neuroscience has struggled to define the neural circuitry behind decision-making processes due to the complex structure of the mammalian brain.

For this work we turn to the honeybee for inspiration. With a brain composed of approximately 10^6 neurons and sized at a tiny $1mm^3$, it may be assumed that such an insect produces mere 'programmed' behaviours, yet, the honeybee exhibits a rich, elaborate behavioural repertoire and a large capacity for learning in a variety of different paradigms. Indeed, the honeybee has been identified as a powerful model for decision-making.

Sequential sampling models, originating in psychology, have been used to explain rapid decision-making behaviours. Such models assume that noisy sensory evidence is integrated over time until a threshold is reached, whereby a decision is made. These models have proven popular because they are able to fit biological data and are furthermore supported by neural evidence. Additionally, they explain the speed-accuracy trade-off, a behavioural phenomenon also demonstrated in bees.

For this work we examine honeybee choice behaviour in different levels of satiation, and show that hungry bees are faster and less accurate than partially satiated bees in a simple choice task. We suggest that differences in

choice behaviour may be attributed to a simple mechanism which alters the level of the decision threshold according to how satiated the bee is. We further speculate that the honeybee olfactory system may be a drift-diffusion channel, and develop a simple computational model, based on honeybee neurobiology, with simulations that match behavioural results.

Contents

List of Figures

Chapter 1

Introduction

This research aims to unite computational modelling with biological experimentation to examine decision-making in perceptual choice tasks, emphasising the importance of the widely applicable 'accumulator' models of decision-making and their contributions to understanding choice behaviours. To this end, the study examines the role of inhibitory circuits within these models and the impact they have upon decision processes. Additionally, it also examines the influence of satiation on perceptual decision-making with respect to the mechanisms of the aforementioned accumulator models. The results here bring together computational neuroscience with biology and will hopefully encourage future studies to do the same.

Decision-making has been well studied behaviourally in a wide variety of different animals and contexts. All animals need to make decisions in their day-to-day lives and some of these will be inherently more complex than others. For example, a foraging animal must continually decide where to search for food and how to carry out this process efficiently, such that the costs incurred are minimised (Marshall et al. 2015). This type of decision-making often requires discrimination between alternative options. For example, a foraging bee will need to discriminate between a rewarding flower and the alternatives which are similar in colour or odour (Dyer and Chittka 2004a). Foraging is one example which is shared across species, however, other levels

of decision-making can also be observed. This becomes particularly apparent in human choice behaviour. When instructed to analyse a photograph, a human will explore differing methods of gathering visual data, implementing eye saccades in a manner which optimises the analysis (Walker-Smith et al. 2013). For example, asking someone to estimate a subject's age from a photograph will invoke initial eye saccades over the face. Deciding to invest in a long-term commitment, such as a house, involves a longer and more complex decision-making process, in which emotions and past experiences play a role. Regardless of the complexity, however, it can be asserted that decision-making is indeed a *process*. In 'Multiple Criteria Decision Making', Milan Zeleny described the process as a '*dynamic and interrelated unity of predecision, decision and postdecision stages*' (Zeleny 1998, p. 84), identifying that decision-making is not simply an act nor is it static nor absolute. Indeed, choice behaviour is heavily influenced by multiple internal and external factors, such as internal state and task difficulty. What kind of components make up the decision-making process? Are we able to identify them from behavioural data? Can modelling the decision-making process provide any insight?

Fast, robust decision-making between alternative options has been well studied in laboratory settings, with a heavy focus on accuracy and response time (Chittka et al. 2009). This type of decision-making has often been modelled using sequential sampling models, which assume that noisy evidence for a stimulus is accumulated over time until a threshold or boundary is met (Laming 1968, Smith and Vickers 1988, Usher and McClelland 2001, Ratcliff and Smith 2004, Bogacz et al. 2006, Ratcliff and McKoon 2008, Brown and Heathcote 2008, Purcell and Palmeri 2017). The level of the threshold is variable and denotes the amount of evidence that is required to trigger a decision. Although there are now many variations of these models which employ different mechanisms (for example, various forms of inhibition), perhaps the most-well known of these is the drift-diffusion model, which was developed in 1978 by Roger Ratcliff and has been proven to be optimal in the sense that, for a

given accuracy, a decision is reached within the shortest possible time and, for a given reaction time, the highest possible accuracy is obtained (Ratcliff 1978). Differing choice behaviours can be induced by varying the components of sequential sampling models, in particular, a mediation of the threshold denotes how long a decision-maker should wait before deciding. A low threshold corresponds to a fast decision which is more likely to be inaccurate due to an insufficient amount of evidence being integrated, whilst a high threshold allows for more time for evidence integration, thus the decision-maker will be slower but more accurate (Bogacz et al. 2010b). This relationship between speed and accuracy is known as the speed-accuracy trade -off and it is a behavioural phenomenon which has been shown to exist in many animals (Chittka et al. 2009). Modification of sequential sampling models' properties such as the threshold level has been shown to replicate findings from the speed- accuracy trade-off. As such, these models offer an explanation for certain choice behaviours. As well as the threshold level, the drift-diffusion model also incorporates the average rate of accumulation, also known as the drift rate, into the process. Sequential sampling models have been successfully fitted to behavioural data obtained from two-alternative forced-choice tasks and so they have seen a rise in popularity (Ratcliff et al. 2016). It is important to note that these models are applicable only to fast, perceptual decision-making. Other types of decisions, such as which candidate to vote for in an election, are higher level and inherently more complex, engaging other mechanisms within the brain to facilitate consideration over some time (Kanai et al. 2011, Jost et al. 2014).

It might be suggested that the complexity of decision-making and action selection is attributed to the complexity of an agent's neural network, but this is not always the case. A human must innately decide which foot to walk with first in a similar fashion to a quadruped. However, simpler agents may be *confined* to simpler decisions, whilst more complex agents encounter decisions that require more internal debate. How exactly a process is carried out within the brain is difficult to pinpoint, due to the inherently intricate design of neural

circuitry. It is here we might be drawn to study the simpler organisms, the nematode worm *C. elegans* for example, with a nervous system consisting of only 302 or 381 neurons (depending on if the worm is hermaphrodite or male, respectively) (Bono and Villu Maricq 2005). The connectome (the map or diagram of the neural connections within the nervous system) of *C. elegans* was first completed in 1986 (White et al. 1986), thereafter being studied and built upon to produce a more complete picture. Analysis of the wiring diagrams led to the realisation that, functionally, neuronal circuits are dynamic and subject to change due to the influence of neuromodulators such as dopamine or serotonin (Bargmann 2012). In addition to *C. elegans*, other systems such as that of *Drosophila* have been examined extensively. For this work, inspiration is drawn from *Apis mellifera*, the honeybee.

Honeybees have been shown to have a rich behavioural repertoire with a robust capacity for learning and memory. They have a remarkable ability to navigate long distances with a low resolution visual system, and they are able to solve discrimination problems of varying difficulties (Guerrieri et al. 2005). Recently, it has become clear that this insect exhibits behaviours widely regarded as cognitive (for example, the ability to learn two abstract concepts simultaneously (Avarguès-Weber et al. 2012)) and that it may be useful as a model of decision-making (Menzel 2012, Giurfa 2013). The honeybee has a relatively simple neural architecture, with simplicity here being defined in terms of number of neurons. With a brain $1mm^3$ in size that contains one million neurons (Menzel and Giurfa 2001), honeybee neural circuitry is thus far simpler than that found in mammalian brains (there are approximately 10^{10} neurons in human brains, for example) but more complex than other model animals such as *Drosophila* (which has $100,000$ neurons). Since the honeybee has been physically constrained in terms of its body size, it has needed to survive with a smaller brain (and thus a smaller number of neurons). Despite this limitation, the honeybee demonstrates a high capacity for decision-making; it has evolved to be 'intelligent' using a smaller amount of neural circuitry,

giving rise to adaptive decision-making behaviours which have ensured its survival. Many behavioural studies have been performed with honeybees, both in the field and within laboratory settings, and they have demonstrated that honeybees are capable of many forms of learning and are able to apply their knowledge to novel stimuli or situations (Zhang et al. 2012b). For example, bees that were trained to learn the concepts of 'sameness' and 'difference' were able to apply their knowledge to novel stimuli that hadn't been presented before, even outside of the sensory domain they had been trained on (Giurfa et al. 2001). Bees that had been trained on colours were able to perform in transfer tests which made use of black and white gratings, and bees trained on the gratings were also able to perform in tests which instead made use of colours. In a study on maze learning, bees trained to follow certain colours through a maze were able to apply their training to novel mazes they hadn't navigated through before (Zhang et al. 1996). These results demonstrate that honeybees are not simply hard-wired or preprogramed. Its ability to survive despite its simplicity is what makes this insect ideal as a model of decision-making; that these insects also contribute to collective decisions within the hive makes them even more interesting for study.

Over the past few decades, there have been many turning points in honeybee research. The pioneering work of Karl von Frisch in the 20th century was perhaps the first of these turning points. In 1946, von Frisch published 'Die Tänze der Bienen' which documented how honeybee foragers communicate the location of resources to other hive mates by means of the now well known 'waggle dance' (**?**). This ground-breaking discovery in honeybee 'language' won him the Nobel Prize in 1973. Prior to this, the notion that animals - especially insects - could be capable of such intricate communication would not have been entertained. Karl von Frisch's work shone a much-needed light on the inner workings of the honeybee hive, demonstrating for the first time the complexity of the workers within. More recent research has demonstrated that honeybees make use of other signals for their collective decision-making. One

example is the stop signal, which was first proposed to be a begging signal by Harald Esch in 1964 (Esch 1964). After a few decades, research by James Nieh in 1993 indicated that the signal functioned as a way to halt waggle dancing, acting as negative feedback (Nieh 1993). This finding was later confirmed in 2005 by Pastor and Seeley; their study also found that the signal encouraged dancing bees to stop (Pastor and Seeley 2005). To this day research into honeybee communication continues, with many other observed behaviours still not entirely understood.

Another turning point came from the discovery that honeybees could learn Pavlovian associations, a significant finding that arose from Kimihisa Takeda's development of the 'proboscis extension reflex' paradigm in 1961 (Takeda 1961). This paradigm made use of the bee's natural reflex to extend its proboscis to a rewarding stimulus, such as sucrose solution, and the Pavlovian conditioning protocol (Pavlov 1927). It was further solidified by an influential paper by Bitterman and colleagues in 1983 (Bitterman et al. 1983), and the success of the paradigm resulted in many laboratories using it to address a wide range of questions (Giurfa and Sandoz 2012). More complex forms of learning in bees were starting to be discovered in the early 2000's, a notable study being that of Giurfa and colleagues in 2001 which demonstrated that bees are capable of learning abstract properties such as 'sameness' and 'difference' (Giurfa et al. 2001). The study made use of the delayed-matching-to-sample task where bees were required to respond to a stimulus which matched the one shown on the entrance to a maze, and also the delayed-non-matching-to-sample task where bees were required to choose a stimulus which didn't match the one shown. After training, honeybees were shown to transfer what they had learnt to novel stimuli, indicating learning of the concept itself.

It was once assumed that there were great differences between the vertebrate and invertebrate species, such that insects were inherently useless for investigating any aspect of mammalian cognition. However, this assumption began to lose its strength when behavioural studies on honeybees began to

reveal surprising insights about their nature, noted by James Gould, an evolutionary biologist of Princeton University, in the late 1990's: '*This picture has changed over the past decade; honey bees, at least, turn out to be more like birds and mammals...*' (Gould 1986). More recently, it has been suggested that the insect central complex and mammalian basal ganglia are homologous, sharing evolutionary ancestry that results in similar topography and function within those brain regions (Strausfeld and Hirth 2013). Although a somewhat controversial claim, research is nonetheless accumulating that points to the honeybee as being an ideal model for understanding cognition, perhaps capable of crossing the invertebrate border. The honeybee is thus a central theme throughout this thesis.

1.1 Thesis Structure

The thesis is structured in the following way. Chapter 2 gives a detailed literature review which aims to introduce the core principles which have driven this work. Discussed is computational neuroscience and a few sequential sampling models, honeybee neurobiology and the proboscis extension reflex paradigm, the role of inhibitory signals in decision-making, and how motivation can impact decision-making behaviours. Following this are three results chapters.

Chapter 3 describes an abstract model of decision-making. More specifically, it is a model of behavioural switching which can be applied to foraging animals looking to balance their nutritional intake. This balance requires the decision of whether to consume one specific type of nutrient or another. This model builds upon an older model of behavioural switching and combines it with an inhibitory neural mechanism which is well known in computational neuroscience and has been documented heavily in previous research. The introduction of this mechanism is shown to improve decision-making and assist the modelled animal in foraging efficiently. The model makes the prediction that animals should switch between alternatives irregularly if they wish to reduce the costs they encounter.

Chapter 4 presents a behavioural experiment using the famous proboscis extension reflex paradigm, which was devised to examine how honeybee decision- making changes according to motivational state. Differences in behaviours, reaction times and accuracies are attributed to differences in the satiation level of the bees. The behavioural data obtained can be described by the DDM. This model has been applied to humans as well as other mammals; this chapter shows that it can also be applied to honeybees, suggesting that simple perceptual decisions in invertebrates may be solved by the use of mechanisms which also exist within the mammalian brain. This is the first time the drift-diffusion model has been applied within a motivational context.

Chapter 5 describes a new computational model of decision-making which is based on the honeybee brain. The model implements the olfactory system of the bee as a higher-level network, with a particular focus on how groups of neurons interact and how specific groups contribute to the decision-making process. This model predicted the behavioural data in Chapter 4 successfully and indicates what neural circuits may be crucial for decision-making. It builds upon the neural mechanism discussed in Chapter 3 as well as a previous model of decision-making which also implemented this mechanism.

Finally, Chapter 6 summarises the work that has been undertaken here and discusses the limitations. Future directions of research are also proposed.

At the time of writing, some of the work described in Chapter 3 has been published (Marshall et al. 2015) and the results of Chapter 4 and 5 have been combined into a single manuscript, which is in preparation. Some of the ideas in Chapter 4 and 5 were published in another paper which reviewed invertebrate decision-making in light of vertebrate studies (Barron et al. 2015).

Chapter 2

Literature Review

Over the years, computational models have proven to be remarkably valuable to many fields of research and especially in the area of perceptual decision-making. Used to make predictions and replicate real world data, models have provided important insights into how systems work. Indeed, they have been described as '*increasingly essential to systems neuroscience*' (Cleland and Linster 2005, p. 801). A well-known example is the integrate- and- fire neuron model, which aims to simulate spiking neurons and demonstrates how the membrane potential of a neuron changes in the presence of excitatory or inhibitory inputs (see Burkitt 2006 for a review). When the membrane potential of the neuron reaches a certain threshold, an action potential (spike) is fired. The simplicity of this model has made it useful for investigating neural dynamics, as noted in the review by Burkitt: '*Focusing on the subthreshold membrane properties and excluding the mechanisms responsible for generating action potentials... has proven to be a powerful tool in understanding the information processing capabilities of neurons*' (Burkitt 2006, p. 1). Other neural models of varying complexity exist, such as the biologically plausible Hodgkin–Huxley model (Hodgkin and Huxley 1952), and a neural model which offers both the plausibility of the Hodgkin–Huxley model and the efficiency of the integrate-and-fire model has been proposed (Izhikevich 2003). As with neuron models, there are numerous models of decision-making that vary in their complexity

and biological plausibility.

2.1 Decision-Making Models

The field of psychophysics is credited to Gustav Fechner and has been traditionally concerned with the accuracy of decision-makers without considering other behavioural elements such as reaction time. In particular, experimenters have been interested in the discrimination capabilities of subjects and how this changed when stimuli were systematically varied to be increasingly similar or dissimilar.

For a detailed account of the experiments that were performed before the rise of sequential sampling models, see the review by Heitz (Heitz 2014). As noted in this review, the first experimental account of the relationship between decision time and accuracy was provided by Henmon in 1911 (Henmon 1911), but it was the work in statistics which really progressed the conception of the framework of the SAT. Stone (Stone 1960) is credited for the first mathematical model model of the decision process, which made use of the SPRT developed by Wald (Wald et al. 1948) and applied it to the assumption that a decision-maker will accumulate evidence in a task. Perhaps the most compelling experimental evidence for the sequential sampling framework is the work of Shadlen and Newsome in the 1990's; their neural recordings from monkeys in a visual task heavily suggest that the decision-making process in the brain resembles sequential sampling (see Chapter Four).

Mathematical models known as sequential sampling models (which implement the non- probabilistic sampling technique sequential sampling as their underlying principle (Gold and Shadlen 2001, Busemeyer and Johnson 2004, Usher and McClelland 2004, Teodorescu and Usher 2013) thus assume that decision-making is a process whereby noisy evidence is accumulated over time. They have been used to explore the underlying neural mechanisms of decision-making and the consequences of such mechanisms on an animal's speed and accuracy. For example, evidence is assumed to be integrated to a boundary

or threshold, and a change in threshold level has been taken to be '*the basis of speed-accuracy trade-off*' (Ratcliff 1978, p. 65). These models originate from psychology, the most well-known of them being the drift-diffusion model (DDM) which was pioneered by Roger Ratcliff in the 1970's (Ratcliff 1978). Since then there have been many different models proposed which vary in their biological plausibility and can be divided into certain categories such as linear integrators, race models or non-linear attractor models (Stone 1960, Vickers 1970, Link and Heath 1975, Ratcliff 1978, Smith and Vickers 1988, Ratcliff and Rouder 1998, Van Zandt et al. 2000, Usher and McClelland 2001, Wang 2002, Mazurek et al. 2003, McMillen and Holmes 2006, Ratcliff et al. 2007, Brown and Heathcote 2008, Niwa and Ditterich 2008, Ratcliff and McKoon 2008, Wang 2008, Albantakis and Deco 2009, Bogacz 2009, Ditterich et al. 2010, Krajbich and Rangel 2011). Indeed, sequential sampling models have proven to be popular for modelling choice behaviour. As noted by Smith and Vickers, '*the attraction of sequential sampling models is that they provide a description of the relationship between sampling time and performance accuracy, and hence are natural candidates for modeling speed-accuracy tradeoff effects*' (Smith and Vickers 1988, p. 135). Many of these models have been shown to fit certain experimental data well, however, many have failed to explain (and in some cases, have even predicted the opposite of) the results of other experiments (Teodorescu and Usher 2013). Due to the broad scope of models that have been proposed, it is difficult to differentiate between them and ascertain what properties or mechanisms are indeed crucial for the decision process. The schematics of some well-known classical models are given in Fig 2.1.

Usually, these models are designed to fit the two-alternative forced-choice, or 2AFC, paradigm. Here, a subject is presented with two options simultaneously, one of which is the 'correct' choice, and must choose one. The two options can also be presented sequentially (in which case the paradigm becomes the two-interval forced-choice or 2IFC). In the free-response paradigm, a decision is reached in the decision-maker's own time, which corresponds to

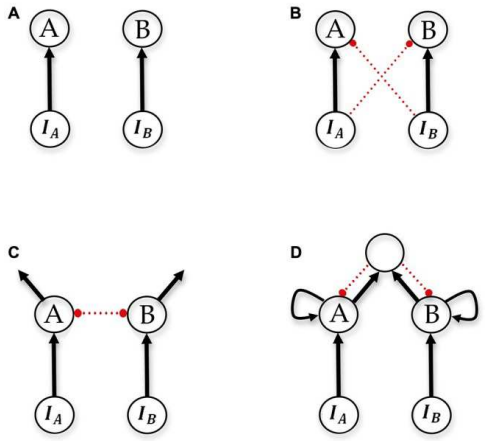

Figure 2.1: Schematics of several accumulator models. A: the race model (Vickers 1970) with uncoupled pathways. B: feed-forward inhibition (Ditterich et al. 2003). C: leaky competing accumulators (Usher and McClelland 2001). D: pooled inhibition (Wang 2002). Circles denoted with 'A' and 'B' are the neural populations which are integrating noisy sensory evidence. These are known as the decision populations as their activity determines whether or not a decision is triggered. 'I_A' and 'I_B' are the sensory populations which present the initial evidence. Note that these models make use of inhibition in varying ways. Black arrows denote excitatory connections, red dotted lines with circles denote inhibitory connections. Figure from Barron et al. 2015.

the decision threshold being reached. In the interrogation paradigm, the decision threshold is typically thought of as being discarded, and the decision is dependent on the accumulator that has integrated the most evidence when the trial ends. The development of the 2AFC and indeed the field of psychophysics is attributed to experimental psychologist Gustav Theodor Fechner. The 2AFC paradigm can be used to test the discrimination abilities of a subject and can be made harder by presenting two stimuli which are perceptually more similar. It has been observed that a focus on decision-making in the binary domain is a simplification of real world decision- making, which will undoubtedly present more than two alternatives at a time. However, it has been argued that two-alternative forced-choice tasks are still representative of the choices that animals will make in their everyday lives (Bogacz et al. 2006) and, furthermore, they can be well studied in laboratory studies. In an attempt to generalise sequential sampling models to a broader range of

tasks, some multi-alternative models have been developed (Bogacz et al. 2007, Tsetsos et al. 2011), however, they are not covered in this literature review. We here instead focus on binary decision-making tasks as they are simpler to model and, as mentioned, sequential sampling models have mostly been developed with the 2AFC in mind (Bogacz et al. 2006). In the following section, the models are presented in order of their complexity, starting with the one-dimensional (single integrator) Drift-Diffusion Model.

2.1.1 Drift-Diffusion Model

Remarked as being 'one of the cornerstones of modern psychology' (Milosavljevic et al. 2010, p. 437), the drift-diffusion model (DDM) is a statistically optimal (Laming 1968) model of decision-making (Ratcliff 1978, Stone 1960) that demonstrates the trade-off between accuracy and decision-speed. The DDM is an implementation of the sequential probability ratio test (SPRT), which has been shown to implement optimal decision-making (Wald et al. 1948) in the sense that, for a given accuracy, a decision can be made within the shortest possible time. In the same way, for a given decision-speed, the model will achieve the greatest possible accuracy. The behavioural task determines whether speed or accuracy will be optimised; if the optimisation of speed is required, the decision-maker's accuracy will decrease, demonstrating the speed-accuracy trade-off (likewise, if accuracy needs to be optimised, the decision-maker's speed will decrease).

Other models which are more complex and are able to be reduced down to the DDM are thus optimal in this sense (Bogacz et al. 2006). This particular model has kindled some interest recently as it has been shown to better fit reaction time data than other non-optimal models (Ratcliff and Smith 2004), and, furthermore, it can be applied to value-based decision tasks (Milosavljevic et al. 2010, Krajbich and Rangel 2011). Conversely, other researchers propose that the model is too simple to be useful and put forward experimental data that the model cannot account for, thus highlighting its limitations (Pirrone

et al. 2014).

The DDM is a one-dimensional decision-maker, *viz.* it has one integrator which accumulates the noisy difference between two alternatives over the duration of a trial. A decision is made when the accumulated difference crosses one of two thresholds; the value of these thresholds being set by the experimenter (but they are generally thought to be positive and negative in value, since the point of equal evidence is at zero). Varying the decision thresholds can demonstrate the trade-off between speed and accuracy; a threshold that is closer to the initial integration point induces a shorter reaction time and a threshold that is further away induces a longer reaction time, as a higher amount of evidence accumulation is required to reach it. Generally, accuracy is decreased and increased respectively. The decision-maker itself can be set to be more accurate (more experienced) or less accurate (less experienced); the former is able to make correct decisions more quickly, the latter tends to make more mistakes.

An alternative implementation of this model is known as the single bound drift- diffusion model, and this can represent go/no-go tasks instead of 2AFC tasks (see Cohen-Gilbert et al. 2014 for a recent example with human participants). In go/no-go tasks, the subjects are required to either respond or not respond to a stimulus, as such the negative threshold (representing the response to an alternative stimulus) is removed or set to zero. The process of evidence accumulation remains the same. A simulated example of a single bound DDM is given in Fig. 2.2. In this figure, the lower threshold is set to zero; in each simulation, the modelled animal can either respond to the presented stimulus or instead refrain from responding.

Whilst the single bound variant of the drift-diffusion model remains to be useful for the examination of decision-making in go/no-go tasks, its implementation generates a problem with ambiguity. It is not possible to ascertain whether the subject has refrained from responding due to the fact that it was not able to reach a decision in the given time (i.e, was still deciding), if the

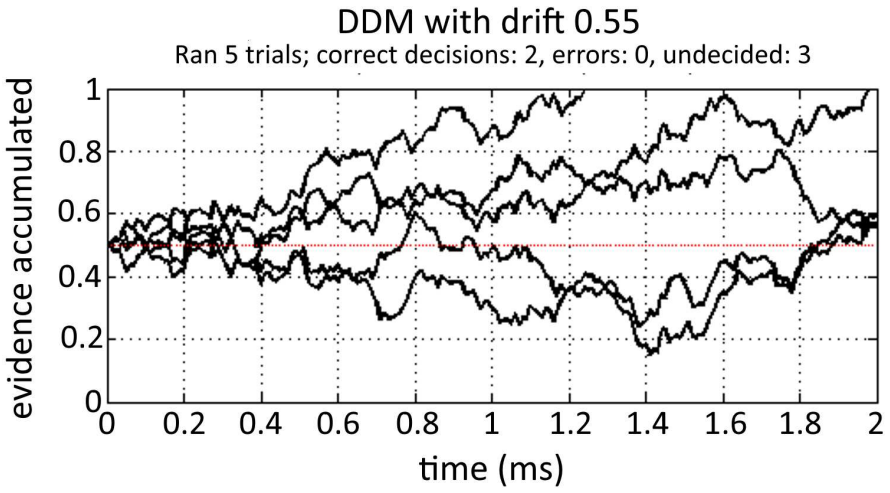

Figure 2.2: An example of a one-dimensional DDM simulation, with absorbing boundaries (author's own simulation). Here, the lower decision threshold has been set to zero and the higher decision threshold has been set to one. Over five individual trials, the decision-maker has made two correct choices (shown by the integrator reaching the top threshold twice) and has remained undecided for the other three trials (where the integrator has not reached a threshold within the time limit). The integrator's preference for an alternative demonstrates the strength of the drift as the difference in the evidence is integrated, with small fluctuations towards each threshold. The bias of this decision-maker is slightly towards being accurate (with a drift of 0.5 denoting no bias), however, this can be increased or decreased via alteration of the drift parameter (see main text).

subject had specifically chosen not to respond, or if the subject refrained from responding for some other reason, such as uncertainty or disengagement. As such, it is not possible to define whether or not the subject has made an error in each simulation.

Mathematically, the drift-diffusion model is defined by the following differential equation (Bogacz et al. 2006):

$$dx = Adt + cdW, x(0) = 0 \tag{2.1}$$

where x represents the difference in evidence that has been accumulated, Adt represents the average evidence increase and cdW represents noise, with dW representing a Wiener process (a continuous, stochastic diffusion process, also called Brownian motion). The initial amount of evidence integrated is usually set to zero, but it can be set to a non-zero value in order to encode a pre-

experimental bias towards one of the alternatives, thus making the likelihood of choosing one alternative higher than the likelihood of choosing the other. This can alter the accuracy of the decision-maker, and in such situations, the decision-maker is usually able to choose an option within a shorter time.

A variant of the DDM, known as the attentional drift-diffusion model (aDDM) applies the model to a task involving attentional fixation (Krajbich et al. 2010, Krajbich et al. 2012). As such, evidence accumulation is entirely dependent on the visual attention of the decision-maker. In this way, evidence for an alternative can only be integrated when the subject is looking at said alternative. When an alternative is out of focus, evidence is not accumulated or accumulated to a lesser extent. Consequently, a subject spending a lot of time looking at an alternative is integrating a lot more evidence for that alternative, as such, it is more likely to be chosen. The aDDM was found to be predictive of reaction times within a purchasing task (which combined visual stimuli with numerical) (Krajbich et al. 2012).

2.1.2 Ornstein-Uhlenbeck

The Ornstein-Uhlenbeck (OU) model is a variant of the DDM (Busemeyer and Townsend 1993). The OU process itself was proposed many decades before (Uhlenbeck and Ornstein 1930), originally as an alternative to the diffusion process model. It introduces the additional term λ , which controls the magnitude of how dependent dx is on the current value of information accumulated. Mathematically, it is defined in the following way (again using the equations derived from Bogacz et al. 2006):

$$dx = (\lambda x + A)dt + cdW, x(0) = 0 \tag{2.2}$$

with the parameters as before for the drift-diffusion model. The parameter λ alters the rate of accumulation; setting this to zero thus reduces the OU model to the DDM as the rate would be equivalent to that of the DDM. As such, this model is also statistically optimal. Setting $\lambda \neq 0$ accelerates x towards

one of the two thresholds, dependent on the value of λ.

2.1.3 Race Model

The DDM and OU models both incorporate one accumulator to integrate evidence. In contrast to this, the race model (Vickers 1970) assumes that evidence for each alternative is being integrated by independent accumulators. As such, evidence for n alternatives is being integrated separately by n accumulators. The accumulators are not coupled in any way and are not leaky, and they 'race' to reach the decision threshold. Thus, the outcome of the decision is determined by which accumulator reaches the threshold first. The race model cannot be reduced to the DDM (Bogacz et al. 2006) and as such it is not an optimal decision- maker. Despite this limitation, race models have been praised for their biological plausibility: '*they offer a parsimonious, generic and neurobiologically plausible mechanism by which actions can be selected on the basis of either a cued single action (specified) or a cued choice between actions*' (Rowe et al. 2010, p. 893). Formally, the race model is defined as:

$$\begin{aligned} dy_1 &= I_1 dt + c dW_1 \\ dy_2 &= I_2 dt + c dW_2 \end{aligned} \tag{2.3}$$

where I_n denote input strengths and cdW_n denote noise, as with the DDM, however, the noise is independent for each accumulator. Here, there is no 'drift', A. The likelihood of the alternatives is typically equal and the decision-maker is assumed to have no initial preference for either, i.e., $y_1(0) = y_2(0) = 0$.

The race model has been used to predict perceptual and behavioural decision-making and has also been tied in with action selection (Rowe et al. 2010). In the study of Rowe and colleagues, an experiment with humans was coupled with the development of a race model with competitive accumulators. Human subjects were asked to select actions without any indication of which

actions were correct, a paradigm based on an earlier one using monkeys. Using fMRI scans, brain activity of the subjects was recorded and the activity of certain areas correlated with the model predictions. Thus, the race model was used in determining parts of the human brain associated with action selection.

Recently, the race model has also been applied to a behavioural experiment which explored the role of confidence in value-based decision-making (De Martino et al. 2013), where confidence was defined as the '*degree of subjective certainty in having made the best choice, which equates to choosing the higher valued item*' (De Martino et al. 2013, p. 105). In their behavioural experiment, twenty human participants chose between food items that they would eat later, and indicated their confidence in their choice. The model predicted that an increase in confidence would result in a decrease in reaction time, and this was indeed found to be the case. The results showed that the race model was able to describe the relationship between confidence, reaction time and the difference in value of the alternatives and matched the observed behavioural data.

2.1.4 Pooled Inhibition

The pooled inhibition model (Wang 2002) was proposed as a more biologically plausible accumulator, originally introduced as a model for the posterior parietal cortex, an area in the mammalian brain which guides saccadic eye movement. For a TAFC task there are two pools of *decision neurons* which accumulate evidence for each alternative. These pools of neurons compete with each other (with the activity of one pool directly influencing that of the other pool) and have *self-excitatory recurrent projections* (see Fig 2.1), such that the neurons are able to maintain activity without stimulus input for some time, and also decay in activity at a slow enough rate for integration to take place (Bogacz et al. 2006).

In addition to these decision neurons there is also a pool of shared inhibitory neurons. Both pools of decision neurons are able to excite these inhibitory

neurons, and the inhibitory neurons in turn inhibit the decision neurons via recurrent connections (Bogacz et al. 2006). This was introduced as a way of inhibiting the decision neurons' activity without them inhibiting each other directly. With these three pools of neurons, the pooled inhibition model is a three-dimensional model and is formally defined in the following way (Bogacz et al. 2006):

$$\begin{aligned} dy_1 &= (-ky_1 - wy_3 + vy_1 + I_1)dt + cdW_1 \\ dy_2 &= (-ky_2 - wy_3 + vy_2 + I_2)dt + cdW_2 \\ dy_3 &= (-k_{inh}y_3 + w'(y_1 + y_2))dt \end{aligned} \tag{2.4}$$

with dy_1 and dy_2 representing the decision neurons, dy_3 representing the inhibition neurons, v denoting the self-excitatory recurrent projections, w' denoting the weights between the decision neurons and inhibition neurons and k denoting the decay of activity (with k_{inh} representing the decay of activity in the inhibition neurons). I_1 and I_2 denote input units with strength I. Finally, cdW_1 and cdW_2 denote white noise.

2.1.5 Limitations

The classical computational models of decision-making, although incredibly useful, are not without their flaws. Taking the DDM as an example, whilst this model can quite elegantly show the trade-off between speed and accuracy, it fails to demonstrate reasonable behaviour under certain conditions. For instance, given a choice between two alternatives, both of which are equivalent (and assuming that the decision-maker is not biased towards either alternative if value is taken into account), the model will be reduced to simply integrating noise over time until randomly ending up at one decision threshold or the other. A real world decision-maker would be expected to instead immediately choose either at random (Pais et al. 2013). This limitation prompted the development of more complex models which were able to solve this problem

(Brown and Holmes 2001, Pais et al. 2013).

2.2 The Honeybee as a Model of Cognition

It was noted by James McConnell in his 1966 review on invertebrate learning that '*man is generally more interested in man than in any other animal...*' (McConnell 1966, p. 107) and the decision-making abilities of human subjects has been thoroughly studied within a wide range of tasks. Over the years, it has been heavily debated whether or not invertebrates are capable of learning, however, the accumulation of research to date heavily suggests that they are. Recently, it has become clear that the honeybee in particular has a rich behavioural repertoire. Research has shown that this insect is not merely hardwired; that is to say, it is able learn and adapt its behaviour to new situations. This capacity for plasticity was noted in Randolf Menzel's recent review on honeybee research: '*Honeybees contradict the notion that insect behaviour tends to be relatively inflexible and stereotypical*' (Menzel, 2012, p. 758). Indeed, recent studies into honeybee learning and memory have shown that this insect may be an ideal candidate as a model for learning. They are able to rapidly learn Pavlovian associations (for example, Bitterman et al. 1983) and are capable of other types of non-elemental learning such as contextual learning (see Giurfa 2003b for a detailed review).

Karl von Frisch was famous for his work on honeybee behaviour and his discovery of the intricate waggle dance, which honeybee foragers use to communicate to other bees within the hive the location of food sources (Von Frisch 1967). Indeed, the honeybee is a eusocial creature which is required to make decisions both independently and as part of a group, with one example of collective decision- making being house hunting (Seeley et al. 2012). Foraging requires the bee to navigate several kilometres in search of flowers, thus it must not only find its way to food sources once outside the hive but must also return safely; and it is for this reason that honeybee navigation and spatial learning has been so greatly studied (for example, see Srinivasan et al. 1996), along

with other invertebrates that demonstrate such learning, such as the digger wasp (Tinbergen 1972) and desert ant Wehner 2003. Furthermore, honeybees see in colour and have demonstrated impressive feats in visual learning and memory (for a detailed review see De Ibarra et al. 2014 as well as Zhang et al. 2012b). Additionally, they are able to sense odours acutely and they perform remarkably well in discrimination tests (Guerrieri et al. 2005). For their everyday success and ability to survive despite their limited neural circuitry, '*the insect brain must therefore provide intelligent solutions to a wide range of ecologically relevant problems...*' [p. 62](Menzel and Giurfa 2001) and, as such, they function well as models of decision-making and learning.

2.2.1 The Proboscis Extension Reflex

Classical or Pavlovian conditioning, originally studied through behavioural experiments with dogs, was pioneered by Ivan Pavlov back in the late 1920's (Pavlov 1927). It is used to train an animal to associate an initially neutral stimulus (the conditioned stimulus or 'CS') with another stimulus which is naturally rewarding (usually, some kind of food reward, referred to as the unconditioned stimulus or 'US'). After a single reinforced trial, the animal begins to associate the reward with the neutral stimulus, and, as a consequence of this, it will display behaviours towards the neutral stimulus which otherwise would not have been invoked (in the case of the original experiments, salivating at a bell when no food was present). Today, this form of conditioning remains an important tool in the study of animal behaviour, and it has been adapted to suit many different experimental tasks within different sensory modalities. Indeed, it has also been applied in the invertebrate realm.

Work on classical conditioning with honeybees started in the late 1950's when Matsutaro Kuwabara discovered that they could be trained to associate coloured lights with a sugar reward (Kuwabara 1957). This discovery was made by use of the Proboscis Extension Reflex (PER) paradigm, an accessible approach to the study of honeybee learning which requires bees to be har-

nessed rather than free-flying (Kuwabara 1957, Takeda 1961, Bitterman et al. 1983, Felsenberg et al. 2011, Giurfa and Sandoz 2012). In this paradigm, an experimenter gently touches the antennae of a bee with a toothpick or similar object which has sucrose solution on. This will invoke the proboscis extension reflex from the bee, which is a behavioural response whereby the bee extends its proboscis (tongue). After a few trials of classical conditioning, where a stimulus is presented just prior to the presentation of the sugar, the bee will extend its proboscis to the neutral stimulus alone. Unfortunately, using the PER paradigm with stimuli in the visual domain has proven difficult as it requires the bees' antennae to be cut. For quite some time this was overlooked; consequently the results from Kuwabara's work were unreproducible (Giurfa and Sandoz 2012). Furthermore, cutting the antennae for visual conditioning results in a drop in learning performance, as shown by the experiments conducted by Hori and colleagues (Hori et al. 2006; 2007). This is a result thought to arise from a reduction in bees' responsiveness to sucrose (de Brito Sanchez et al. 2008). Furthermore, in order for the bees to learn the association, the training phase needed to be extended to span across two days.

A popular alternative to visual stimuli in PER conditioning is olfactory stimuli. In 1961, Kimihisa Takeda developed the PER protocol which used odourants instead (Takeda 1961) and it was later built upon by Bitterman and colleagues who implemented proper controls (Bitterman et al. 1983). Experiments over the years have shown that bees are able to perform well in olfactory conditioning trials, rapidly learning the association between odour and sugar reward after a single trial and reaching high performance rates after around five trials. For this reason, PER is usually performed using olfactory as opposed to visual stimuli. However, recent studies have shown that conditioning with visual stimuli is possible in bees, but a change in harnessing is required (Dobrin and Fahrbach 2012). Another study examined the role of motion cues in visual conditioning and found that the performance of bees improved when the trained colours were presented along with visual stimuli (Balamurali et al.

2015). Interestingly, this study also found that intact bees performed equally as well as bees that were deprived of their antennae. Another study in visual PER demonstrated that bees are capable of learning visual associations but that they are unable to discriminate between colours well, even though the colours used in the experiments had shown to be discriminable by bees in free-flying experiments (Niggebrügge et al. 2009).

The protocol for olfactory PER conditioning has now been standardised (Felsenberg et al. 2011). Standardisation is important for ensuring that significant differences found between results have not arisen due to differences in experimental protocol. In addition, results derived from studies making use of the standardised protocol should be directly comparable. PER conditioning has enabled researchers to examine honeybee decision-making and discrimination behaviours with a high degree of control over the experimental subjects, to the point where neural recordings can be made as they respond to odours (Smith and Menzel 1989). For example, an individual's level of satiation can be manipulated precisely in PER studies; this sort of control is not possible in free-flying experiments. A very thorough study by Guerrieriand colleagues serves as an example of how discrimination behaviours can be evaluated in great detail using PER (Guerrieri et al. 2005). In this study, the paradigm was used to evaluate honeybee discrimination using sixteen different odours. The odours were classified in terms of their functional group (alcohols, aldehydes and ketones) and chain length (between six and nine carbon atoms), which made task difficulty very easy to manipulate systematically. For example, odours within the same functional group will have a higher degree of similarity between them. Odours with closer chain lengths will also be more similar to each other. Thus, an aldehyde of chain length six (Hexanal) is more similar to an aldehyde of chain length seven (Heptanal) than eight (Octanal). In their experiment, the authors used a total of 1,457 bees to produce a behavioural matrix which demonstrated how the bees were generalising across odours. The bees were each trained to a single odour using the PER

paradigm and then tested with four alternatives, one of which might have been the trained odour. The more bees that responded to the alternatives, the more perceptually similar to the trained odour the test odour was considered. The full generalisation matrix shows the percentage of proboscis extensions recorded across all sixteen odours and demonstrates that, whilst the bees mainly responded to the trained odourant, they also responded to perceptually similar ones. Furthermore, the authors went on to show that the perceptual distances of the odours were correlated with those found from neural recordings: '*The correlation between both datasets was highly significant, thus indicating that odours that are encoded as physiologically similar are also perceived as similar by honeybees*' (Guerrieri et al. 2005, p. 719).

Despite the attractiveness of the PER paradigm, it does not come without its limitations due to how it has been applied across different labs. Over the years, researchers have adapted the methodology to suit their own experiments and introduced small changes in inter-trial interval (ITI), the length of time the odour is presented for, and a multitude of other parameters, which in turn may impact experimental results. This was noted by Giurfa and Sandoz: '*...subtle modifications in experimental parameters such as inter-trial intervals ... among others, may lead to radically different conclusions, some of which may be misleading*' (Giurfa and Sandoz 2012, p. 58). In a recent review on visual conditioning in bees, Avarguès-Weber and Mota also write: '*...the literature provides results obtained with different visual conditioning protocols that are rarely comparable and sometimes conflictive.*' (Avarguès-Weber and Mota 2016, p. 108). The recent standardisation of the procedure is a step towards minimising this drawback. Other influences on PER results, such as the effects of the season, ITI and the number of trials used in training, have been considered and discussed thoroughly in a publication by Frost, Shutler and Hillier (Frost et al. 2012).

Despite its limitations, the PER paradigm has proven to be invaluable to honeybee research, especially in the evaluation of learning and memory (for

example, see Scheiner et al. 1999). Since its establishment in 1961, it continues to be widely used to date and has even been adapted for use in experiments with other insects such as ants (Guerrieri and d'Ettorre 2010). It is a paradigm well suited for the examination of olfactory discrimination behaviours and is here used in our behavioural experiments with bees.

2.2.2 The Speed-Accuracy Trade-off

In decision-making tasks, the accuracy of an individual under different experimental conditions has been studied extensively. In humans, it has been shown that there is a relationship between the time taken for a person to complete a decision- making task and their accuracy. Subjects who make quicker decisions are more prone to error whilst those who take longer are more accurate. This result has also been found in other animals, and more generally it suggests that the sampling time required to solve a decision task is related to the accuracy of an agent. This is known as the speed-accuracy trade-off (SAT). It has also been shown that there is a relationship between the *difficulty* of a decision-making task and the time taken by a subject to complete it, though this is not always the case. In a task where accuracy is critical and errors are penalised, an animal will sacrifice speed in order to make a more reliable decision, however, within a different task this sacrifice may be costly, as such sampling time depends heavily on both the individual and the context of the task. The speed-accuracy trade-off has recently been shown to exist in bees within a variety of different contexts such as foraging and house hunting. Indeed, according to the theory behind the drift-diffusion model, the SAT is an unavoidable phenomenon for optimal decision-makers.

In 2003, Chittka, Dyer and Dornhaus demonstrated that bumblebees exhibit a speed-accuracy trade-off (Chittka et al. 2003). Within their experimental paradigm, a colour discrimination task using projected virtual flowers, bees that made decisions quickly made more errors whilst bees that were slower were more accurate. Furthermore, bees were shown to have consistent

individual differences in how quickly they made a decision, which indicates that bees have their own individual foraging strategies. Interestingly, it has been shown that faster bees collect nectar more efficiently than slower bees, though this holds only when foraging patches have a higher proportion of flowers that are rewarding (Burns 2005, Burns and Dyer 2008). It has been suggested that these individual differences are beneficial to a colony as multiple foraging strategies will reduce the variability in nectar collection rate and in turn promote colony fitness (Burns and Dyer 2008). If the natural diversity of a colony's environment is taken into account, it is perhaps logical that having individuals employing their own foraging behaviours will be beneficial as they will be targeting different patches of flowers. Thus, individual differences in foraging strategies are an important factor when designing discrimination tasks; such differences will have an impact on an animal's accuracy (Burns 2005).

The experiment by Dyer and Chittka in 2004 built upon the premise that a colour discrimination task becomes increasingly difficult as colour differences are reduced, with colour difference defined as '*the Euclidean distance between stimuli loci in colour space*' (Dyer and Chittka 2004b, p. 761). They investigated how bees decided to forage on certain flowers with relation to how hard they perceived the task to be, also supporting the hypothesis that bees alter the time they take to make decisions depending on task difficulty.

In their visual discrimination task, bees were trained to forage from artificial flowers (plastic discs), with distractor flowers delivering water as opposed to sugar water. Task difficulty was controlled by colour distance from the target stimulus; distractor flowers that were more similar in terms of colour thus corresponded to a harder decision task. They found that bees very quickly learned to discriminate flowers that were low in similarity, needing fewer visits to flowers in order to solve the task. For harder tasks, there was a sharp drop in accuracy and bees needed more visits to flowers before they could begin to learn to discriminate between them. Their experiment confirmed

that between individuals there is a speed-accuracy trade-off; for harder tasks, bees that made decisions more quickly were also less accurate. Additionally, they analysed the bees as a group and found that they were slowing down to complete harder tasks.

The results from these studies suggest that there exists a speed-accuracy trade- off in bees; difficult discrimination tasks result in longer sampling times for the animal, which suggests that assessment of stimulus characteristics in the brain improves over time, perhaps due to noise averaging (Wright et al. 2009). Moreover, bees can be trained to perform better and discriminate more effectively in harder tasks by using differential conditioning.

In discrimination tasks it has often been assumed that an animal will desire to perform exceptionally well, and, as a consequence of this, errors have been taken to show the limit of an individual's capabilities. As stated by Chittka and colleagues, '*It is not clear whether low accuracy actually reflects the limits of discrimination*' (Chittka et al. 2003, p. 388). Indeed, this is often not the case (Chittka et al. 2009, Chittka et al. 2003, Giurfa 2004). In particular, discrimination tasks using differential conditioning have shown that an increase in cost for errors will often result in the animal slowing down to accumulate more evidence in tasks. For example, when foraging induces a risk of predation, bees take longer to inspect flowers and ensure that they are safe before landing (Ings and Chittka 2008). In an experiment by Chittka and colleagues, distractor flowers were developed to deliver aversive quinine solution as a punishment and in response to this the bees improved their accuracy (Chittka et al. 2003). When the punishment was removed from the task the accuracy of the bees decreased, indicating that the improved performance was not due to experience. The results suggest that an animal's discrimination capacity is not static, rather, it depends on both the individual as well as the context. As such, how bees are trained becomes an important contributing factor when examining their behaviour in discrimination tasks.

2.3 Honeybee Neurobiology

Advances in digital technology have encouraged the formation of 3D brain atlases, and there has been a rise in research studying invertebrate brains in an attempt to create complete neural maps. Perhaps the most well mapped invertebrate brain is that of *Drosophila*. The connectome or neural map of this insect has recently been constructed using a database of 23,579 neuron images (Shih et al. 2015). The honeybee brain has also been studied in great detail, and a computational atlas known as the Honeybee Standard Brain (HSB) has been developed (Brandt et al. 2005, Rybak et al. 2010). The HSB was created from the data of twenty individual worker honeybees, consisting of twenty-two neuropils, and serves as a representation of the average honeybee brain. It is a virtual, 3D map which is widely accessible, and it contains detailed information about the locations of specific neurons within the bee brain.

The olfactory system of the honeybee has been well studied (for example, see Zwaka et al. 2016), along with other brain regions which are suggested to play a role in learning, memory, decision-making or action selection. Here, the focus is on the olfactory regions but other higher order regions are briefly discussed. Where gaps are apparent in the honeybee literature, research in other insects such as *Drosophila* is substituted, if available, instead. It is argued that this is acceptable because the basis of decision-making and action selection, i.e., the convergence of competing motor commands for execution at a specific neuropil, is assumed to be conserved between the different species (Perry and Barron 2013b). Considering that action selection is an evolutionarily 'old' problem, this assumption should indeed hold. Strausfeld noted the similarities of the higher order regions, the mushroom bodies, in his paper: '*comparisons between insect groups suggest that within an order there are highly conserved features of the mushroom body shape and lobe arrangements*' (Strausfeld et al. 1998, p. 15). Indeed, in Shih and colleagues' recent paper on the *Drosophila* neural map, the authors suggested that, even for this simplistic insect, '*...the overall organizational scheme showed fundamental similarities to the network*

structure of the mammalian brain' (Shih et al. 2015, p. 1249). Strausfeld and Hirth recently suggested that vertebrate and invertebrate action selection centres (the basal ganglia and central complex, respectively) may be homologous, however, this suggestion has been seen as controversial (Strausfeld and Hirth 2013). If similarities exist between *Drosophila* and mammalian brains, these similarities are expected to be found in honeybee brains. Undoubtedly however, there will be specific differences in the differing brain morphologies that may determine behavioural differences between similar insects. Thus, it can be argued that a review of the literature for multiple insects is a worthwhile cause, as common traits or mechanisms as well as species-specific differences can be identified. Indeed, the importance of 'comparative connectonomics', the '*quantitative study of cross-species commonalities and variations in brain network topology*', has very recently been brought to light (van den Heuvel et al. 2016, p. 345). It is worth mentioning that previous models of insect neurobiology have also used data from more than one inveretebrate (for example, see Smith et al. 2008 and Cope et al. 2016).

Due to a vast number of researchers publishing neurobiological data from many different arthropods, conflicting or otherwise confusing terminology has arisen as an unfortunate consequence. As such, different names have been given to equivalent brain regions across species. For example, the fan-shaped body and ellipsoid body are also referred to as the central body upper division and central body lower division, respectively. Recently, a group of researchers known collectively as the 'Insect Brain Name Working Group' have attempted to define a standard that can be applied across insect species (Ito et al. 2014), This nomenclature will no doubt help to ensure that future studies use the same terminology.

2.3.1 The Honeybee Olfactory System

In the honeybee, there are around 60,000 olfactory receptor neurons (ORNs) which provide the olfactory system with information about a presented odour's

identity. The ORNs then innervate the antennal lobes (AL), sites of olfactory processing (Galizia 2014), via four different tracts, which are named T1 - T4 (Abel et al. 2001, Kirschner et al. 2006, Nawrot 2012). The antennal lobes are structures composed of around 160 glomeruli. Each glomerulus collates odour information from ORNs which express the same olfactory receptors. Within the glomeruli, local interneurons (LNs) project to multiple other glomeruli but are constrained within the antennal lobes. Also within the antennal lobes are dendrites of projection neurons (PNs). Both excitatory and inhibitory PNs (ePNs and iPNs, respectively) output processed neural signals to higher order brain centres such as the lateral protocerebrum (LP) (Galizia 2014). The LP is an understudied site of convergence for processed olfactory input; it is thought to be pre-motor in nature and is indicated to play a major role in action selection (Galizia 2014, Barron et al. 2015). Projection neurons innervate these higher order brain regions via five different antennal lobe tracts, which are named the median (m-ALT), the lateral (l-ALT) and the medio-lateral (ml-ALTs) tracts (of which there are three) (Abel et al. 2001, Kirschner et al. 2006). A schematic diagram of the honeybee brain is given in in Fig 2.3, where the m-ALT, l-ALT and ml-ALT1 tracts are shown.

Some of the projection neurons receive olfactory signals from one glomerulus and are thus referred to as uniglomerular. There are estimated to be around 900 of these neurons, from both the m-ALT and l-ALT (Zwaka et al. 2016). Other projection neurons collate neural signals from several glomeruli and are thus referred to as being multiglomerular. These multiglomerular PNs have been shown to be mostly GABAergic (thus, inhibitory). These inhibitory projection neurons, the iPNs, project from the antennal lobes directly to the lateral horn (LH - a sub-region of the lateral protocerebrum indicated to play a role in innate decision-making (Heimbeck et al. 2001, Gupta and Stopfer 2012) via the three ml-ALTs, bypassing the mushroom bodies entirely. The uniglomerular, excitatory projection neurons (ePNs) innervate the mushroom bodies as well as the lateral horn via the m-ALT and l-ALT tracts. The former

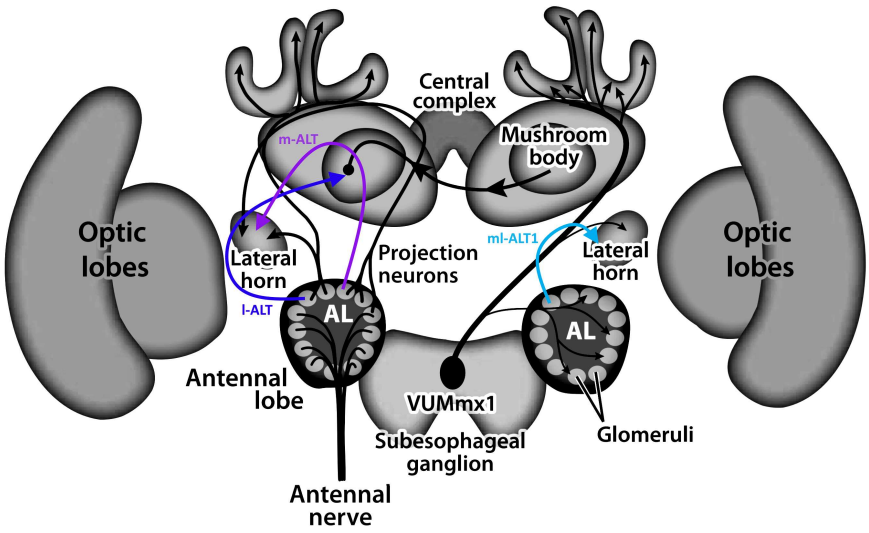

Figure 2.3: Schematic of the honeybee brain from the frontal position. Shown in colour are three antennal lobe (AL) tracts: ml-ALT1 (where projection neurons innervate the lateral horn (LH) from the antennal lobes), m-ALT (where projection neurons innervate the mushroom bodies and then the lateral horn) and l-ALT (where projection neurons innervate the lateral horn and then the mushroom bodies). Local interneurons within the antennal lobes are not shown. Antennal lobe tracts ml-ALT2 and ml-ALT3 are also not shown. It should be noted that the projections shown here are bilaterally symmetrical. Figure adapted from (Perry and Barron 2013b).

tract innervates the MB first before heading to the LH; the latter does the reverse and innervates the LH before the MB. Recently, some of the differences between the m-ALT and l-ALT neurons have been observed. For example, they differ in volume as well as in neural length (Zwaka et al. 2016). It should also be noted that the l-ACT and m-ACT neurons differ in latencies, with l-ACT neurons being quicker to respond to stimuli but also being 'unspecific' (Müller et al. 2002). As the l-ACT neurons project to the LH first, it might be that these neurons are for innate responses or attention. Additionally, l-ACT neurons have been shown to habituate faster. The slower responding m-ACT neurons project to the mushroom body calyces (cup shaped structures that have been identified as input regions to the mushroom bodies, with inner regions termed as the basal ring, the collar and the lip (Mobbs 1982)) first, perhaps encoding odours specifically for identification (Müller et al. 2002, Galizia 2014). It has been suggested that the l-ACT neurons may be involved

with learning (Müller et al. 2002).

Within the invertebrate literature, the paired neuropils known as the mushroom bodies were first brought to light in 1850 (Dujardin 1850). Since their discovery, they have been the primary point of focus in many experiments on learning, memory and cognitive abilities in insects (for a recent review, see Menzel 2014). The discovery that honeybees do not simply exhibit pre-programmed behaviour led to a rise of experiments focusing specifically on learning and memory, which arose in the 1970's and 1980's (for example, Menzel and Erber 1978, Erber 1981, see also Menzel and Muller 1996 and Giurfa 2007 for more heavily detailed reviews); this coupled with the result that the cooling of the mushroom bodies impaired memory has led to research mostly focused on these regions (see Heisenberg 1998). The mushroom bodies have been suggested to be involved with decision-making (Mizunami et al. 1998b, Grünewald 1999, McGuire et al. 2001, Schröter and Menzel 2003, Akalal et al. 2006). Indeed, they have been found to integrate multiple forms of sensory information (Menzel and Muller 1996). That they are sites which take input from more than one sensory domain adds strength to the hypothesis that they play an important role in decision-making.

It was found in the early 1980s that cooling the lateral protocerebrum had no effect on memory (Menzel 1983) and as such this neuropil (and the ml-ALT) has been heavily overlooked, despite the fact that it exists as a pre-motor area (i.e., Gupta and Stopfer 2012, Galizia 2014) and seems therefore to be crucial in decision-making processes too (Barron et al. 2015). It is within the lateral protocerebrum that evidence accumulating pathways are hypothesised to converge (Barron et al. 2015), and the activity of this region is thus suggested to determine motor output. The convergence of evidence pathways is of particular interest as it may be comparable to the action selection circuit of the basal ganglia.

As mentioned before, recently proposed by Strausfeld and Hirth (Strausfeld and Hirth 2013) was the idea that the vertebrate basal ganglia and invertebrate

central complex (another region thought to be important to action selection, see Plath and Barron 2015) are homologous, based on comparisons of gene expression and anatomical similarities, for example. Although a controversial suggestion, these similarities are nonetheless worthy of investigation.

We have seen that an animal's performance in a choice task is dependent upon how sensory information processed by the brain. There are many factors which affect this, one such example being task difficulty. However, it is also the case that behaviour can be influenced by internal state, indeed, without this perhaps we would not see behaviour at all: '*The field of motivation is concerned with what animates living organisms, that is, what makes them go*' (Wright 2016, p. 16). What the term 'motivation' precisely describes is still being debated upon. In humans, three categories of motivation have been proposed, one of which is the motivation to survive (Reeve 2016). Evidently, this category of motivation does not just apply to humans. Within this category is hunger, which undoubtedly also plays an important role in laboratory experiments where subjects are responding to food rewards (for example, in the PER paradigm). As such, we next explore how an animal's level of satiation can influence its behaviour.

2.4 The Impact of Satiation

Studies in animal decision-making have shown that choice behaviours are robust and can be influenced by the internal state of the animal. Hunger is perhaps the most well studied of all the possible motivational influences as it is somewhat simpler to control. Indeed, an animal's level of satiation has been shown to mediate certain behaviours, especially those associated with feeding. The influence of hunger is crucial to study, as animals being observed within choice tasks are often responding to food rewards, as such, differences in their motivational levels could cause differences in their behaviours. For humans, hunger has been termed as one of multiple 'visceral factors' which have the ability to influence and impair rational decision-making (Loewenstein

1996), which suggests that starvation or extreme hunger can heavily impact an individual's decision-making process.

The impact of satiation level can be well studied for a variety of animals. Indeed, starvation has been shown to produce behavioural switching and induce searching behaviours in insects (Bell 1990). The sea slug, *Hermissenda*, was shown to consume food faster when starved (Avila et al. 1998). In threespine sticklebacks, extreme hunger caused fish to invest in more 'risky' behaviours, showing an increased tendency to perform inspections on possible predators (Godin and Crossman 1994). As with the sea slug, threespine sticklebacks were also found to feed at an increased rate. In crucian carp, hungry fish were more likely to sacrifice safety than satiated fish, again demonstrating more 'risky' behaviours (Pettersson and Brönmark 1993). Rather than opting for safety, the fish showed a preference for open, unsafe habitats which offered a feeding area. Moreover, the carps' foraging behaviour was shown to impact predator-avoidance, demonstrating that several internal states may be interacting during the time a decision is made. Indeed, starvation has also been shown to alter mating behaviours. In *Microvelia austrina*, a semi-aquatic insect, starving males were shown to mate for shorter periods of time than satiated males (Travers and Sih 1991). In *Mesocyclops edax*, a genus of crustaceans, starving organisms were shown to alter their prey preference (Williamson 1980). It was shown that satiated organisms initially avoided *Bosmina*, a type of water flea. However, after several days of starvation, *Mesocyclops edax* began to demonstrate a preference for *Bosmina* instead. In the crayfish *Orconectes virilis*, starved animals demonstrated an increase in activity (Hazlett et al. 1975). Furthermore, animals that had been starved for one week tended to be more aggressive. After being starved for two weeks, both the physical activity and the aggressiveness of the starved animals decreased, most likely because the animal has undergone extreme starvation and has either reached a critical point wherein it is no longer able to function or is dying. Recent studies into the unicellular slime mould, *Physarum poly-*

cephalum, show that even this brainless organism exhibits behaviours that are mediated by internal state (Latty and Beekman 2011). Again, as with the sea slug and threespine sticklebacks, starving slime moulds displayed more 'risky' behaviours and ventured into well lit (aversive) environments in order to reach a more plentiful food source, whilst non-starving organisms preferred to remain in darker environments despite less food being available.

Taken together, all these studies demonstrate how internal state, in this case more specifically the level of satiation, can alter animal behaviours. This research also suggests that the driving force of hunger can even interact with other internal states, demonstrating that multiple internal states can have an impact on behaviour. Indeed, it is clear that an animal's hunger will have a profound impact on its decision-making, especially if the animal is responding to a food reward.

With honeybees, it has already been shown that sucrose responsiveness has an impact on learning performance. Individuals which show a high sucrose responsiveness demonstrate a higher level of acquisition than individuals with a low responsiveness (Scheiner et al. 2004) and it has been shown that food intake can mediate sucrose responsiveness (Pankiw et al. 2001). It can be inferred from this result that food intake can impact learning and decision- making as a whole. Furthermore, food intake has been shown to impact memory formation. Using the Proboscis Extension Reflex paradigm, Friedrich and colleagues showed that learning performances were highly dependent on when individuals were fed (Friedrich et al. 2004). Their study linked the cAMP-PKA cascade with satiation and demonstrated that memory formation is dependent on satiation level; it was found that feeding individuals four hours prior to conditioning would impair memory formation. These studies heavily suggest that there is a link between satiation, learning and memory, which in turn suggests that the mushroom bodies may play a crucial role. Additionally, the latter study measured the basal PKA activity from the central brain, as such this higher order centre may also play a role in conveying satiety sig-

nals: '*Only satiation status, not sucrose responsiveness, affects the basal PKA activity in the brain tissue*' (Friedrich et al. 2004, p. 4464). Notably, a single interneuron 'VUMmx1' has been shown to respond heavily to sucrose and learnt CS presentations in another study which used the PER paradigm; this neuron innervates the lateral protocerebrum as well as the lips and basal rings of the mushroom bodies (and the glomeruli of the antennal lobes) (Hammer 1993b).

We have now covered how decision-making and action selection can be influenced and how this can be observed at the behavioural level (for example, animals are riskier when hungry and are less accurate in more difficult discrimination tasks). Another topic of interest is how the brain processes stimuli during decision-making and how this processing can be influenced by neural mechanisms. As such we now proceed to investigate an inhibitory mechanism which is employed by the brains of both vertebrates and invertebrates, which has been hypothesised to aid in decision-making tasks.

2.5 Lateral Inhibition

Research into the existence of lateral inhibition in arthropods can be traced back to studies on horseshoe crab visual processing in the 1950's. One study demonstrated that optic nerve fibres within the eye could be inhibited by the activity of neighbouring fibres: '*...the frequency of the discharge of impulses in a single optic nerve fiber is decreased and may even be stopped by illuminating areas of the eye in the neighbourhood of the sensory element from which the fiber arises*' (Hartline et al. 1956, p. 651). Lateral inhibition has also been shown to exist in the fly (Zettler and Järvilehto 1972, Kirschfeld and Lutz 1974, Strausfeld and Campos-Ortega 1977) and cat, through an experiment which examined orientation detectors (Blakemore and Tobin 1972). By the 1970's, lateral inhibition in the visual system was referred to as a 'common phenomenon' (Kirschfeld and Lutz 1974).

These experiments all conclude that lateral inhibition exists within the

visual system. Research has also been performed on other systems within the brain, and some studies have found that a similar mechanism exists for the olfactory system. Indeed, lateral inhibition has been shown to exist within the olfactory bulb as well as the antennal lobe (Urban 2002, Mori et al. 1999). Importantly, some of these studies suggested that lateral inhibition was not static and could be modified by the inputs presented. For example, a study using *Drosophila* found that the strength of lateral inhibition between glomeruli scales with ORN strength, such that it is reduced in the case of weak ORN input (Olsen and Wilson 2008).

A study by Wilson and Laurent suggests that temporal patterns in neurons are dependent on what odour has been presented (Wilson and Laurent 2005). The study focused on *Drosophila* and demonstrated that GABAergic inhibition increases the differences between neural representations of odours, essentially acting as a decorrelator. Importantly, the authors note that two different odours can induce very similar activation levels in a projection neuron initially, but that this similarity decreases later on with $GABA_B$-mediated inhibition. It was suggested that inhibition arises due to the activation of other glomeruli. The study also shows that temporal patterns are not present in ORN responses to odours, thus it has been hypothesised that they arise due to antennal lobe processing.

The temporal differences in neural responses prompted another study to examine how this might be happening within the brain. The research conducted by Linster and colleagues (Linster et al. 2005) built on older studies both in mammalian olfactory bulb and insect antennal lobe, where, as mentioned before, it has been shown that inhibitory networks are important for olfactory processing. The authors raised the point that the exact mapping of these inhibitory networks had not been clearly defined or organised. Their work suggests that it is the response profiles of the glomeruli that determine the lateral connectivity between these neural structures. Indeed, their computational model of the honeybee antennal lobes points to lateral inhibition being

mediated by the response profiles of the glomeruli; their network wherein lateral inhibition is proportional to the similarity of the response profiles was able to reproduce experimentally derived results (namely the output of PNs). The study reiterates the point that odour representations in the brain are both spatially and temporally defined since representations are more dissimilar when leaving the antennal lobes than when entering them as ORN input. Here, it is proposed that the antennal lobes are not merely suppressing a lesser activated glomerulus and that lateral inhibition is not based on spatial location or proximity, rather, it is a function of glomerulus response profiles.

The publication by Linster and colleagues was cited in a recent review by Galizia, in which the processing of the antennal lobes is discussed heavily (Galizia 2014). Here, the results of the study are reinforced and it is again suggested that lateral inhibition within the antennal lobes is not uniform. Galizia also cited a study by Chou and colleagues which focused on *Drosophila* and showed that glomeruli with narrower tuning properties are less innervated by other glomeruli LNs. The fact that some glomeruli have narrower tuning also suggests that they have less overlap with the response profiles of other glomeruli, which in turn may highlight a decrease in lateral inhibition necessary for the discrimination of odours (and hence less innervation).

In 2007, Schmuker and Schneider inferred that '*insect and vertebrate olfactory systems can be subdivided into three stages of functional organization*' (Schmuker and Schneider 2007, p. 20285). They argue that the initial stage is where stimulus attributes are encoded by neurons into neural signals, the second is where the stimulus representations are decorrelated, and in the third stage the signals are associated with specific qualities. They then asked if such a system could be generalised. In their publication, they were able to design a computational model that could process chemical information using only these three computational principles. They too support the hypothesis that the decorrelation of neural signals in the second stage assists in stimulus classification. In their model, they implemented correlation-based lateral inhibition

and agreed that processed output should be more dissimilar as a result.

A study in *Drosophila* in 2008 also suggested that neural responses to odour stimuli have both spatial and temporal properties, however, they did not find evidence for 'response-sharpening', which may be attributed to experimental limitations (the odours they used activating fewer glomeruli, for example) or a difference in the olfactory processing systems of bees and fruit flies (Silbering et al. 2008).

It is important to mention that behavioural evidence of lateral inhibition has even been found in honeybee collective decision-making. It was shown that lateral inhibition - implemented via headbutts which are referred to as stop signals - functioned as a deadlock breaker which allowed the swarm to make a choice as a whole (Seeley et al. 2012). A computational model of value sensitive decision-making reinforced the results from this study, whereby the strength of lateral inhibition was the deciding factor in whether or not a decision was made (Pais et al. 2013). The model was also sensitive to the relative value of the alternatives; if the options were equally poor then the model would wait in case superior alternatives were presented later, whilst if the options were equally good, one would be chosen at random. Taken together, we can see that lateral inhibition exists as an important mechanism not only in the brain, but also in collective decision-making behaviours.

2.6 Summary

In this literature review, the fundamentals of various different topics have been covered. Though quite broad, all these topics have contributed to this work. In particular, the impressive abilities of honeybees have been highlighted, along with their efficient neural circuitry. Additionally, the motivation for modelling the decision process as an accumulation of evidence as has been presented, and a range of classical decision-making models described. The impact of satiation on choice behaviour was covered briefly, emphasising that decision-making processes are highly influential. Finally, the importance of lateral inhibition

was introduced. This inhibitory mechanism is explored in the next chapter, which presents an abstract model of action selection where a modelled animal needs to choose between two alternatives whilst foraging. The mechanism of lateral inhibition is adapted from one classical decision-making model in particular, the leaky competing accumulator model, and is applied to this model of foraging to demonstrate how this mechanism is beneficial to choice behaviour.

Chapter 3

The Role of Inhibition in Decision-Making

During its lifetime an animal will need to make important decisions, some of which will be critical for its survival. Taking the honeybee as an example, this animal must decide not only where to forage but also which flowers it should target. There will be various factors that will influence these decisions, such as context (for example, the time of day), environmental or social information (for example, choosing to forage in a location that has been communicated by a hive mate), and internal motivational state. If a honeybee, or indeed any animal, is presented with two alternatives, how do the motivations of the animal contribute to its choice? The question perhaps invokes the paradox of Buridan's ass, wherein a donkey finds itself exactly halfway between food and water. The donkey is equally hungry as it is thirsty, and since its motivational drives are equivalent it should choose the closer option. Stuck in an infinite loop which can never be resolved, the animal remains undecided and eventually dies.

The question of how foraging animals solve the problem of choosing between alternatives remains to be fully answered. Previous research has indicated that foraging behaviours can be modelled using the 'geometric framework' (Simpson and Raubenheimer 1993, Simpson et al. 2004), whereby an animal

attempts to balance its nutritional intake. The animal moves through n-dimensional nutritional space and acts so as to bring its current nutritional state closer to that of a target state. This target has been defined as the nutritional state that will ensure the highest reproductive success for the animal (Houston et al. 2011) but it could also be defined as the state which maximises growth rate and development (Helm et al. 2017), dependent on the developmental state of the animal. The foraging problem thus becomes how the animal decides to move through nutritional space in order to reach the target, with performance determined by the distance from its current state to that of the target. When the nutrients the animal is consuming do not interact in any way, this can be found simply by taking the Euclidean distance.

Behavioural switching, defined here as moving from the consumption of one type of nutrient to another, is a crucial aspect of this problem as it allows an animal to reach the target state. Choosing between food and water, with the assumption that both of these options offer only one nutrient, can thus be imagined in two dimensional space, and the decisions of the animal should move it closer to the target point. How might it choose to move through this nutritional space? The motivation of the animal, here defined as 'the tendency to eat or drink' (Marshall et al. 2015), will sway it one way or another: an animal that is more hungry than thirsty will be expected to forage for food as opposed to water. Upon finding nourishment the animal should then feed until its corresponding motivation has been reduced sufficiently. As the animal becomes satiated it will experience the motivation to drink; as such the animal will, eventually, switch from feeding to instead seeking water. As the animal consumes one type of nutrient, it reduces its deficit for that nutrient. In previous work, a deficit has been defined as 'the quantity of food or water that the animal will ingest, under ad libitum conditions, until it is satiated' (Sibly 1975). Here, a food or water deficit is thus the shortage of the respective nutrient from the target. Consuming one type of nutrient will move the animal along a linear path through nutritional space, as such switching

from one behaviour to another can be imagined as a series of linear foraging bouts.

With this scenario in mind, it could be assumed that an animal will simply act on its strongest motivation, however, there are costs associated with decision-making that should first be considered. Depending on the environment, food and water resources may not be in nearby locations, thus switching from one activity to another will incur costs in terms of using resources to travel from one location to another. It becomes apparent that there is a trade-off between an animal reducing its motivations and reducing the cost of behavioural switching. Of course, in real world foraging situations there are other costs to consider. In the case of the honeybee, a novel rewarding flower may hide a predator; as such choosing to forage on it would increase the risk of predation. Here, a model of animal choosing between two nutrients is implemented to address the question of how the cost of behavioural switching can be reduced.

A foraging animal choosing between two alternatives is comparable to a two- alternative forced-choice task (2AFC), a paradigm originating from psychophysics which was developed to analyse a subject's choice and discrimination behaviours (Fechner 1966). The subject is presented with two alternatives and is required to choose one of the two (as such, choosing neither is not an option). Usually implemented within a laboratory setting, this paradigm allows for a thorough analysis of the subject's reaction time, among other behavioural traits, and can be easily manipulated such that the task is easier or harder. Although a simplification of real world decision-making tasks, the 2AFC paradigm is useful as it is analytically simple and is still representative of some of the decision problems that could be made by an animal. Perhaps one of the main differences between real foraging situations and a 2AFC task is the cost of switching between options. Due to time and energy investments, a foraging animal will incur a cost if it chooses to switch from one alternative to the other. However, in 2AFC tasks, the cost of switching will either be

minimal or the animal will not be given such an option. In some 2AFC tasks, once the animal has made a decision the task ends and another is presented, as such switching is not possible. Nevertheless, the foraging task modelled here can perhaps be described as an ecological 2AFC task with switching costs.

Due to the simplicity of the 2AFC paradigm, recent advances in computational neuroscience have provided insights into what types of neural circuitry may give rise to certain behaviours in binary decision-making tasks. Indeed, a variety of decision-making models, known as sequential sampling models, have been described and analysed in great detail (Ratcliff and Smith 2004, Bogacz et al. 2006) and are able to fit data obtained from behavioural experiments, making the underlying theories behind them applicable to other models, such as the foraging model discussed here. One of the main findings from these models is that the coupling of evidence accumulation pathways is beneficial for decision-making. The mechanism of lateral inhibition, or cross-inhibition, whereby the competing accumulators inhibit one another mutually, has been shown to be especially beneficial. Within the sequential sampling models, this was a prominent feature of the leaky competing accumulator (LCA) model, discussed in more detail below. This mechanism is here adapted and applied to the motivations of the modelled animal, as such, the benefits of inhibitory coupling is examined within a foraging context which can be comparable to a 2AFC with switching costs.

It is important to note that, for this particular model, lateral inhibition is not being applied between neurons or neural units. Although comparisons can be made between the LCA model and this model of foraging (they both apply lateral inhibition as a form of coupling to benefit decision- making, for example), they are not equivalent. Here, the benefits of lateral inhibition are examined by analysis of how the animal reduces its deficits over time and its ability to reduce the costs it incurs as it switches from one alternative to another. Despite these differences, lateral inhibition as a form of coupling has been well explored within the realm of sequential sampling models, as such a

brief overview of the LCA is given here.

3.1 The Leaky Competing Accumulator

The leaky competing accumulator (LCA) is a biologically-inspired model of decision-making (Usher and McClelland 2001). This model is notably different from earlier decision-making models such as the race model (Vickers 1970) due to the nature of the way that the accumulators interact. In the race model, evidence accumulating pathways are independent, such that they never interact, and are 'racing' towards a joint decision threshold; when one accumulator reaches this threshold, a decision will be made. However, it has been shown that this implementation is suboptimal (Bogacz et al. 2006). The race model is outperformed by other models which make use of coupled pathways, and these models may furthermore be more neurally plausible. Indeed, the race model is unable to approximate the well-known drift-diffusion model (Ratcliff 1978), another sequential sampling model which has been shown to be optimal (optimality here is achieved by obtaining the highest accuracy for a given reaction time, or the fastest reaction time for a given accuracy). Under certain parametric conditions, however, the LCA can approximate the drift-diffusion model and so exhibit optimality (Bogacz et al. 2006).

The LCA model couples the evidence accumulation pathways through a mechanism known as 'mutual inhibition' or 'cross inhibition'. Here, accumulators are connected so that they inhibit each other, which introduces competition into the network. After the model is allowed some time to integrate sensory evidence, the end result is that one accumulator will have a very high firing rate whilst the other will not, which is beneficial to a decision-maker (note that the higher its activity, the more it suppresses the alternative accumulator). This mechanism has been shown to aid decision-making previously (Bogacz et al. 2006, Marshall et al. 2009) and, remarkably, it has also been documented in insect collective decision-making (Seeley et al. 2012, Pais et al. 2013, Reina et al. 2015). The LCA model furthermore incorporates a leak or

decay to the accumulated evidence; as such the removal of sensory evidence will cause a drop in accumulator activity and the firing rate will return to the baseline. Recently, cross-inhibition has been shown to aid decision-making in a model of animal foraging (Marshall et al. 2015), the methodology and results of which will be presented in this chapter.

Formally, the dynamics of the LCA model are described by two coupled differential equations:

$$\begin{aligned} dy_1 &= (I_1 - ky_1 - wy_2)dt + c_1 dW_1 \\ dy_2 &= (I_2 - ky_2 - wy_1)dt + c_2 dW_2 \end{aligned} \tag{3.1}$$

where y_i denotes the activity levels of the accumulators, I_i denotes the mean activity of the sensory neurons, k denotes the strength of the neural leak, w denotes the strength of the mutual inhibition and dW_i denotes Gaussian-distributed noise with zero mean and root-mean-square strength c_i (Bogacz et al. 2006).

For a two-alternative forced-choice task, the LCA is a two-dimensional model with two accumulators, each integrating evidence in support of one alternative or the other. The model can also be extended for multiple alternatives (Usher and McClelland 2004, McMillen and Holmes 2006) however, this is not discussed here. This model can be applied to decision tasks within both the free-response and interrogation paradigms (Bogacz et al. 2006). In the former, the accumulators continuously integrate evidence until one crosses a decision threshold. The first accumulator to reach this threshold is the accumulator that determines the final decision. In the latter paradigm the accumulators are allowed to integrate evidence for a certain amount of time, which is decided by the experimenter. At decision time, the trial ends and the accumulator with the highest integrated evidence determines the decision. Both of these experimental paradigms were made use of in a recent experiment using a biomimetic robot (Lepora et al. 2012). For the task, the robot

was required to identify the shape and location of a presented object using a biologically inspired sensing system: artificial 'whiskers' which imitate those of rodents. Within the free-response paradigm, the robot was able to make a decision in its own time, once it had accumulated enough evidence. Within the interrogation paradigm, the robot was expected to make a decision after a pre-set number of whisks. The study found that the robot was able to obtain similar decision accuracies in both of these paradigms.

3.2 Materials and Methods

The LCA mechanism of lateral inhibition is here applied to a previous model of behavioural switching (Houston and Sumida 1985) in which an agent is deciding between seeking water or seeking food. This model attempted to solve the problem whereby an animal services one motivation (for example, thirst) such that its 'drive' is reduced to a value just below that of the alternative drive. This would result in the animal switching to the alternative activity and then servicing the alternative motivation for a short time before switching again (since the two competing motivations will be unequal but in close proximity to each other). This behaviour was referred to as 'dithering' or 'chattering' and is an example of costly behavioural switching (Houston and Sumida 1985). The model implements the notion of positive feedback, such that performing an activity would not only reduce its associated drive but also increase its motivation. The rise occurs quickly after the agent has switched to a new activity and is mediated by an upper bound. This way, a drive is pushed to its maximum when the animal switches and is not close to the value of the other. This solution prevents the problem of dithering, and also the problem whereby the drives are equivalent and the model cannot choose between them, akin to the paradox of Buridan's ass.

Since the mechanism of lateral inhibition, as described with the LCA model, has been shown to improve decision-making, it is here applied such that the two motivations of the animal are coupled. Unlike the LCA model,

the model here is deterministic (as such noise has been removed). These motivations are formally described by a pair of partial differential equations, as with the original LCA model:

$$\begin{aligned} dv_1 &:= (c_1 dx_1 + c_2 x_1 + c_3 v_1 + c_4 v_2)dt \\ dv_2 &:= (c_1 dx_2 + c_2 x_2 + c_3 v_2 + c_4 v_1)dt \end{aligned} \tag{3.2}$$

where v_i denotes the level of the i-th motivation (for example, seeking food), x_i is the level of the i-th deficit and dv_i and dx_i denote their rates of change. Additionally, c_1 denotes the influence of dx_i on dv_i, c_3 denotes the strength of motivational decay and c_4 denotes the strength of cross-inhibition. To see the equivalence between this model and the LCA, note that the first two terms of the equation denote the sensory information (I_i), the third denotes the leak of the accumulators (k) and the fourth denotes the strength of inhibition (w). Here, cross-inhibition is introduced by setting $c_4 < 0$; setting $c_4 > 0$ thus removes it. It should be noted that this model, unlike the LCA model, is deterministic (without noise). Under certain parametric conditions ($c_2 > 0, c_3 = -c_2, c_4 = 0$) the model is equivalent to the original model of Houston and Sumida (Marshall et al. 2015). Thus it is intuitive that using these parameters and setting $c_4 < 0$ is equivalent to introducing cross-inhibition into the original model.

In each simulation, the agent initially begins equidistant from two locations which offer either food or water. The agent then moves to the location corresponding to its strongest motivation. An agent motivated by hunger will thus move towards the location where food is present. There, it reduces its corresponding deficit (feeds) until the strongest motivation changes. Whilst the agent is performing action i, the corresponding deficit (x_i) decreases at the rate of $dx_i = -(g - h)dt$, where g is the rate of reduction of the deficit and h is the rate of increase of the deficit. When the agent is not performing action i, the deficit x_i increases at the rate of $dx_i = h$. As in the original model, the

dynamics of increasing deficits are ignored and it is assumed that $h = 0$ (Houston and Sumida 1985). Once the agent is more strongly motivated by another deficit, it ceases its current activity and moves from its current location to the other. Whilst moving, it cannot reduce either of its deficits. This behavioural pattern continues until the agent has fully reduced its deficits. It is assumed that the animal will be interrupted during foraging, and the probability of this interruption per-unit-time is set to λ. After interruption, the agent is scored on its performance according to a penalty function, which is defined as the following:

$$p := x_1^2 + x_2^2 \tag{3.3}$$

A lower penalty thus determines a higher performance score for the agent. From the model simulations, parameters were identified which generated the optimal performance of the agent, for various switching costs and interruption probabilities. Additionally, c_1 was systematically varied with the cost of switching between two alternatives, τ. In order to maintain equivalence with the original model of Houston and Sumida, only c_1, c_3 and c_4 were varied, leaving $c_2 = -c_3$ and with the constraint that $c_3 < 0$. Performance variations due to λ and τ were deemed to construe a sensitivity analysis.

3.3 Results

It should be noted that these results are presented in Marshall et al., (2015) and that they are the work of the other co-authors of this publication. The analysis of the foraging bouts of the animal (presented in Fig 3.6) is original work and was used as supplementary material for the publication. Additionally, the non-linear extension of this model and the results that are presented in Section 3.4 are also original work and feature an adapted version of the model.

Simulations of the model were carried out whereby the cost of switching

and the rate of interruption parameters were varied. For each of these different set-ups, parameters c_3 and c_4 were varied systematically. At the end of each simulation, the penalty score was determined along with the foraging bout durations. Furthermore, the model dynamics were analysed in order to show how the deficits and motivations of the animal changed as time progressed. Taken together, these results ascertained which parametric set-ups gave the optimal performances. The penalty scores shall first be discussed. A low penalty is derived when the modelled animal is able to reduce frequent, costly behavioural switching, thus moving its nutritional state to that of the target without incurring costs. An example 2D plot of the penalty scores is given in Fig 3.1.

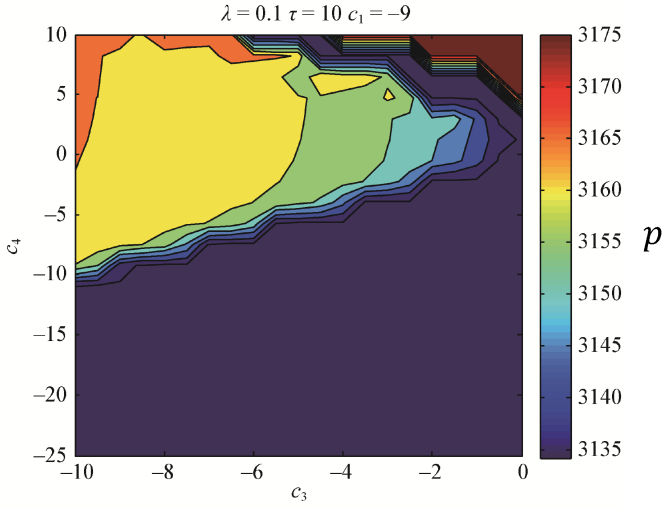

Figure 3.1: A 2D slice which shows the penalty scores obtained for the given values of λ (interruption probability), τ (switching cost) and c_1 (influence of dx_1 on dv_i). Parameter c_2 was set to $-c_3$, and parameters c_3 and c_4 are both varied systematically. Darker blues denote the lower penalty scores and thus a better performance from the animal. From Marshall et al. 2015.

The best penalty scores are obtained when the animal is modelled with cross- inhibition (i.e., where $c_4 < 0$) and also when $|c_4| > |c_3|$, as shown by the dark blue areas of the graph. Setting $c_4 > 0$ results in an increase in the penalty score. This was the case for all the examined variations of interruption rate and switching cost, demonstrating that the relationship between the c_i

parameters is conserved across parametric space. A more complete sensitivity analysis for the model is shown in Fig 3.2, where penalty scores are shown for different values of λ and τ, as well as c_1.

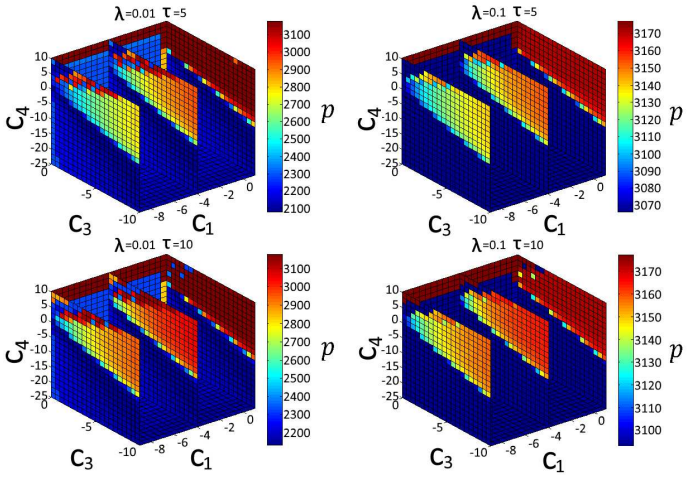

Figure 3.2: The full sensitivity analysis which shows the derived penalty scores for varying switching costs and interruption rates. As in the previous figure, darker blue areas denote lower penalties and thus better performance. From Marshall et al. 2015.

The full sensitivity analysis makes it clear that cross-inhibition generates the lowest penalties for the entire parameter space, again shown by the dark blue areas. As before, the dependence on c_3 is highlighted; when $|c_4| > |c_3|$ the decision penalties are greatly reduced. Furthermore, the result is held for changing values of c_1, λ and τ, demonstrating the importance of cross-inhibition and rate of decay on the model performance. The reason why these two parameters impact the penalty scores can be determined by analysis of the modelled animal's motivations and deficits over time. The results of the model show that the foraging behaviour changes when cross-inhibition is introduced. This is shown in Fig 3.3. Time in these plots begins at the top- right (as such, when $t = 0$ the motivations are at their highest) and increases along the x-axis. The original publication also makes use of these plots (Houston and Sumida 1985).

In the top-left plot of Fig 3.3, the motivational state of the animal without cross-inhibition fluctuates regularly across the behavioural switching line (dashed

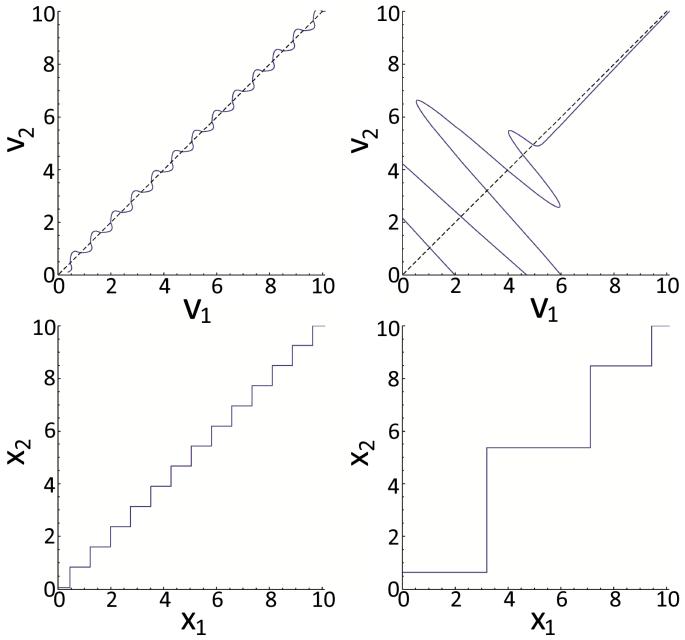

Figure 3.3: The animal's motivations and deficits in two different parametric situations. The top row depicts the motivations, v_i, and the bottom row the deficits, x_i, of a modelled animal with cross-inhibition (right) and without (left). Without cross-inhibition, the animal frequently switches between the two activities. The introduction of inhibitory coupling induces irregular foraging bouts and a reduction in behavioural switching. Time begins at the top-right of plots where motivations and deficits are highest, and ends at the bottom-left, as in (Houston and Sumida 1985). Parameters: $c_1 = -1, c_3 = -1$. Left (without cross-inhibition): $c_4 = 0$. Right (with cross-inhibition): $c_4 = -1$. Initial deficits and motivations: $x_1(0) = v_1(0) = 10, x_2(0) = v_2(0) = 10.1$. From Marshall et al. 2015.

line in the figure, where $v_1 = v_2$), which corresponds to the animal frequently switching back and forth between the two activities. In the top-right plot, the motivations are far more irregular and they allow for the animal to spend more time attending to a single nutrient. This is also shown in the bottom plots which depict the deficits. The introduction of cross- inhibition results in irregular foraging bouts so that the animal spends less time switching. This behaviour in turn produces a lower penalty score. The dynamics of the motivations and deficits can be examined in more detail with regard to the decay parameter c_3. This is shown in Fig 3.4.

It is obvious that the relationship between c_3 and c_4 produces differences in the animal's motivations and consequent deficit reduction. As demonstrated

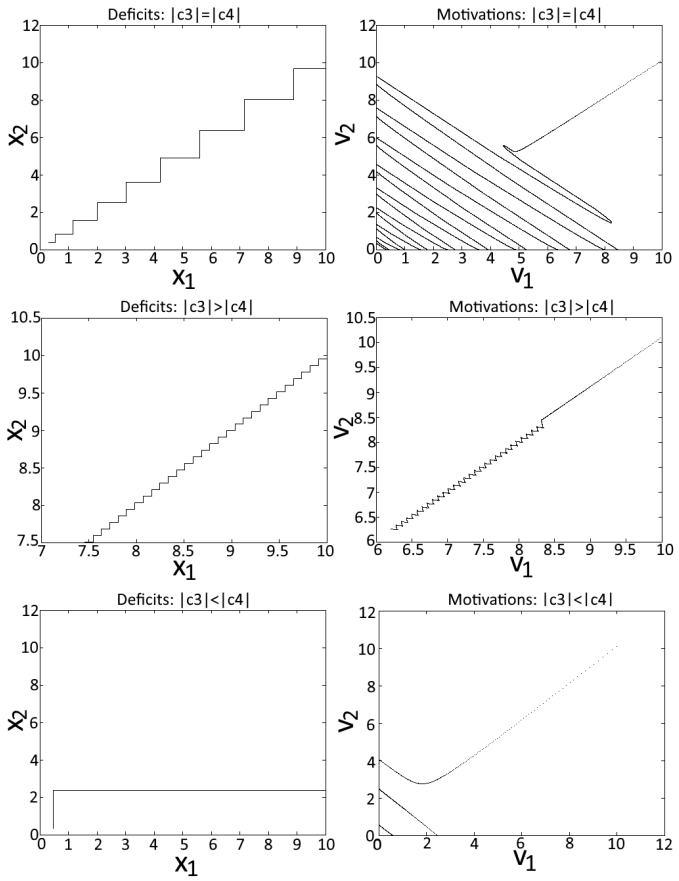

Figure 3.4: The animal's motivations and deficits for different combinations of c_3 (strength of decay) and c_4 (strength of cross-inhibition). Deficits are given on the left, motivations on the right. When $|c_3| > |c_4|$, the animal switches between the two activities frequently. If $|c_3| \leq |c_4|$, the animal instead exhibits irregular foraging bouts and a reduction in frequent switching. Top row: $c_3 = -5, c_4 = -5$, middle row: $c_3 = -10, c_4 = -2$, bottom row: $c_3 = -5, c_4 = -20$. All plots: $c_1 = -5$.

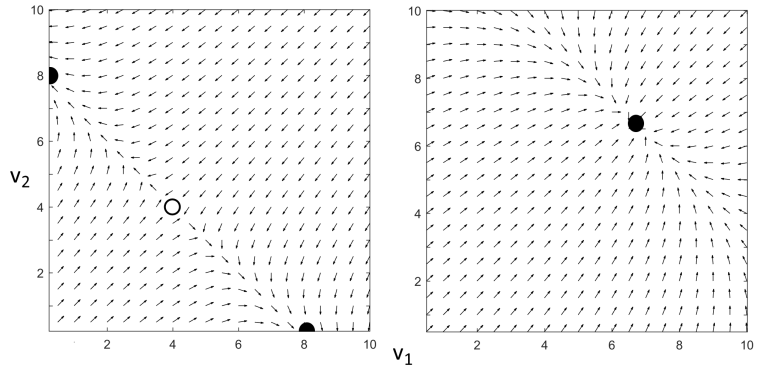

Figure 3.5: Plot to show how the dynamics of the modelled motivations change according to variations in the c_3 (strength of decay) and c_4 (strength of cross-inhibition) parameters. When $|c_3| < |c_4|$, the motivations of the animal end up in a state where one is suppressed whilst the other is not, which is beneficial to decision-making. Stable points in motivational states are denoted by black filled (semi) circles, unstable points are shown by open circles. Parameters: $c_1 = 0, c_3 = -2$. Left: $c_4 = -3$. Right: $c_4 = -1$.

in the previous figure, when $|c_3| > |c_4|$ the performance of the animal is impaired due to frequent switching, shown in the middle row. Setting $|c_3| \leq |c_4|$ drastically improves performance, especially when $|c_3| < |c_4|$. An explanation as to why the motivations of the animal change in such a way can be attributed to the model dynamics. To give further insight into these dynamics, we show them graphically in Fig 3.5.

The dynamics of the model are very different according to the relationship between c_3 and c_4. When $|c_3| < |c_4|$, the motivational states converge onto one of two stable points wherein one motivation is high whilst the other is suppressed. The motivations are accelerated to this state due to the single unstable point between them. However, when $|c_3| > |c_4|$, the motivations instead converge onto a single stable point where the they are very similar to each other (right plot in Fig 3.5). This in turn causes the animal to switch frequently.

In addition to the penalty scores, the animal's foraging bout durations were analysed over the running time of each simulation, in order to ascertain whether or not the derived behavioural data were plausible. In previous studies

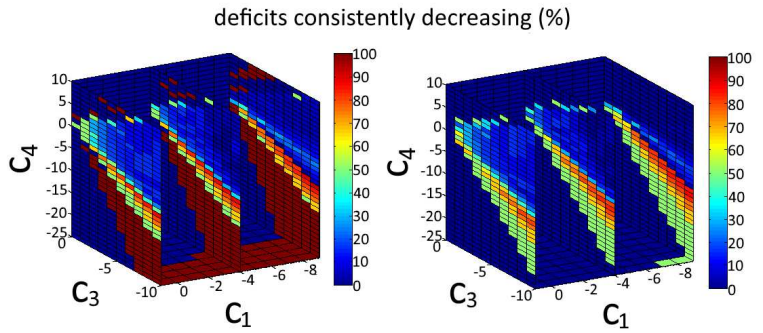

Figure 3.6: The full results for the foraging bouts analysis for systematically varied values of c_1, c_3 and c_4. The matrices show the proportion of bouts that are being consistently decreased over the simulation time, as such, the best performance is 100%, shown in darker red. If this result is obtained, the agent has successfully decreased the length of each foraging bout following the previous one, spending less time consuming nutrient i as deficit x_i decreases. The results are given for each deficit individually, on the left is the result for the x_2 deficit, on the right, the x_1 deficit. The results indicate that the best performances are derived when cross-inhibition is present within the model ($c_4 < 0$). The consistency of foraging bout reduction is also dependent on the c_3 parameter (strength of decay), as with previous results.

examining animal 'vilgilance', defined as alertness to stimuli (Dukas and Clark 1995), it has been found that such alertness cannot be sustained over long periods of time. In a model of foraging which incorporates vigilance as a parameter, it was found that vigilance should decrease over time as the animal forages and that the length of the animal's foraging bouts should decrease accordingly (Dukas and Clark 1995). As such, the modelled animal here should be reducing its bout durations as time progresses and its deficits are reduced. Hence, a more successful parameter setup will result in the agent being able to consistently reduce its bout durations over the entire period of the simulation. Fig. 3.6 shows the full results matrices whereby c_1, c_3 and c_4 are being varied.

It is apparent that the agent is able to fully decrease its bout lengths in some simulations for the x_1 deficit, however, it is unable to fully decrease its bout lengths for the other deficit. Despite this difference in performance, a similar pattern can be noted in both matrices. When cross-inhibition is introduced (where $c_4 < 0$), the proportion of bout length durations being reduced increases. As the strength of inhibition grows, the agent is able to fully reduce its bouts in a greater number of simulations. This indicates that

cross-inhibition not only improves decision-making but also is necessary for the agent to implement rational foraging behaviour.

3.4 The Non-linear Model

The model described thus far is a deterministic linear model, however, decision-making is a non-linear process (Bogacz et al. 2007). As such, a non-linear extension to the model was developed which builds upon the original model by altering the formal dynamics, such that the interaction between the motivations is no longer linear. The extension was originally planned to be non-deterministic, however, preliminary results suggested that the introduction of noise was not beneficial for the model. As such, even though noise would make the model more biologically plausible, it remains deterministic.

The non-linear model functions in precisely the same way as the linear model, where the modelled animal must choose between two alternative nutrients and act so as to minimise the distance between its current nutritional state and that of the target state. With the entirety of the model parameters being held constant, the only change is in the formal definition of the model. Like the linear model, it is defined by two coupled differential equations:

$$\begin{aligned} dv_1 &:= c_1 dx_1 + c_2 x_1 + c_3 v_1 + c_4(v_2 - (1/(v_1)^2)) \\ dv_2 &:= c_1 dx_2 + c_2 x_2 + c_3 v_2 + c_4(v_1 - (1/(v_2)^2)) \end{aligned} \tag{3.4}$$

Here, there is an additional coupling present in the inhibitory pathways. A single accumulator's activity is still inhibited by the other accumulator, however, this inhibition itself is stunted by a small amount of the accumulator's *own* inhibition. This coupling is demonstrated diagrammatically in Fig. 3.7. This introduces additional competition in the inhibitory pathways as well as the decision pathways. Thus, the strength of cross -inhibition is modulated or regulated by the agent's own motivations. This change was hypothesised to be beneficial to the model. In the original model, one motivation would

be accelerated and would inhibit the other accordingly; as such a very motivated animal would perhaps supress the competing motivation too rapidly. Mediation, however, would ensure that the inhibition is never too high.

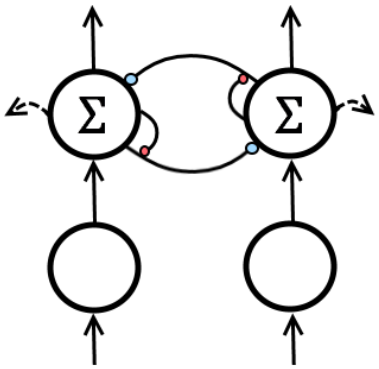

Figure 3.7: Schematic of the non-linear extension of the original model. Here, lateral inhibition between accumulating units is mediated by inhibition. Excitation denoted by arrows, leak denoted by dashed arrows. Inhibitory connection to competing units denoted by small blue circles, inhibition of lateral inhibition denoted by small red circles.

3.5 Results

As described before with the linear model, simulations were carried out where the model parameters were varied and the penalty scores determined. In addition, we also analysed the model dynamics and foraging bout durations of the animal to contrast with the results of the linear model. The full sensitivity analysis for the penalty scores is given in Fig. 3.8, where λ, τ and c_1 are varied along with c_3 and c_4.

It is apparent that the penalty scores for the non-linear model are very similar to those which were derived previously; the lowest penalties are obtained from simulations where cross-inhibition is introduced ($c_4 < 0$) and $|c_4| < |c_3|$. This result is not surprising. However, when cross-inhibition becomes particularly strong, there is a slight degradation in the performance of the agent and higher penalties are derived, dependent on the value of λ. Despite this, the 'wedge' of best performance is consistent with the linear model and is retained

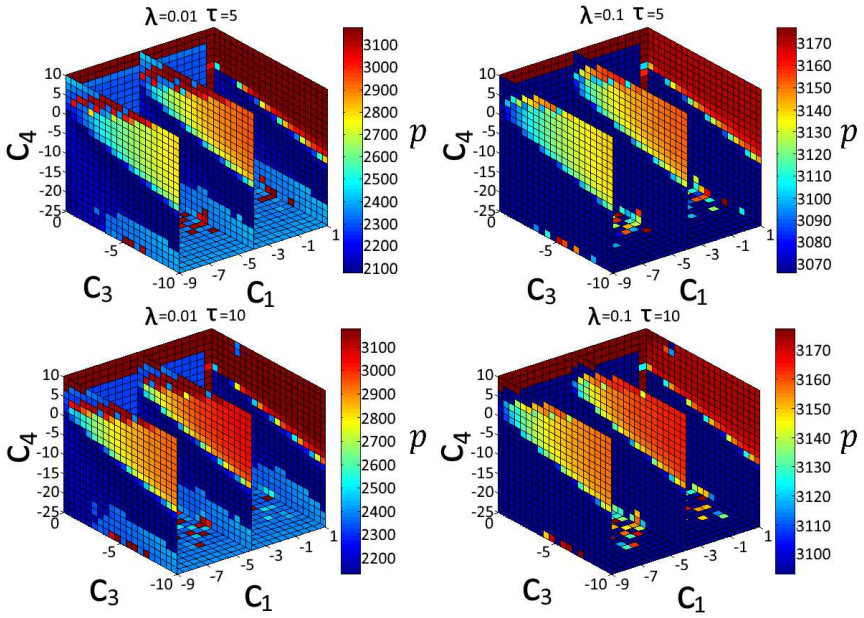

Figure 3.8: Full sensitivity analysis for the non-linear model for varying switching costs and interruption rates. As with the linear model, darker blue areas denote lower penalties and thus a better performance.

throughout the various parametric set-ups. This suggests that the motivations and deficits are changing in a manner very similar to that of the linear model. Indeed, Fig. 3.9 indicates that the dynamics are very similar.

Here, the same result is obtained. With cross-inhibition, the motivations of the agent are irregular and allow for the agent to spend a longer amount of time reducing its deficits earlier on in the trial, with foraging bout durations clearly decreasing in length as time progresses. Without cross-inhibition, the agent again fluctuates across the switching line and so switches between the two nutrients frequently. To build upon this result, a parameter analysis was also carried out to show more clearly how the motivations and deficits changed according to the c_3 and c_4 parameters. This is shown in Fig. 3.10.

As shown before, having $|c_3| > |c_4|$ results in frequent behavioural switching. Instead, setting $|c_3| \leq |c_4|$ results in irregular motivational states. When $|c_3|$ is particularly smaller than $|c_4|$, the motivations are initially very high before being reduced quite quickly. As this result is also similar to that of the

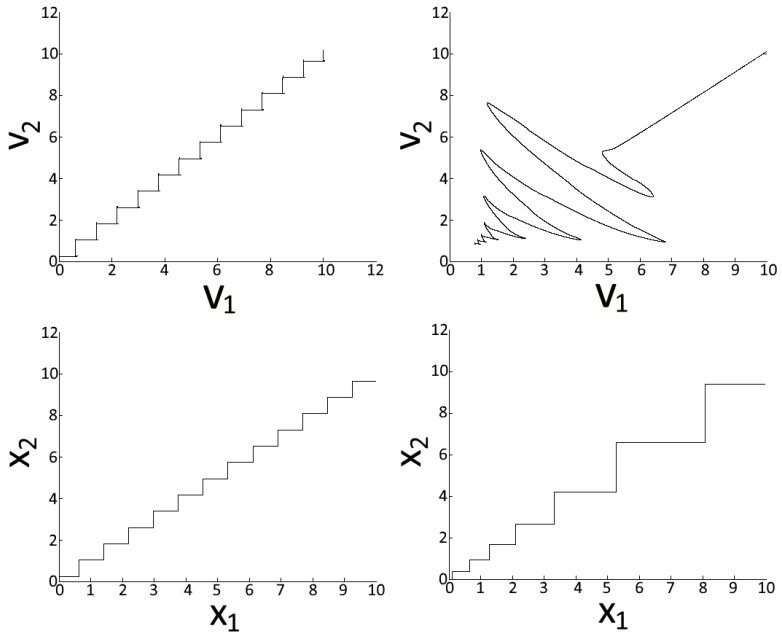

Figure 3.9: The animal's motivations and deficits for two different parametric situations. The top row depicts the motivations, v_i, and the bottom row depicts the deficits, x_i, of the modelled animal with cross-inhibition (right) and without (left). As with the linear model, the introduction of cross-inhibition reduces costly behavioural switching. Again, time begins at the top-right where the deficits are at their maximum level. The parameters used for this figure are the same as those for Fig. 3.3.

linear model, it is implied that the same holds for the model dynamics. These are given in Fig 3.11.

The dynamics are equivalent to those that were found before. When $|c_3| < |c_4|$ the motivations are pushed to one of the two stable points, separated by the unstable point. When $|c_3| > |c_4|$ there is only one stable point where the motivations converge, and at this point the two motivations are very similar to each other.

Finally, the agent's foraging bout durations were analysed. The full results for this are shown in Fig. 3.12.

Here, the agent exhibits foraging behaviours that are similar to that of the linear model. The agent is able to fully reduce its x_2 bout durations in some simulations but not its x_1 deficit bouts. The pattern seen before can be noted again; when $c_4 < 0$ the proportion of bouts being consistently reduced

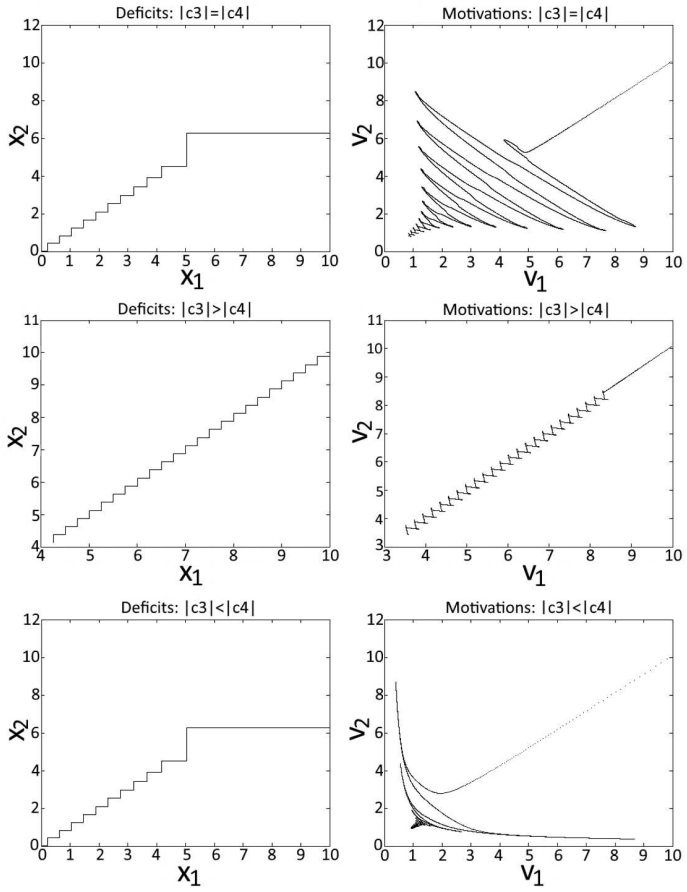

Figure 3.10: The animal's motivations and deficits for different combinations of c_3 (strength of decay) and c_4 (cross- inhibition). Deficits are shown on the left, motivations are shown on the right. As with the linear model, when $|c_3| > |c_4|$ the modelled animal frequently switches. Parameters used for these plots are the same as those used in Fig. 3.4.

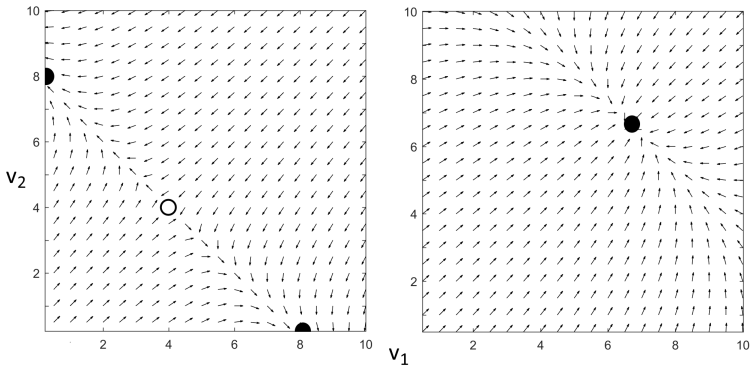

Figure 3.11: As with the linear mode, the changes in the dynamics are dependent on the c_3 (strength of decay) and c_4 (strength of inhibition) parameters. Again, when $|c_3| < |c_4|$, the motivations of the animal end up in a state where one is suppressed whilst the other isn't. Stable points in motivational states are denoted by black filled (semi) circles; unstable points by open circles. The parameters used are as in Fig 3.5.

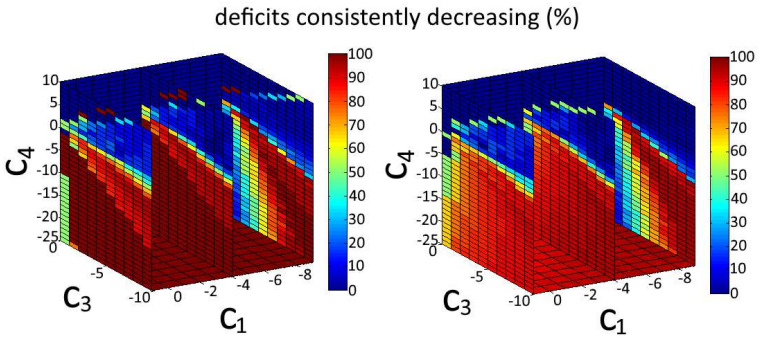

Figure 3.12: The full results for foraging bouts analysis for the non-linear model. The matrices show that proportion of bouts that are being consistently decreased over the simulation time. Again, the best performance that may be obtained is 100%, denoted by the red areas of the plots. The results are given for each deficit individually, on the left is the result of the x_2 deficit, on the right is the result for the x_1 deficit. The results show that the best performances are derived when cross-inhibition is present within the model ($c_4 < 0$). Parameters used are the same as those used in Fig. 3.6.

increases, and as the strength of inhibition grows the number of simulations where the agent is able to reduce 100% of its bouts also increases. When $c_4 > 0$, the agent is unable to consistently reduce its bouts over the course of a simulation, indicated by the blue areas of the matrices. The relationship between cross-inhibition and agent performance is perhaps highlighted to a greater degree in the non-linear extension, as the agent is able to reduce 100% of its bouts in more simulations when $c_4 < 0$. This indicates that the introduction of secondary competition in the inhibitory pathways improves the behavioural performance of the agent, allowing it to reduce its bouts to a further extent (compare results from Fig. 3.6 with Fig. 3.12).

3.6 Discussion

The model presented in this chapter examines an animal's decision-making in a foraging task wherein the animal must choose to consume one nutrient or the other. The model builds upon a previous model of behavioural switching developed by Houston and Sumida (Houston and Sumida 1985) and introduces cross-inhibition as proposed in the leaky competing accumulator model of Usher and McClelland (Usher and McClelland 2001). The results suggest that cross-inhibition improves decision-making, allowing the agent to reduce costly behavioural switching and also reduce its bout length durations over time. Interaction between motivations thus produces more optimal foraging behaviours. Although cross-inhibition has been shown to be beneficial to decision-making models before, it was here shown for the first time in a model of behavioural switching, originally proposed to minimise switching via the positive feedback mechanism. The introduction of cross-inhibition improved the original model by improving the modelled animal's foraging efficiency and reducing frequent switching (Marshall et al. 2015).

There are a few limitations to this model. Firstly, it makes the assumption that each resource will only provide a single type of nutrient, however, in many cases this assumption will not hold (Marshall et al. 2015). As an example,

many food sources will also provide water as well. Secondly, the original model made use of linear dynamics, however, decision-making is not a linear process. As a possible solution to this limitation, a non-linear extension was proposed whereby the interactions between the agent's motivations are not linear. This model made use of the inhibition of lateral inhibition and was able to produce similar results to that of the linear model, with improvements to the agent's trajectories in nutrient space. However, there was also a slight degradation in penalty scores for certain parametric set-ups.

We hypothesise that although the dynamics of the non-linear model were shown to be almost equivalent to those of the linear model, they are such that for particularly large values of $|c_4|$, the agent returns to switching frequently, although it begins the simulation by spending more time at one location before switching to another. As such, the correctness of the behavioural data (where the agent is reducing the duration of its foraging bouts as time progresses) comes with the reintroduction of some costly switching, for certain parameters. Such a result indicates that the non-linear model needs some revision. Alternatively, it could also indicate a trade-off.

Under certain parametric conditions, the interactions of the agent's two competing motivations gave rise to an increase in its penalty score. The dynamics of the interactions were analysed over the course of several simulations. The analysis revealed that, at the beginning of these simulations, one motivation was very large whilst the other had been supressed and was much smaller. Shortly after the initial stage, the motivations decreased to the point where one was only just larger than the other. Finally, the motivations converged to a point where they were almost equivalent and they remained in this state until the end of the simulation, which resulted in the reintroduction of behavioural switching. These dynamics are only present when $|c_4|$ becomes too high. Nonetheless, both linear and non-linear models predict that an animal should exhibit irregular foraging bouts in order to forage optimally.

Can a model of behavioural switching be applied to honeybees? Recently,

the geometric framework has been used to examine honeybee worker nutrition and gives some insight into how honeybees need to balance protein and carbohydrate intake (Paoli et al. 2014). Furthermore, the framework has also been applied to larval honeybee nutrition, and it has been demonstrated that the ratio of protein and carbohydrate levels impact the growth rate and development of honeybees in their larval state (Helm et al. 2017), indicating that there is a nutritional balance that optimises larval growth and minimises mortality. The framework has also been applied to adult worker bumblebees (Stabler et al. 2015) and *Drosophila* larvae (Rodrigues et al. 2015). Furthermore, the geometric framework may also be applied to collective decision-making in honeybee colonies (Bose et al. 2017) as it has been in ant colonies (Dussutour and Simpson 2009). Taken together, all these studies suggest that invertebrates are working to optimise their nutritional intake and that switching from the consumption from one nutrient to another is necessary. In which case, irregular foraging bouts from honeybees is predicted.

To conclude, although this model of behavioural switching may be simplified, it can be applied to many foraging contexts and demonstrates the need for irregular foraging bouts to optimise nutritional intake. Furthermore, the model reiterates the need for coupling in decision-making pathways; a foraging agent without cross-inhibition between its motivations was shown to be unable to reduce costly behavioural switching. The geometric framework has proven very valuable in predicting an animal's decision-making behaviours in light of its motivational state. Indeed, the level of satiation may have a profound impact on an animal's foraging strategies and its ability to choose between two alternatives. This is explored experimentally in the next chapter.

Chapter 4

Is Honeybee Decision-Making Described by a Drift-Diffusion Process?

Fast, robust decision-making in response to sensory input has typically been studied using the two-alternative forced-choice task. Originating from psychophysics, this paradigm involves the presentation of two different stimuli to a subject that must discriminate between them. The similarity of the stimuli is experimentally controlled such that they are quite disparate or almost indistinguishable, corresponding to an easy or hard task, respectively. A classical example of such a task is that of the random-dot paradigm, where a monkey is trained to observe a group of moving dots on a screen and indicate the overall direction of their movement by an eye saccade (Shadlen and Newsome 1996, Kim and Shadlen 1999, Shadlen and Newsome 2001, Roitman and Shadlen 2002, Huk and Shadlen 2005).

Shadlen and Newsome identified the lateral intraparietal (LIP) area of the mammalian brain as a higher-level region which may play a role in the decision process (Shadlen and Newsome 1996). They used the random-dot paradigm to determine whether or not this was the case, recording from neurons in LIP

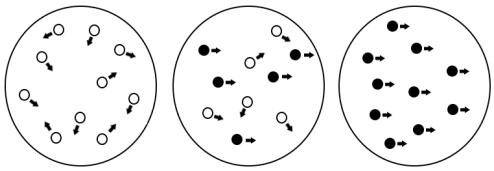

Figure 4.1: The random-dot paradigm. Left: 0% coherence where all the dots move in random directions. Middle: 50% coherence where half the dots move in the same direction (filled black circles). Right: 100% coherence where all the dots move in the same direction.

whilst a monkey made its decision. They found that the responses of the LIP neurons could be used to predict what the monkey would decide, so much so that '*an experimenter could generally predict decisions "on the fly" during an experiment simply by listening to the neuron's activity on the audio monitor*' (Shadlen and Newsome 1996, p. 630). They also found that the predictions became increasingly reliable as time progressed; the neural responses were indicative of a system that integrated sensory evidence over time in order to make a decision. Furthermore, it was later found that the neural accumulation process stopped when it reached a certain threshold (Roitman and Shadlen 2002, Huk and Shadlen 2005). These experimental findings are in agreement with the evidence accumulation theory applied in sequential sampling models: noisy sensory evidence is integrated over time (as such, a higher integration time allows for a greater build-up of evidence and thus a higher degree of accuracy from the animal) until a threshold level is met, at which point a decision is made. These results suggest that for simple perceptual decisions, mammals may be using a drift-diffusion process in order to make a choice. Indeed, the drift-diffusion model (DDM) of decision-making (Ratcliff 1978; 1988, also see Chapter 2 for an overview) has been fitted successfully to experimental data (for example, see Ratcliff and Rouder 1998). Furthermore it can explain the speed-accuracy trade-off (SAT), a behavioural phenomenon whereby faster decision makers are less accurate. Recently, the speed-accuracy trade-off has been described as the benchmark of the decision process (Heitz 2014). Since it is so prevalent within so many species, and even within group

decision- making (Franks et al. 2003, Passino and Seeley 2006, Marshall et al. 2009), it may be inferred that the drift-diffusion process - which is statistically optimal - may capture general features of sensory evidence accumulation. If this is the case, we would expect the phenomenon of the speed-accuracy trade-off to be demonstrated in most decision-making systems, including those of non-mammalian model systems. Indeed, this has already been shown to be the case for honeybees (Burns and Dyer 2008, Wright et al. 2009) as well as bumblebees (Chittka et al. 2003, Dyer and Chittka 2004b, Ings and Chittka 2008, Riveros and Gronenberg 2012) (also see Chittka et al. 2009). As such, we here hypothesise that the drift-diffusion model is able to describe perceptual decision-making behaviours in honeybees. We furthermore hypothesise that changes in motivational state impact the process of evidence accumulation, and are translated as changes in the decision threshold of the drift-diffusion model, as shown in Fig 4.2.

If honeybees approximate a drift-diffusion process in their decision-making, a change in motivational state should thus correspond to a change in accuracy and reaction time, as predicted by the drift-diffusion model. We here test this hypothesis by observing how the choice behaviour of honeybees differs according to two different levels of satiation, using the proboscis extension reflex (PER) paradigm. We first present the PER paradigm in detail before discussing the expected behavioural changes.

4.1 The Proboscis Extension Reflex

The proboscis extension reflex paradigm (Takeda 1961, Bitterman et al. 1983, Felsenberg et al. 2011, Giurfa and Sandoz 2012) makes use of classical Pavlovian conditioning (Pavlov 1927) in order to train honeybees to associate an initially neutral stimulus (known as the conditioned stimulus or 'CS') with positive or negative, biologically relevant, reinforcement (known as the unconditioned stimulus or 'US'). The US will invoke an innate behavioural response and, after training, the animal will exhibit this response to the CS. Touching the

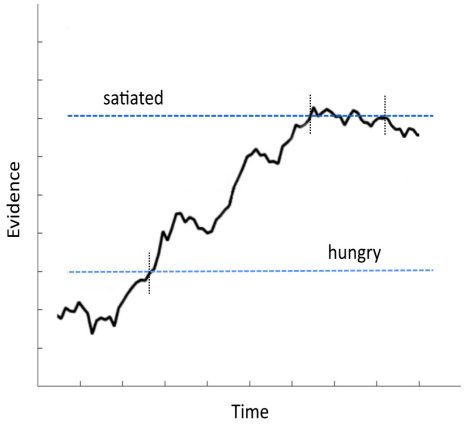

Figure 4.2: Diagram to show a theoretical evidence accumulation process during a decision-making trial. The two dashed horizontal lines depict decision thresholds at two different levels; these levels have been determined by the internal state of the decision maker. Vertical dashed lines indicate when an agent shows a behavioural response. If the evidence accumulated is crossing from below the threshold to above, the decision maker will reach a decision and respond to the stimulus. In our case, a bee will extend its proboscis. If the evidence accumulation is crossing from above to below the threshold, this corresponds to the withdrawal of the proboscis. Here, we can see that a satiated animal should take longer to respond and should also spend less time responding, either in the event of an error or if the stimulus is removed. The lower threshold has been removed since the task given in this experiment is more of a go/no-go task.

antennae of a bee with sucrose solution will elicit the proboscis extension reflex (where the bee extends its proboscis); olfactory conditioning will result in the bee showing this same response to trained odours. Experimental results have shown that honeybees are very efficient learners, with high learning rates after two to five associative trials (Bitterman et al. 1983).

This paradigm has given experimenters the chance to implement a decision-making task that is suitable for invertebrates with a high degree of control over the animal. Although many mammalian choice tasks use visual stimuli, PER conditioning with honeybees requires moving from the visual to the olfactory domain. Previous studies in olfactory PER conditioning have shown that honeybees perform very well and are able to discriminate between odours (Guerrieri et al. 2005). Using visual stimuli with PER is only possible if the antennae of the bees are cut, however, this results in drastically reduced

learning rates (Hori et al. 2006; 2007). This may be attributed to the fact that bees with removed antennae are less responsive to sucrose (de Brito Sanchez et al. 2008). Even after an extensive number of learning trials, many more than what is required for olfactory conditioning, performance is heavily impacted. Furthermore, the ability to discriminate colours, even those which should be easily distinguishable, is diminished (Niggebrügge et al. 2009). Interestingly, a recent study has reported success in honeybee visual conditioning by using an alternative method of restraint (Dobrin and Fahrbach 2012), however, since we here use the classical method of individual tubes and tape to hold the bee in place (see Materials and Methods), it is more desirable to use odours in our conditioning (for a review of visual conditioning protocols, see Avarguès-Weber and Mota 2016).

The PER paradigm has now been standardised (Felsenberg et al. 2011) and it has proven to be popular in honeybee research over the last few decades, especially for the examination of choice behaviour and discrimination ability (e.g., Guerrieri et al. 2005) and also for the analysis of learning and memory. We here use it to ascertain honeybee discrimination ability within different motivational settings, and furthermore compare the behavioural results with predictions from the drift-diffusion model of decision-making. Since the PER paradigm allows for a high degree of experimental control, motivational state (more specifically, the level of satiation) is rendered easy to manipulate. Once restrained, bees can be fed with measured amounts of sucrose solution and then held without food for precise intervals of time.

4.2 The Impact of Motivation

Decision-making is a cognitive process which can be influenced by both external and internal states. An example state is that of hunger, which has been shown to cause observable changes in animal behaviour (for example, Susswein et al. 1978, Weiss et al. 1982, and also see Marshall et al. 2015, a model of behavioural switching with a motivational element). This particular internal

state plays a vital role in PER conditioning studies because the animals are responding to food rewards; a fully satiated animal is unlikely to respond to the conditioning phase of the paradigm as it will have no motivation to extend its proboscis to sucrose. Experimental studies which examine specific behaviours of animals, but do not ensure that the subjects are equivalent in terms of their hunger states, may encounter differences in the behavioural results obtained.

When considering how a bee's satiation level may impact the way it makes choices, it is first necessary to review the components of the hypothesised decision process. Since we view this from a drift-diffusion perspective, we thus find that a change in motivation could alter the following properties:

- the level of sensory noise
- the rate of evidence accumulation (drift)
- the level of the decision threshold

Within a foraging context, a starving bee (or indeed, any animal) will need to act quickly in order to ensure its survival; as such it will make decisions rapidly and perhaps respond to alternatives it might not have if it were satiated, including options which introduce a higher chance of predation (Pettersson and Brönmark 1993, Godin and Crossman 1994, Latty and Beekman 2011). This could be translated as 'risky' behaviour. Within a two-alternative forced-choice task, the bee may respond to unrewarded or punished odours, which would be deemed as being inaccurate. This is intuitive from the drift-diffusion perspective, as faster decision-making results in a higher proportion of errors. Theoretically, the component to control the speed of the decision-maker is the decision threshold. We then hypothesise that, if honeybees are indeed using a drift-diffusion process to make perceptual decisions, then hungry bees should implement a lower decision threshold than satiated bees (corresponding to faster and more inaccurate decision-making). A lower decision threshold determines that the evidence accumulation process will trigger a motor output within a shorter amount of time, thus the discrimination abilities of hungry

bees should be reduced. Within the PER paradigm, we expect to see that hungry bees will respond positively (extend their proboscis) to neutral or negatively reinforced odours and will exhibit a faster response time than satiated bees.

4.3 Materials & Methods

The first PER experiment was conducted between July 2016 and August 2016 at Paul Sabatier University in Toulouse, France. Honeybee foragers, both departing and returning from the hive, were caught from the hive entrance in the morning (using a BioQuip bee vacuum) then chilled on ice until immobilised. Each individual was placed within a small tube and restrained using the classical method of harnessing, such that it could only move its proboscis and antennae. This was implemented using two strips of tape. One strip covered the front and back of the tube so that the body of the bee was secure. Another, thinner strip ran underneath the neck and held the head in place. Groups of twenty to twenty- five individuals were harnessed each day, fed until satiated on 30% sucrose solution and then left to rest for three hours.

Thirty minutes before conditioning, the harnessed bees were tested for their sucrose responsiveness by gently touching their antennae with 50% sucrose solution on the end of a toothpick. Individuals that did not extend their proboscis in response were discarded from the experiment. Each day, the fifteen individuals that showed sucrose responsiveness proceeded to the training phase of the experiment. If there were more than fifteen individuals available, the first fifteen of these were chosen.

The olfactory version of the PER paradigm uses odourants as the conditioned stimuli (CS) and sucrose solution as the unconditioned stimulus (US). Odourant molecules that are within the same chemical group or that have similar carbon chain lengths are generalised more often by bees (Guerrieri et al. 2005). The two odours chosen were Hexanal and 1-Heptanol; these differ in both their carbon chain length and chemical group and should thus be easy to

discriminate (Guerrieri et al. 2005). Hexanal was chosen to be the rewarded odour and 1-Heptanol the unrewarded. During conditioning, presentations of the rewarded odour were paired with the US (50% sucrose solution) and presentations of the unrewarded odour were not reinforced (and as such the odour remained a neutral stimulus). During training, the bees were exposed to each odour five times, resulting in a total of ten trials, which were in a pseudo-random sequence with an inter-trial interval (ITI) of eight minutes (Drezner-Levy et al. 2009). In each trial, individual bees were placed in front of an olfactory stimulus controller (shown in Fig 4.3), which provided a clean airflow, and were allowed to familiarise with the set-up for fifteen seconds. An extractor fan was also positioned behind the bee in order to remove lingering odours. After familiarisation, if the trial was a rewarded trial, the CS was presented alone for four seconds and then presented with the US for a further two seconds. For unrewarded trials, the CS was presented alone for four seconds but no US was present in the following two seconds, as such it was presented alone for a total of six seconds. Finally, the CS was removed and replaced with clean air by the controller, allowing the US to be presented alone for one second (for rewarded trials). Finally, the bee remained exposed to the clean airflow until its trial ended, whereby it was removed from the set-up. In total, a single trial lasted for thirty-two seconds.

After conditioning, an individual was judged to have learnt the associations if they responded correctly in either the last two trials (trials nine and ten) or trials six, seven and eight (see Appendix B for the data sheets that were used for the training). A rewarded odour was presented in trial nine and an unrewarded odour in trial ten, thus if an individual showed a proboscis extension in trial nine and no response in trial ten, they were judged to have learnt. If an individual failed to perform correctly in these two trials, performance in trials six, seven and eight was examined. Trials six and seven were rewarded trials and eight was unrewarded, as such if an individual extended its proboscis in trials six and seven and showed no response in eight, it was judged to

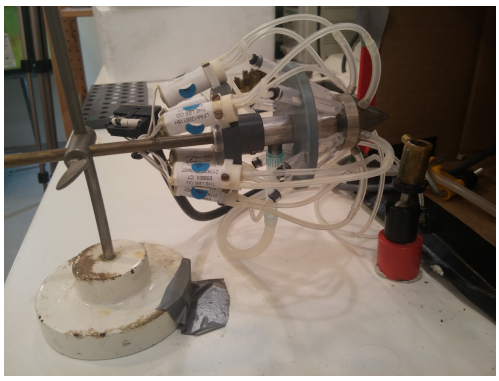

Figure 4.3: Some of the apparatus used for the PER experiment. Shown is the rig which provides a constant, clean airflow during experimental trials; this device is attached to a programmable computer which controls the timing of the odour presentations (not shown). Also shown is the chamber wherein a harnessed bee is held, behind this is an extractor fan which removes lingering odours from the set-up.

have learnt. If an individual did not fulfil either of these two criteria, it was classified as a non-learner and discarded from the experiment. Preliminary experiments showed that some individuals were quicker to learn the associations than others and demonstrated that they had learnt from trial six onwards. Some individuals were slower to learn and demonstrated the correct responses in the final two trials. Some individuals showed they had learnt in trials six to eight but then became unresponsive to the rewarded odour on the ninth trial, however, this did not mean that they had not learnt. It was hypothesised that these individuals had either become demotivated (perhaps due to being gradually satiated throughout the training phase where they were given small amounts of sucrose solution) or perhaps fatigued, and so they were taken through to the testing phase. In this way, the criteria implemented here for learning allows for both slower and faster learners. Spontaneous responders, individuals which responded to the first rewarded trial in the training phase, remained in the experiment. In total, there were three of these (one from the satiated group and two from the hungry group).

After all fifteen bees had undergone the training procedure, those which had demonstrated learning were split into two different motivational groups, hungry and satiated. Both groups were allowed to rest for another three hours

and were then tested on the learnt odours in order to assess their discrimination abilities. The hungry group received no further food whilst the satiated group received a further $5\mu L$ of 30% sucrose solution one hour before the testing phase using an Eppendorf pipette. The bees were placed nearby, but not at, the experimental site. Bees that were fed demonstrated good sucrose responsiveness, however the responsiveness of the hungry bees was not tested. If a bee within the satiated group did not show PER when required to feed prior to the testing phase, it was discarded from the experiment.

During testing, bees were once again placed individually before the airflow and presented with Hexanal and 1-Heptanol over two sequential trials. The bees' responses to the two odours were recorded; bees showing a proboscis extension to the rewarded odour and no response to the unrewarded odour were marked as having perfect accuracy. Bees that failed to respond to the rewarded odour, responded to the unrewarded odour, or responded to the airflow before the odours had been presented, were all marked as incorrect. As such, here an 'error' is defined as an incorrect response in either (or both) of the two test trials. Bees which did not respond in both of the trials were discarded from the experiment. To determine the response time of each of the bees, video footage was recorded during the testing phase and analysed frame by frame. The response time was determined as the length of time between the odour onset (which is determined by a small beep from the olfactory stimulus controller) and the first full proboscis extension exhibited by the bee. The reaction time was measured in the same way for both the rewarded and unrewarded test trials.

Bees responding to the airflow are here marked as incorrect as it is more likely that they are responding to the contextual evidence (that they are within the experimental setup, which itself is not rewarding) rather than the odour presentations. As such, they may be responding randomly. If this is the case, this result aligns with the drift-diffusion model as random responses are more likely when a lower decision threshold has been implemented.

The total number of bees used for this experiment and the analysis of the results is 84 (41 of which were in the satiated group and 43 of which were in the hungry group). A total of fifteen bees were excluded from the analysis of the results due to being completely unresponsive in the testing phase (nine satiated, six hungry), as such the total number of bees successfully trained was 99. Bees faced exclusion from the experiment primarily at five points: post-harnessing (if they did not show PER after being harnessed and could not be fed to satiation), prior to training (if they did not show PER just prior to the training phase), post- training (for being unable to learn the associations), at feeding (if they did not show PER to receive sucrose solution) and post-testing (if they remained completely unresponsive).

For this experiment, we made the following predictions:

- hungry bees will be more inaccurate than satiated bees (or satiated bees will be more accurate than hungry bees).

- hungry bees will have faster reaction times than satiated bees (or satiated bees will have slower reaction times than hungry bees).

These predictions are inferred from the single-bound drift-diffusion model, which predicts that behaviour will be dependent upon the 'positive' threshold (the animal can either respond or not respond). We hypothesise that the level of this threshold can be altered by hunger. More specifically, we hypothesise that a hungry animal will lower their decision threshold. This in turn indicates that they will reach a decision within a shorter amount of time and will consequently become more inaccurate in their decision-making. On the other hand, satiated bees should implement a higher threshold and thus spend more time accumulating evidence before making a decision, as such they will have slower reaction time but will be more accurate.

4.4 Results

Since the odours chosen for the experiment are quite dissimilar, we presented the honeybees with what should have been an easy discrimination task. We first demonstrate that the learning curves for the two groups was similar from the training phase, such that differences in behaviour cannot be attributed to any differences in learning. The acquisition curves are presented in Fig 4.4.

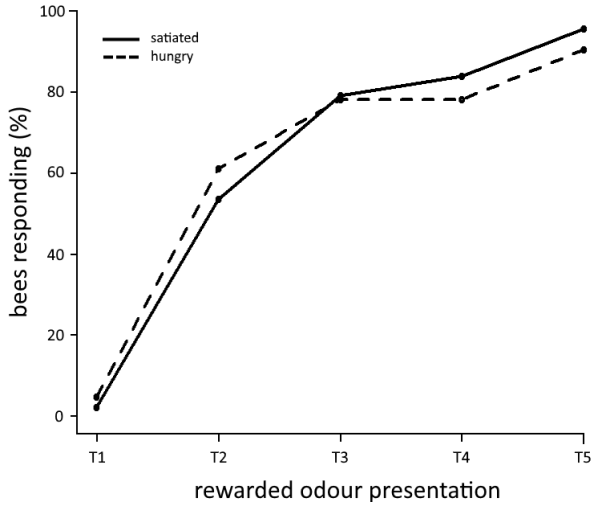

Figure 4.4: The acquisition curves for both groups of bees from the training phase. Trial numbers denote the presentations of the rewarded odour and do not refer to the actual trial numbers in the experiment. All bees were treated the same upon being caught from the hive (i.e., fed until satiated and left to rest) and were given the same training, as such it should be the case that these curves are similar.

During the testing phase, we found that the majority of bees from both groups responded to the rewarded odour presentation (all the hungry bees and 39 of the 43 satiated bees) and that there was no significant difference between the responses. For the unrewarded odour presentation, responses were recorded from 22 of the 41 hungry bees and only 6 of the 43 satiated bees. There was a significant difference between the responses of the two groups (2-sample proportion test, $p < .001$), thus we can conclude that hungry bees were significantly more responsive to the unrewarded odour than the satiated bees. These data are shown in Fig 4.5.

The full data for the testing phase are shown in Fig 4.6. A total of 9

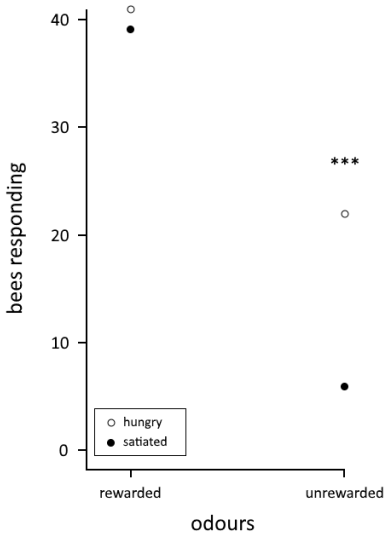

Figure 4.5: Graph to show the percentage of bees from each motivational group that responded to each of the test odours.

satiated bees and 6 hungry bees were discarded for not responding in any of the test trials and are thus not included in the results. We found that most satiated bees were able to discriminate between the two odours correctly, with around 73% of bees completing the choice task without error. However, the hungry bees were not able to discriminate as well, and often mistakenly extended their proboscis to the clean airflow before the odour was presented. Only 30% of bees in the hungry group were able to discriminate without making a mistake (that is to say, the bees both responded to the rewarded odour and did not respond to the unrewarded odour). We found that the difference in accuracies of the two groups was statistically significant (2-sample proportion test, $p < .001$), as such we can conclude that hungry bees are more error prone than satiated bees. For completeness, we also analysed the accuracies of the groups when the bees that responded to the airflow were not marked as incorrect. In this case, only the bees that responded to the punished odour were marked as making an error. This analysis is shown in Fig 4.7. We further analysed the response times of the two groups to determine whether or not hungry bees were responding to the presented odours

quicker than satiated bees. Analysis showed that, for the unrewarded odour presentation, hungry bees were, on average, significantly quicker in extending their proboscis (Wilcoxon-Mann-Whitney test, $p < .05$). Although we did not find significance for the rewarded odour, the average response time for hungry bees was still less than that of satiated bees. The result for the unrewarded trial, taken together with the results for the accuracies, is indicative of a speed-accuracy trade-off: the hungry bees were quicker in making their decisions for the unrewarded odour presentation and suffered a reduction in their accuracy. This result suggests that the bees were not making mistakes due to some other factor; if bees are indeed approximating a drift-diffusion process in their decision-making, a reduction in decision accuracy must be accompanied by a reduction in reactionctime, and these behavioural changes should arise from a reduction in the level of the decision threshold. The results here, in terms of the decision accuracies and the reaction times for the rewarded odour, support the theory that hungry bees are limiting their evidence accumulation by means of a reduced decision threshold.

4.5 Discussion

We used the proboscis extension reflex paradigm to test the discrimination abilities of honeybees in two different motivational states, keeping the training protocol for the groups equivalent and instead focusing on how internal state can impact decision-making processes. Our experiment showed that hungry honeybees are significantly more prone to error than satiated bees. We also found that hungry bees were more likely to extend their proboscis to the airflow upon being placed in front of the PER rig. This implies that hungry bees may be completing their evidence accumulation even before the odour onset, and consequently reaching the decision threshold before they could identify which odour was presented. This also indicates that hungry bees were making quicker decisions. Another interpretation of this result is that the hungry bees were responding randomly, however, randomness is also indicative of a

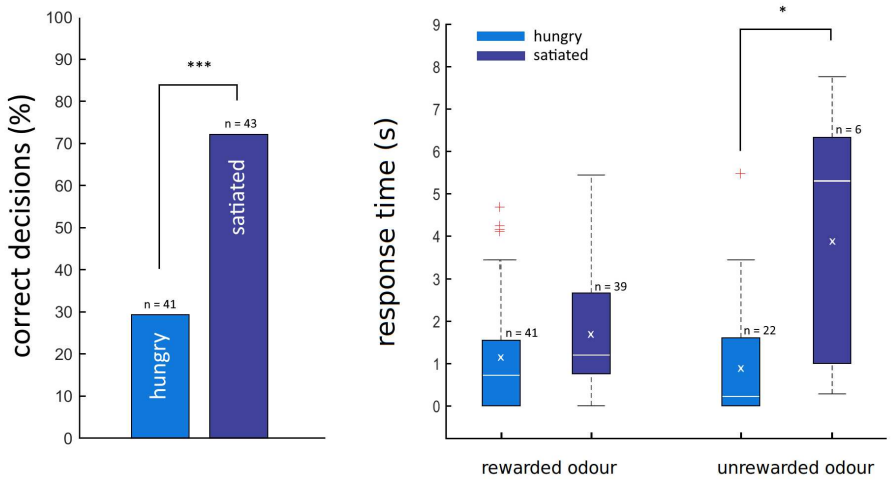

Figure 4.6: Honeybee decision-making analysis. Left: percentage of bees from each group that were able to discriminate without making an error. Right: response times for both the rewarded and unrewarded odours. Medians are denoted by white lines, averages are denoted by white crosses, and outliers are denoted by red crosses. Asterisks denote significance values. *: $p < .05$, **: $p < .01$, ***: $p < .001$. Statistical test used for decision accuracies: 2-sample proportion test. Statistical test used for reaction times: Wilcoxon-Mann-Whitney test.

lower decision threshold (fluctuations below and above the threshold due to noise causing extensions and retractions, which is more likely to happen when the threshold is at a lower level, see Fig 4.2). As such, either behavioural explanation (that hungry bees were more random or that they had reached a decision before the odour onset) can align with inferences from the drift-diffusion model.

In addition, we compared the response times of hungry bees to satiated bees after the odour had been presented. We found that hungry bees took less time to extend their proboscis than satiated bees (had a quicker response time). For the unrewarded odour, this result was significant. This behavioural result supports the hypothesis that honeybees may be implementing a drift-diffusion process for simple perceptual decisions, where the level of satiation mediates the decision threshold and consequently the level of evidence that is accumulated before a decision is reached. The results suggest that a hungrier animal will have a lower decision threshold allowing it to respond to

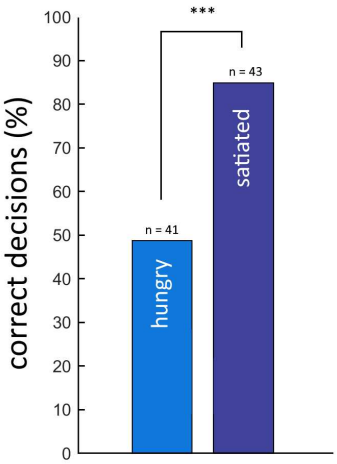

Figure 4.7: Honeybee decision-making analysis with the definition of error adjusted to no longer count bees responding to the airflow as incorrect. Graph shows the percentage of bees from each group that were able to discriminate without error. Asterisks denote significance values; ***: $p < .001$. Statistical test used: 2-sample proportion test.

stimuli more quickly, even if a speed-up in response time will result in more errors. This result, taken together with the decision accuracies, demonstrates a speed-accuracy trade-off (hungry bees made more mistakes and were faster in responding to odours). It should be noted, however, that there was no significant difference for the rewarded odour presentation. There may be several reasons for this, including an insufficient sample size or an insufficient difference in the hunger levels of the two groups. These limitations provide options for future work; this experiment could be replicated but with more than 40 bees in each group, and an additional hour (or more) could be added to the resting time for the bees so that hunger is more pronounced. It should be noted that bees in the hungry group should not be starved to the point wherein their health is impacted as this would undoubtedly change the behavioural results.

We have here used the drift-diffusion model to make inferences about the behaviours of the bees in different internal states. For one of the presented odours, the satiated bees are significantly slower in responding than hungry

bees and this was explained by the mediation of the decision threshold. Another explanation is that satiated bees are slower due to their using energy for digestion. This is unlikely to be the case, since they were given a small amount of sucrose solution when fed before the testing phase and were also given one additional hour to rest after this feed (as such, they had already had one hour for digestion). Additionally, although we have here attributed a change within internal state to a change in the level of the decision threshold, it is possible that the rate of acccumulation (drift) is being altered instead. To explain the slower response times, satiated bees would need to be accumulating evidence at a slower rate than hungry bees, as such it would take them longer to reach a static threshold. At this stage, it is not possible to ascertain whether the behavioural changes observed are due to a change in the decision threshold level or the drift. In order to clarify, the mechanism that implements the threshold within the brain of the bee would need to first be identified and then neural recordings taken.

Another explanation of the behavioural differences found could be that the hungry bees became disengaged from the task, which in turn made them more prone to error or more prone to showing PER to stimuli other than the rewarded odour. As mentioned above, if it was the case that hungry bees became disengaged and more 'random' than satiated bees, this behavioural difference still fits the predictions of the drift-diffusion model and the inference that hunger should lower the decision threshold. However, this is unlikely, as all bees were tested for sucrose responsiveness before the testing phase and were seen to be motivated and in good health, as such it is most likely not the case that they were too hungry. However, since this experiment was a go/no-go task, it is impossible to distinguish between a bee that has chosen not to respond from a bee which did not reach a decision. As such, it would be beneficial to translate this experiment into a true two-alternative forced-choice task, such that an error can be properly differentiated from disengagement or a lack of choice. This is yet another opportunity for future work.

4.6 Experiment Two: Materials & Methods

We theorised that the large number of bees extending their proboscis to the airflow in the first experiment may have been caused by the training protocol. During conditioning, 1-Heptanol was an unrewarded CS, as such there was no negative consequence for bees that extended their proboscis to it. Furthermore, it has been shown that aversive conditioning can improve discrimination in free- flying bees (Avarguès-Weber et al. 2010), as such, it may have been the case that conditioning with an aversive stimulus was impacting the discrimination abilities of the bees. We therefore decided to replicate the experiment and alter the training protocol such that 1-Heptanol was instead punished.

The second behavioural experiment was conducted between January 2017 and February 2017 at Macquarie University, Sydney, Australia. It mirrors the first experiment closely and aims to examine the same hypothesis, however, we implemented a few changes that were hypothesised to make the results more concrete. Firstly, we altered the conditioning protocol such that 1-Heptanol was punished as opposed to unrewarded. During training, presentations of 1-Heptanol were paired with saturated salt solution, which was delivered to the bee in the same way as the sucrose solution (on the end of a toothpick, first touching the bees' antennae and then the proboscis). This change was hypothesised to stop the bees from responding to the clean airflow prior to the odour onset.

Additionally, we incorporated a change to the testing protocol. As well as presenting the two odours that were used during conditioning, we made three different compound odours which were composed of ratios of the trained odours (referred to in terms of their ratios; Hexanal:1-Heptanol). As such, the compound mixtures would present the trained odours simultaneously but also introduce a novel component. How exactly bees process compound odours has been debated. Two main theories have been proposed, 'elemental' processing (whereby a compound is regarded as the sum of its parts, $X_{AB} = X_A + X_B$) and 'configural' processing (whereby a compound is an entirely novel stimu-

lus) (Rescorla 1973, Giurfa et al. 2003, Deisig et al. 2003). It has been found that bees use a combination of both (Meyer and Galizia 2012); as such a compound is processed as the sum of its parts, but with the addition of a unique component ($X_{AB} = X_A + X_B + X_C$) (Deisig et al. 2003). Thus, compound mixtures are excellent for controlling precisely how similar or dissimilar an odour should be to the rewarded stimulus; they present both learned and novel features, and they furthermore negate the need to use alternative odours (that differ in chemical group and carbon chain length). Whilst it is expected that the presence of a greater amount of CS+ within a compound odour will generate more responses from bees (as it is chemically more similar to the rewarded odour), the presence of the CS- (and the novel component) indicate that the odour is still not equivalent with the rewarded odour and that a response is thus a mistake. We used five test odours in total:

- the rewarded odour, Hexanal (100:0)
- a compound composed mostly of the rewarded odour (70:30)
- a compound composed with equal amounts of the trained odours (50:50)
- a compound composed mostly of the punished odour (30:70)
- the punished odour, 1-Heptanol (0:100)

The introduction of the compound odours meant that the bees would experience multiple presentations of the trained stimuli without reward. As such, the bees may have, at some point during the testing phase, undergone a type of learning whereby the reinforced CS become neutral again due to the retraction of reinforcement. This type of learning is known as 'extinction'. Thus, the rewarded odour is no longer perceived as rewarded, and the punished odour as no longer punished. Previous experiments have shown that total extinction can occur within as little as five trials and can begin on the second trial (Bitterman et al. 1983), following a conditioning protocol composed of five positive trials. As such, after the first odour presentation during

the testing phase, bees may become less responsive. Additionally, the possibility of extinction makes the order in which the test odours are presented important. For this experiment, the five test odours were arranged into the following pseudo-random sequence: (1) 0:100, (2) 70:30, (3) 50:50, (4) 100:0, (5) 30:70. The odours were presented in this order to each bee, such that each individual encountered the odours in the same sequence. This was to ensure that any observed differences in behaviour were not due to encountering the odours at a different time.

As with the first experiment, the bees' responses to the five odours were recorded. Bees that showed a proboscis extension to the rewarded stimulus (Hexanal, or 100:0) only were marked as discriminating without error. Bees that failed to respond to the rewarded stimulus, responded to any of the compound mixtures of the punished odour, or responded to the airflow before the odour onset, were marked as incorrect. Since there were more chances to make an error and involved the presentation of compound mixtures that would be very similar to the rewarded stimulus (i.e., 70:30), this choice task was harder than the one presented in the first experiment. However, this experiment aimed to examine the choice behaviour of the bees more closely and determine if hungry bees would mistake the compound mixtures as the rewarded odour more than satiated bees.

Video footage of the bees was again recorded during the testing phase. Frame-by- frame analysis was carried out to determine the bees' response time (calculated as before), as well as several other additional response characteristics that were thought to differ between the two groups. We also examined the number of proboscis extensions, or bouts, that the bees exhibited during their trials. For example, a full extension followed by a retraction and then another extension would count as two bouts. We furthermore observed that some bees differed in which direction they extended their proboscis; many bees would extend level to the plane of their head, others would try to extend above the plane. We thus recorded if bouts were at or above the plane. Finally, we

analysed how much time each bee spent with an extended proboscis, and calculated how much of this time was spent before the odour onset or after the odour onset.

The total number of bees used for this experiment and the analysis of the results is 81 (41 of which were in the satiated group and 40 of which were in the hungry group). A total of 23 bees were excluded from the analysis of the results due to being completely unresponsive in the testing phase (18 satiated, 5 hungry), as such the total number of bees successfully trained was 104. As the methodology for the second experiment remains the same (asides from the changes to the test odours and that the unrewarded odour is now punished) the bees can be excluded from the experiment within the same points as listed in Materials & Methods. A total of 15 bees were noted to be spontaneous responders (8 satiated and 7 hungry) and these were kept in the experiment.

Our final predictions for this experiment are as follows:

- from the first experiment, hungry bees will be more inaccurate than satiated bees (or satiate bees will be more accurate than hungry bees).
- from the first experiment, hungry bees will have faster reaction times than satiated bees (or satiated bees will have slower reaction times than hungry bees).
- the closer a compound odour is to the rewarded odour (chemically), the more bees will respond to it. As such, the more similar compound odours should induce more errors in both the hungry and satiated bees.
- hungry bees will spend more time extending their proboscis to an odour than satiated bees (or satiated bees will retract their proboscis more quickly than hungry bees).

The first two predictions are inferred from the drift-diffusion model, as with the first experiment, and are described in the Materials & Methods of that experiment. Additionally, we hypothesise that compound odours that

incorporate more of the rewarded odour than the punished odour are more likely to be responded to than other compound odours (i.e., the 70:30 compound will see more bees responding to it than the 50:50 compound). This is due to the fact that odours which are chemically more similar to a trained, rewarded odour will induce bees to generalise more, as seen in the work of Guerrier and colleagues (Guerrieri et al. 2005). Here, an error is defined as a response to any odour other than the rewarded odour (as such, PER to any of the compound odours or the punished odour will be recorded as a mistake). This is due to the way bees are hypothesised to process compound odours (Meyer and Galizia 2012), however, an alternative theory (which concludes that PER to compounds should not be marked as mistakes) is presented in the Discussion. The fourth prediction is again inferred from the single-bound drift-diffusion model. Since hungry bees are implementing a lower decision threshold, the accumulated evidence should remain above the threshold for a longer amount of time, as such the hungry bees should prolong their responses to the odour presentations.

4.7 Experiment Two: Results

We presented honeybees with a discrimination task that was slightly more difficult than that of the first experiment, however, it should still have been easy to solve. As before, we first show that the acquisition curves for the two groups are similar. This is given in Fig 4.8.

The proportion of bees responding to each of the test odours is shown in Fig 4.9. We found that hungry bees were consistently more responsive to the odours than satiated bees, with a higher proportion of hungry bees extending their proboscis in each case. For three of the odours, the differences in responsiveness between the two groups was significant; analysis showed that significantly more hungry bees extended to the punished odour (2-sample proportion test, $p < .001$), the 70:30 compound mixture (2-sample proportion test, $p < .05$) and the 50:50 compound mixture (2-sample proportion test, $p <$

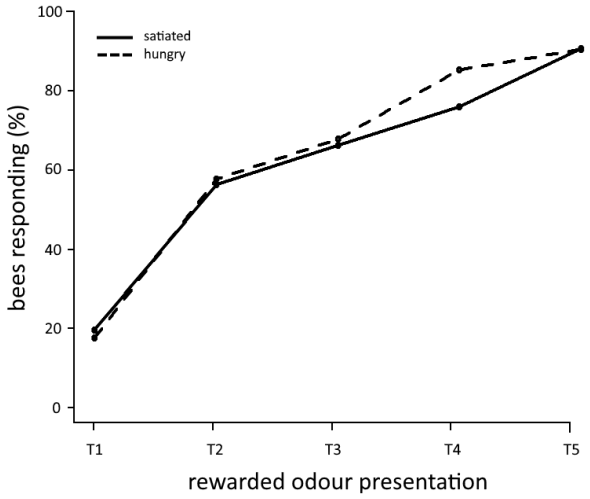

Figure 4.8: The acquisition curves for the two groups of bees from the training phase. As in Fig 4.4, the trial numbers here denote the presentations of the rewarded odour and do not refer to the actual trial numbers in the experiment. As in the first experiment, all bees were given equivalent training and were treated the same upon being taken form the hive.

.05). Bees responded similarly to the rewarded odour and the 30:70 compound mixture. The majority of bees from both of the motivational groups responded to the rewarded odour, and the compound mixtures that were composed of a higher ratio of the rewarded odour invoked a higher proportion of responses.

The compound mixtures were more similar to the rewarded odour than the punished odour, as such these should have been more likely to invoke incorrect responses from the bees, depending on the ratio of the odour's components. Indeed, analysis of Fig 4.9 indicates that more bees responded incorrectly to the 70:30 mixture than the 50:50 mixture, and more to 50:50 than 30:70. The number of responses for the punished odour is quite high, especially for the hungry bees. A theory as to why is presented in the Discussion. Since the compound mixtures were indeed similar to the rewarded odour, but not entirely equivalent, bees that responded to them were noted for making an error. We calculated the percentage of bees from each group that had responded to the rewarded odour only (i.e., obtained perfect accuracy) and then compared this result to that of the first experiment. The accuracy comparison is shown

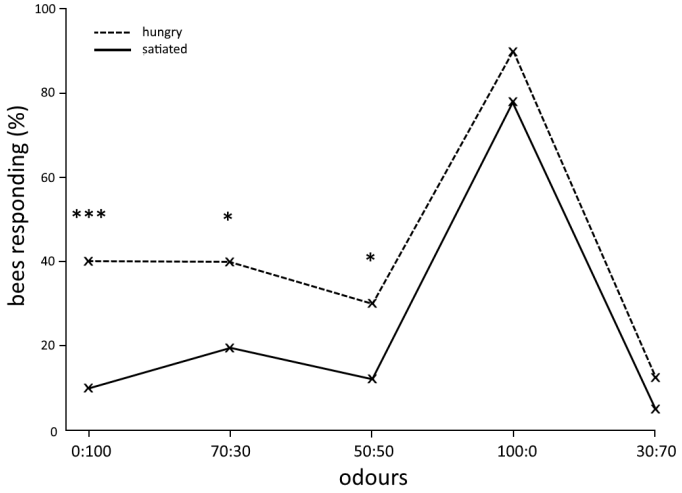

Figure 4.9: Line graph to show what percentage of bees from each motivational group responded to the test odour presentations. Satiated $n = 41$, hungry $n = 40$.

in Fig 4.10.

The accuracies obtained from the two motivational groups are very similar for both experiments. In the second experiment, the satiated bees are slightly less accurate, however, this is not significant. This result confirms that we were able to replicate the first experiment even with slightly different training and testing protocols. In both experiments we found that satiated bees are significantly more accurate than hungry bees (2-sample proportion test, $p < .001$). This can also be seen in Fig 4.9, as less satiated bees respond to the compound mixtures and punished odour. After confirming that the decision accuracies were comparable, we then compared the response times. Since the first experiment only used the rewarded and punished odours in the testing phase, we here only include response times for these two odours from the second experiment. These data are shown in Fig 4.11.

As with the decision accuracy data, we found that both experiments resulted in similar response times for the bees. Hungry bees were again significantly faster in responding to the rewarded odour presentation than satiated bees (2-sample t-test, $p < .05$). Hungry bees were also faster in responding to the punished odour, however, this result was not significant. Nonetheless, of the

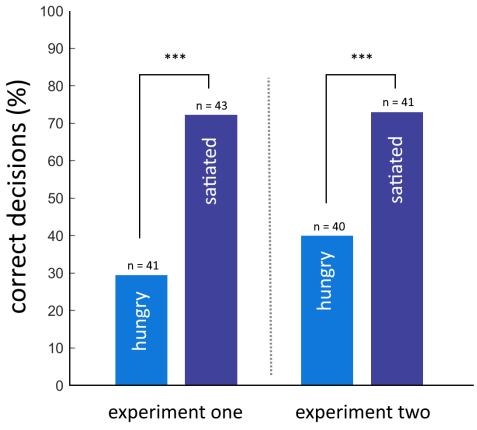

Figure 4.10: Comparison of average honeybee decision accuracy from the first (left) and second (right) experiments. A correct decision corresponds to a proboscis extension to the rewarded odour only. For the second experiment, the given accuracy refers to those bees which responded to the rewarded odour but not the punished odour. In both experiments, there was a significant difference in the accuracies of the two groups (2-sample proportion test, $p < .001$ in both cases.

two reaction time comparisons that were made in each experiment, one in each was found to be significant: hungry bees were faster than satiated bees in responding. These significant results taken with the accuracy data (which show that hungry bees are more error prone) are suggestive of a speed-accuracy trade-off and fit the predictions of the drift-diffusion model of decision-making. We additionally analysed the response times for the novel compound mixtures; the full data are presented in Fig 4.12. Of the 41 satiated bees, 4, 32, 8, 5 and 2 of these responded to the punished, rewarded, 70:30, 50:50 and 30:70 compound odours, respectively. Of the 40 hungry bees, 16, 36, 16, 12 and 5 of these responded to the punished, rewarded, 70:30, 50:50 and 30:70 compound odours, respectively.

We found that hungry honeybees were, on average, faster decision makers for every test odour except for the 30:70 compound. In this case, hungry bees were almost equivalent. We found that one test odour shows a significant result. However, the data show that hungry bees are indeed consistently responding quicker than satiated bees, for both the trained odours and novel compounds. This is not the only behaviour that was consistently different in

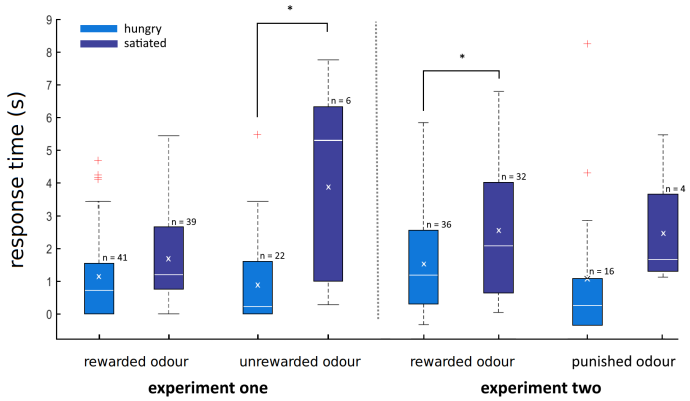

Figure 4.11: Comparison of average honeybee response time from the first (left) and second (right) experiments. See Appendix A for statistical tests used.

hungry bees, however. We performed a more in-depth analysis of the honeybees proboscis extension behaviours to determine whether or not hunger would cause any other behavioural modifications. From here on we thus separate the data in terms of the test odours and present the results for each one. Firstly, we present the full data for the rewarded odour in Fig 4.13.

We first analysed how many times honeybees extended their proboscis to the odour and in which direction. This analysis highlighted no differences in the behavioural response of the bees; both groups extended, on average, an equivalent number of times, with hungry bees extending above the plane of the head slightly more than satiated bees. However, this result is non-significant. As discussed before, hungry bees were responding quicker than satiated bees. We then looked at how much time the bees were spending responding to the odour. We found that hungry bees were spending slightly more time in total extending their proboscis to the odour, however, the result was non-significant. We found that, despite the presence of the punished odour during the conditioning, some bees were still extending their proboscis to the airflow. Bees from the hungry group spent significantly more time extending before the punished odour onset than satiated bees (2-sample t- test, $p < .05$). This may suggest that, for some bees, the decision threshold is sufficiently low enough that evidence accumulation stops before the stimulus presentation.

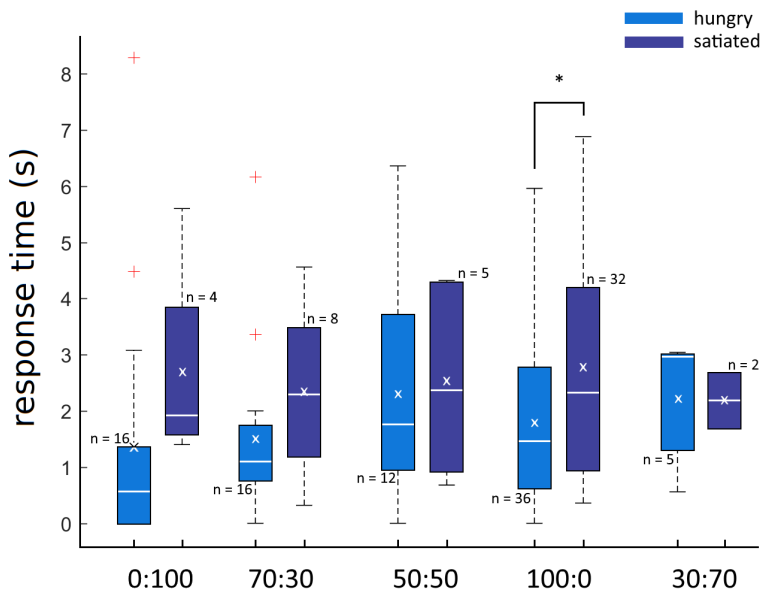

Figure 4.12: Response time data for all test odours. Medians denoted by white lines, averages denoted by white crosses, outliers denoted by red crosses. Odours are given in the order they were presented in during the experiment. See Appendix A for statistical tests used.

Finally, we found no difference in the time bees took extending their proboscis after the odour onset. This result is perhaps expected as both groups will recognise this as the rewarded odour. Following this, we analysed the data for the compound mixtures. The results are shown in are shown in Fig 4.14.

The vast majority of honeybees did not respond to the 30:70 compound mixture (bottom plot in Fig 4.14, hungry $n = 5$, satiated $n = 2$). We found that, for the bees that did respond, hungry bees showed more proboscis extensions both at the plane of the head and above. No satiated bees displayed any extensions above. Both motivational groups were roughly the same in their response times, however, there is a very notable difference in how much time they spent with their proboscis extended, despite the lack of significance. No bees extended before the odour presentation, however, hungry bees spent far more time than satiated bees responding to the presented odour.

A small number of bees responded to the 50:50 compound mixture (middle plot in Fig 4.14, hungry $n = 12$, satiated $n = 5$). We found that, as with the

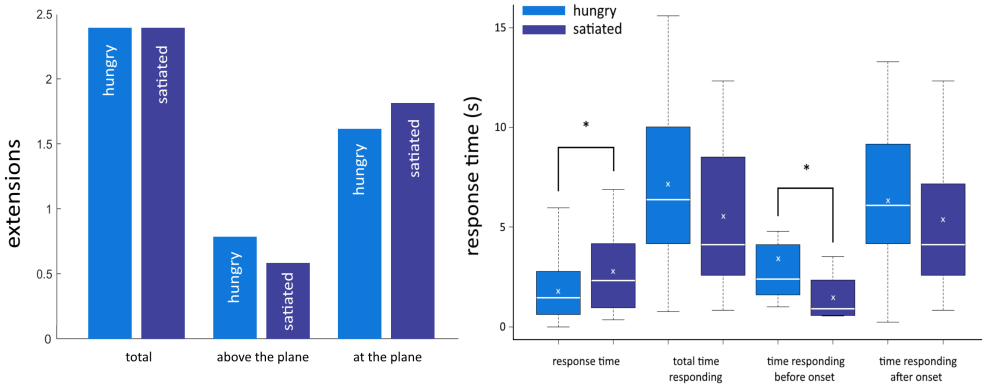

Figure 4.13: Complete data analysis for both motivational groups responding to the rewarded odour. Left: average number of proboscis extensions, in total and separated by the direction. Right: average response time, average total time that honeybees extended their proboscis for during the trial, average time honeybees extended their proboscis *before* the odour was presented, and average time honeybees extended their proboscis *after* the odour was presented.

30:70 compound, there were no significant results. Bees from both groups showed a similar number of proboscis extensions, in total and in both directions. Hungry bees had a slightly quicker response time, but the result is non-significant. For this odour, bees spent more time responding before onset. Interestingly, satiated bees spent slightly more time responding than hungry bees. After the odour onset, hungry bees were again spending more time displaying a proboscis extension, although this result isn't as strong as that of the 30:70 mixture.

The 70:30 compound mixture generated the most responses out of all three mixtures, an expected result as this odour was the most similar to the rewarded odour (top plot in Fig 4.14, hungry $n = 16$, satiated $n = 8$). We found several significant results here. Firstly, hungry bees showed significantly more proboscis extensions than satiated bees (2-sample t-test, $p < .05$). Strikingly, hungry bees also showed more proboscis extensions above the plane of the head than satiated bees (2-sample t-test, $p < .05$). Hungry bees were slightly faster in their decision-making than satiated bees, and spent more time extending their proboscis after the odour had been presented. Hungry bees spent significantly more time extending before the odour onset (Wilcoxon-Mann -Whitney

test, $p < .05$, no satiated bees showed this behaviour for this compound) and also more time extending in total (Wilcoxon-Mann-Whitney test, $p < .05$). This result indicates that satiated bees that made the error in responding to this odour withdrew their proboscis significantly quicker than hungry bees. Finally, the full data for the punished odour are given in Fig 4.15.

Significantly more hungry bees than satiated bees responded to the punished odour (2-sample proportion test, $p < .01$). We found that hungry bees displayed more proboscis extensions than satiated bees, both above the plane of the head and at the plane, however, nothing of significance was found. Hungry bees responded faster to this odour than satiated bees, but again this was not significant. However, hungry bees were spending significantly more time extending their proboscis, both before odour onset, (2-sample t-test, $p < .05$), after odour onset (Wilcoxon-Mann-Whitney test, $p < .05$), and in total (Wilcoxon-Mann -Whitney test, $p < .05$)

4.8 Experiment Two: Discussion

Honeybees are required to solve discrimination problems every day of their lives, and many of these will be more difficult than others. They will no doubt encounter these tasks within different internal states, and how their motivation impacts their decision-making is understudied. We here used the proboscis extension reflex paradigm in order to carry out differential conditioning with honeybees and examine their discrimination behaviours in different states. As in the first experiment, we separated the trained bees into two different motivational groups, hungry and satiated, and examined how their internal state impacted their decision performance. We observed honeybee decision-making behaviours and analysed them in a high amount of detail, going further than just obtaining data for accuracy and response time. The results of the experiment highlighted several behavioural differences between hungry and satiated bees.

We noted from our first experiment that the testing phase was more akin

to a go/no-go task as opposed to a two-alternative forced-choice task. The ideal experimental set-up for the latter would involve the presentation of two odours simultaneously, perhaps from either side of the bees' head, such that each odour contacts one individual antenna. Additionally, the set-up would need to allow the bee to move its head from one side to another. The equipment we used (the olfactory stimulus controller and individual tubes) were not compatible with such an experiment. Instead, we tried to incorporate some changes into the testing phase in our second experiment, and included odours that were composed of components of both the trained odours. In this way, although still not a true two-alternative forced-choice task, the two odours used for training would be presented simultaneously.

We were able to replicate some of the results from the first experiment despite altering the training protocol and introducing novel compound odours to the testing phase. We again found that satiated bees are significantly more accurate than hungry bees (even when the responses to the compound odours are not taken into account and thus not marked as errors) and are also, on average, slower decision makers. However, the difference in response times was only significant for the rewarded odour presentation. This result, taken together with the difference in accuracies, is indicative of a speed-accuracy trade-off, however, four of the five test odours showed no statistical significance for response times. The reason for this may be due to the same problems encountered in the first experiment: insufficient sample sizes or the hungry group not being sufficiently hungrier than the satiated group. We can conclude that one result from the response time data fits the inferences of the drift-diffusion model of decision-making and is indicative of the SAT.

Hungry bees were, on average, faster decision-makers than satiated bees in each of the test odour presentations (asides from the 30:70 compound, however, there were only 2 satiated bees that responded to this odour). This observation is reassuring and we hypothesise that a larger sample size together with a possible change to how hungry the bees are allowed to become (i.e.,

resting them for four or five hours instead of three) will emphasise what has been observed here and will further produce more significant results. For a particularly hungry animal, acting quickly will be of great importance in order to avoid mortality. Responding quickly and making the wrong choice will perhaps be less costly than not acting at all. It is thus intuitive that a simple, efficient mechanism can control decision-making behaviours and alter them according to motivational state. A simple change in the decision threshold can cause a 'cautious' animal to become more 'risky'.

In the first experiment, we noted that hungry bees were more likely to show a response prior to the odour onset. Here, we examined how much time the bees were spending with their proboscis extended both before and after the odours had been presented. We found that, for several odours (rewarded, 70:30 compound, punished), hungry bees spent significantly more time responding to the airflow before odour onset. This result gives some insight into the evidence accumulation process. For some bees, this process begins upon being placed in front of the olfactory stimulus controller, and hungry bees were more likely to come to a decision before the odour was even presented, even though an extension may have resulted in punishment, depending on the odour. This behaviour, which perhaps could be seen as 'risky' as it is a response to an unknown, possibly negative stimulus, also fits the theory that hungry bees are implementing a lower decision threshold. Rather than waiting for the odour onset, the bee begins to accumulate contextual evidence as soon as it is placed in front of the airflow, and the implementation of a lower decision threshold results in extensions of the proboscis. For satiated bees, which should be implementing a higher decision threshold, a higher amount of evidence is required to trigger a response. Therefore, even if satiated bees begin evidence accumulation as soon as they are placed within the airflow, they are more likely to refrain from responding than hungry bees.

We also found behavioural differences in the amount of time bees spent extending their proboscis after the odour onset. For all odours except the

rewarded odour, hungry bees spent much more time extending their proboscis than satiated bees. For the punished odour, this difference was significant. This result suggests that evidence accumulation continues after an initial decision has been made, which gives the animal a chance to change its mind or correct an error. For satiated bees with a higher decision threshold, the evidence accumulated has a greater chance to fall below it, and is more likely to do so if, for example, the bee has identified that the stimulus presented is not the rewarding stimulus. Behaviourally, this would correspond to satiated bees withdrawing their proboscis after a shorter amount of time. On the other hand, with lower decision thresholds, hungry bees are far more likely to continue responding once they have reached their decision, as the evidence accumulation will take more time before it drops back below the decision threshold. Indeed, fluctuations (due to noise) or a reduction in the presented evidence will remain above a lower decision threshold for longer. This is represented visually in Fig 4.2.

The significant results found (hungry bees more inaccurate, hungry bees faster in responding to the rewarded odour, hungry bees spent more time responding to the 70:30 compound, hungry bees spent more time responding before the odour onset for the 70:30 compound, hungry bees showing a greater number of bouts to the 70:30 compound, hungry bees spending more time responding to the punished odour and to the airflow before the odour onset) are all compatible with the drift-diffusion model of decision-making. Whilst there are many results which are non-significant, none have been found that conflict with the inferences of this model and the hypothesis that hungry bees are implenting a lower decision threshold. We also discovered significance in a result not predicted: for the 70:30 compound odour, hungry bees were making a significantly greater number of proboscis extensions than satiated bees, and also displayed significantly more 'above the plane' extensions than satiated bees. Thus, internal state has shown to cause a difference not only in the decision accuacy and response time but also the total time responding

(before and after odour onset), the total number of proboscis extensions and even the direction in which bees will extend. Here, it is proposed that bees are extending upwards as, during conditioning, the toothpick with rewarding sucrose solution would come from above their head, touch their antennae and finally touch their proboscis. As such, the bees are extending in the direction the reward should be coming from.

In our first experiment, we theorised that the introduction of the punishment during training would stop bees from responding to the clean airflow before the presentation of the odour. The results show that this is not entirely the case, as many bees (mostly from the hungry group) still extended their proboscis. For the rewarded odour presentation and 70:30 compound, hungry bees spent significantly more time responding to the clean airflow. Interestingly, discrimination performance was not enhanced despite the fact that we introduced the punishment. This may be because the testing protocol was also made more difficult than that of the first experiment, and perhaps if we had used this testing protocol originally, the decision accuracies would have been reduced there.

An unusually large number of bees responded to the punished odour (mostly bees from the hungry group, but also a few satiated). This may have been because this odour was presented first in the testing phase. The 30:70 compound mixture was presented last, following the rewarded odour, and this generated the fewest responses. Since this odour was presented last, extinction may have had an effect, and this would explain why more bees responded to the punished odour than the 30:70 compound odour, even though the 30:70 compound had components of the rewarded odour and was thus more similar to it than the punished odour. This result also further indicates that the sequence in which odours are presented is very important. The first odour presented may generate more responses than usual as it will be the first time the bees experience the experimental set-up post conditioning. For both groups, bees were more likely to respond to the 70:30 compound than the 50:50 compound. These

two odours were presented in the middle of the sequence and the responses to these odours were predicted, as the former is more similar to the rewarded odour than the latter. We have seen in previous experiments that more similar odours invoke a higher amount of generalisation (Guerrieri et al. 2005) and this is indeed the case here.

Many of our results were consistent but insignificant. An example is response time; although hungry bees were, on average, faster than satiated bees, only in one case was this a significant difference. We also found this in the first experiment. We propose two reasons for this lack of significance despite consistency and sometimes quite large differences between groups (for example, see the bottom plot in Fig 4.14. There is quite a big difference in the total time the bees were responding, as well as the total number of proboscis extensions they exhibited). Firstly, in most cases, we lacked statistical power due to small sample sizes. As many of the bees didn't respond to the punished or compound odours (especially 30:70), we only had data from a few bees. To obtain more significant results, this experiment would need to be performed with $n > 80$ bees. Secondly, bees may need to be made hungrier and thus starved for a longer period of time. In this experiment, hungry bees went three hours without food before they were tested. If this interval is lengthened to perhaps six or more hours, we expect that the results will be more exaggerated (although it will be important not to starve bees to the point where they become unhealthy, as this will be an unfair comparison).

As mentioned before, the testing protocol used here is more akin to a go/no-go task than a two-alternative forced-choice task. Even with the simultaneous presentation of the trained odours in the compound mixtures, the honeybee can only either respond or not respond. Can the drift-diffusion model account for this sort of task? Indeed, the process of evidence accumulation should remain the same. There is a theoretical difference in terms of the thresholds used to trigger motor responses, however. A two-alternative forced-choice task requires the use of two thresholds; the accumulated evidence crossing

one threshold corresponds to a specific behaviour from the animal whilst the crossing of the other threshold corresponds to another behaviour. For a harnessed insect responding to odour presentations from the left and right, this may translate as a head turn in one direction or the other. For the go/no-go task, only one threshold is required, corresponding to a response (for example, a proboscis extension) or lack thereof. A drawback to using this paradigm is that it is unclear if a lack of response is due to the evidence accumulation process not yet terminating or if the individual has actively made the decision not to respond. Nonetheless, the go/no-go task is still suitable as the drift-diffusion model may be adapted to suit it (using one threshold) and still keep the underlying assumptions.

An alternative theory to be considered is that the responses to compound odours should not be marked as errors (however, even when the data are analysed with this criteria, the hungry bees are still significantly less accurate than satiated bees). Instead, those which respond should be deemed more accurate, since they are responding to the presence of the rewarded odour. With this theory in mind, it could instead perhaps be concluded that hungry bees are significantly more likely to generalise than satiated bees. However, seeing as significantly more hungry bees responded to the punished odour here, this most certainly allows us to draw the conclusion that hungry bees are making more errors. With the concept of generalisation in mind, however, it would be interesting to replicate the experiment of (Guerrieri et al. 2005) and examine how two different internal states impact the rates of generalisation. This is another possibility for future work.

All the data obtained from these experiments heavily suggest that honeybees are using a drift-diffusion process for evidence accumulation, and furthermore that the level of satiation can mediate the level of the decision threshold and in turn cause behavioural changes. This is in contrast to another theory of why response probability in bees changes according to hunger, which instead proposes that hunger causes a change in sucrose responsiveness (Page Jr et al.

1998). If we conclude that satiation can mediate the decision threshold, we can ask the following questions: where might this process take place within the bee brain, and can we imagine the olfactory system to act as a network that is integrating evidence over time? Furthermore, what sort of circuit is required to implement this efficiently, and could a computational model provide any insight into what properties such a network might need?

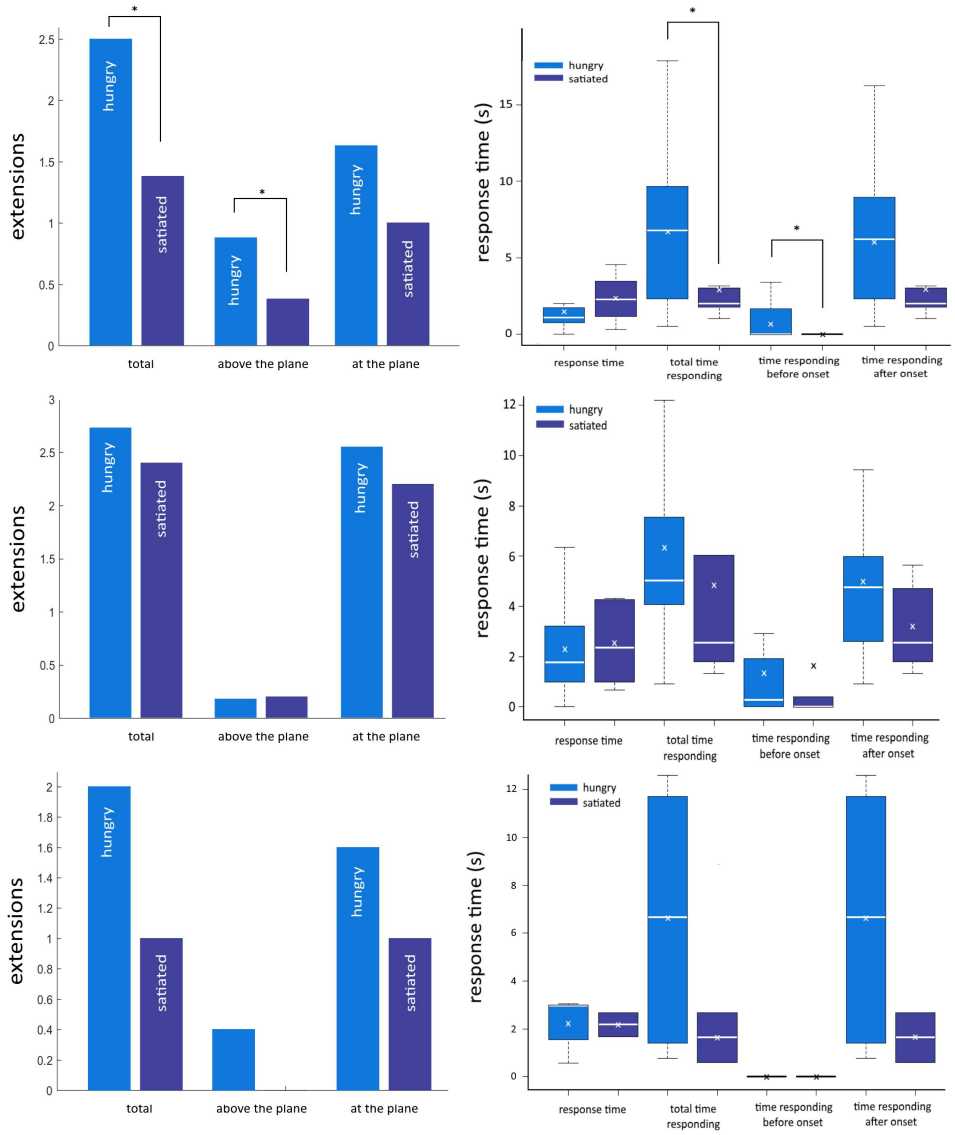

Figure 4.14: Complete data analysis for both motivational groups responding to the 70:30 (top), 50:50 (middle) and 30:70 (bottom) compound odours. Data presented as in Fig 4.13. Medians denoted by white lines, averages denoted by white (or black) crosses. Asterisks denote significance values. *: $p < .05$. See Appendix A for statistical tests used.

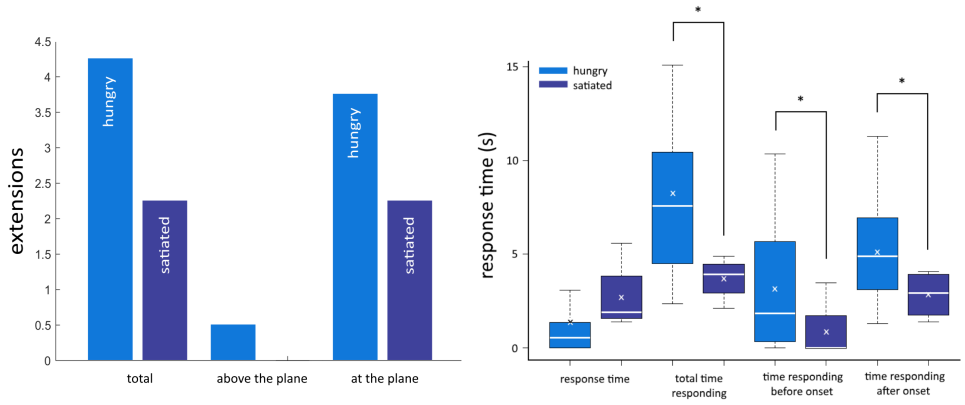

Figure 4.15: Complete data analysis for both motivational groups for the punished odour. Data presented as in Fig 4.13. Medians denoted by white lines, averages denoted by white crosses. Asterisks denote significance values. *: $p < .05$. See Appendix A for statistical tests used.

Chapter 5

A Computational Model of Decision-Making

How do animals make fast and robust decisions? Previous research has attempted to answer this question through a combination of behavioural and neurobiological experiments in conjunction with computational modelling. In particular, sequential sampling models that focus on the decision accuracy and decision time of an agent have been used to infer what neural mechanisms may be driving choice behaviours. Originating from psychology, the first and most well-known of these models is the drift-diffusion model (Ratcliff 1978). Since this model was developed a multitude of others have built upon it which vary in their complexity and biological plausibility (see Ratcliff and Smith 2004). One model in particular, the leaky competing accumulator model (LCA, Usher and McClelland 2001), introduces an important biological mechanism which was introduced in Chapter 3. We here use the LCA model, and an extension of it (Brown and Holmes 2001) which is described in detail later in this chapter, as inspiration for a biologically inspired model of decision-making. To give an idea of where the model fits in comparison with other sequential sampling models, a chart of the different models and their categories is given in Fig 5.1, though there are many others which are not depicted here.

Sequential sampling models have been successfully fitted to the behavi-

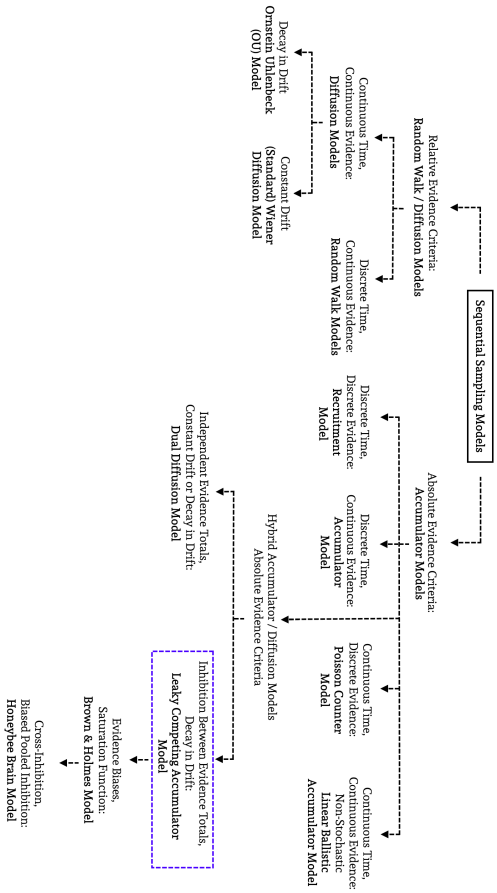

Figure 5.1: A range of sequential sampling models with the LCA model highlighted. Adapted from Ratcliff et al. 2016 with permission.

oural data obtained from experiments within a range of decision-making tasks (Ratcliff et al. 1999, Ratcliff and Rouder 2000, Ratcliff and Smith 2004, Ratcliff et al. 2004). They are usually developed with three underlying assumptions about decision-making processes (Bogacz et al. 2006): the first assumption is that evidence for each choice is being accumulated over time by the decision-maker, and that the evidence presented is in favour of each choice. Second, it is assumed that the accumulation process is noisy. Finally, sequential sampling models employ the mechanism of the decision threshold to ascertain when a decision has been made. The decision threshold determines the level of evidence that is required for the initiation of a response. As such, a decision can only be made when the decision-maker has sampled a sufficient amount of

evidence. A low decision threshold invokes a fast response and typically yields a reduction in accuracy. A high threshold allows for a slower response but a gain in accuracy since the model will be acting upon a level of evidence which is not so limited. This is known as the speed-accuracy trade-off (SAT), a behavioural phenomenon which has been shown to manifest in a variety of animals (Chittka et al. 2009).

In choice behaviour, the speed-accuracy trade-off has been shown to be ubiquitous. For example, it has been demonstrated in human (Bogacz et al. 2010a, Starns and Ratcliff 2010) and primate decision-making (Heitz and Schall 2012, Shadlen and Kiani 2013), as well as insects such as monarch butterflies (Rodrigues 2016), bumblebees (Chittka et al. 2003, Dyer and Chittka 2004b) and honeybees (Wright et al. 2009). The SAT has furthermore been demonstrated in collective insect decision-making, governing the behaviours of ant and honeybee colonies (Franks et al. 2003, Marshall et al. 2009). Although the concept may seem trivial, the fact that it is so prevalent across species suggests that it is very much worthy of being studied. Indeed, it has been argued to be the 'benchmark' for the decision process (Heitz 2014) and so understanding the neural mechanisms that govern this phenomenon may solidify our understanding of decision-making. Sequential sampling models offer insights into these mechanisms. As such, they play an invaluable role in decoding the decision-making process, at least for simpler, perceptual decisions where an animal's ability to discriminate between alternatives will impact their accuracy and speed.

What kind of mechanisms might a successful decision-making system employ? Already mentioned is that of lateral inhibition: a circuit especially researched in the visual domain (Goldstein and Brockmole 2016). Empirical evidence for this mechanism can be traced back to experiments with the horseshoe crab, wherein experimenters took advantage of the organisation of the eyes in order to record activity from individual receptors (Hartline et al. 1956). They demonstrated that stimulation of one receptor would cause it to respond,

however, the stimulation of the receptor's neighbours would cause a reduction in its activity. Importantly, the experiment brought forward the proposal that the receptors were not independent of each other, instead they were coupled such that each receptor could directly inhibit those surrounding it. Lateral inhibition is also a well-known circuit in classical decision-making models, used for introducing competition within competing integration channels and for implementing a winner-take-all output (Usher and McClelland 2001, Brown and Holmes 2001, Bogacz et al. 2006). The circuit has been shown to improve decision-making in a behavioural switching context (Marshall et al. 2015) and, remarkably, it has also been observed in collective decision-making in honeybees (Seeley et al. 2012, Pais et al. 2013), acting as a stop signal for bees which are opting for inferior choices. Recently, this circuitry has been suggested to exist in the invertebrate antennal lobes as another form of contrast enhancement, and it has been proposed that this mechanism should be dynamic and dependent on the similarity of the odours presented, such that inhibition is increased when the odours are similar (Linster et al. 2005). This hypothesis introduces the concept of an inhibitory mechanism which is robust and dependent on the composition of the presented stimuli.

To fully answer the question of what other mechanisms might be in place to ensure robust discrimination, it is necessary to study neurobiology. It may be tempting to use human or primate brains as the 'gold standard', however, the neural circuitry behind decision-making behaviours at this level is remarkably complex (Sporns et al. 2005, Azevedo et al. 2009). Instead, it is more reasonable to turn to a simpler animal which is fully capable of solving discriminatory problems. An increasing amount of research over the years has pointed to the honeybee as an ideal model for the study of decision-making and cognitive-like behaviours (Menzel and Giurfa 2006, Giurfa 2007, Menzel 2012). Indeed, this is an animal which has a rich behavioural repertoire and is capable of solving complex problems despite its relatively small brain. It is clear that this animal has evolved to solve decision-making problems efficiently,

making use of simple yet powerful neural circuits which are more amenable to study. It is for this reason that we use the honeybee for inspiration for this model. A particularly well studied domain within the honeybee brain is that of olfaction; a higher-level map of several olfactory circuits has been identified such that information flow from one centre to another can be traced. Indeed, honeybees have been shown to be very good at olfactory discrimination tasks (Guerrieri et al. 2005) and the simplicity of the honeybee olfactory system allows us to examine what is contributing to such a robust decision-making circuit.

The Honeybee Olfactory System

The olfactory system of the honeybee (and other invertebrates) has been well studied (e.g., Galizia 2014). In the honeybee, detection of olfactory stimuli begins at the antennae, where olfactory receptor neurons (ORNs) encode information about the presented odours and then innervate the antennal lobes (AL), which are sites dedicated to the processing of olfactory information. Within the antennal lobes are structures known as glomeruli, which collate odour information. Within the glomeruli are local interneurons (LNs) which project to multiple other glomeruli, and both inhibitory and excitatory projection neurons (iPNs and ePNs, respectively). The projection neurons project to higher-order brain regions such as the mushroom bodies (MB) and lateral protocerebrum (LP), both of which are regions thought to play important roles in decision-making. For a more in-depth review of the olfactory system, see Chapter 2. Within the antennal lobes, an inhibitory mechanism referred to as lateral inhibition has been shown to exist. This mechanism is thought to act as a decorrelator of neural signals, aiding in discriminatory behaviours and introducing contrast enhancement into the system.

Indeed, inhibitory circuits have been shown to be of great importance to decision-making and action selection circuits, both biologically and theoretically (Hensch et al. 1998, Bogacz et al. 2006, Hergarden et al. 2012, Pool et al.

2014, Marshall et al. 2015, Barron et al. 2015). In honeybee neurobiology, the function of the iPNs within the olfactory circuit has only recently been quantified. Recent studies have shown that they act as a form of gain control or contrast enhancement, increasing the differences between neural response maps in order to aid in discrimination (Parnas et al. 2013). This is hypothesised to be implemented by a 'high-pass filter' such that iPN inhibition selectively blocks low-frequency spike trains but allows high-frequency spike trains to pass (Parnas et al. 2013). Additionally, recent work by Liang and colleagues (Liang et al. 2013) suggests that iPNs regulate olfactory information in the lateral horn by suppressing the responses of a specific population of neurons (vlpr neurons) to some odours but not others. Thus, it can be concluded that iPNs target some, but not all, higher-order neurons (i.e., are selective)

Task Difficulty

One parameter that is often manipulated in discrimination tasks is that of task difficulty. An intuitive result obtained from behavioural experiments is that the harder the task, the more a decision-maker's accuracy will decline. Within the olfactory domain, this is thought to be caused by overlapping response profiles within the pre-processing areas of the brain (both in vertebrates and invertebrates), with the neural mechanisms employed to aid in discrimination (e.g., lateral inhibition) unable to sufficiently decorrelate neural signals (for the drift-diffusion model, more difficult choices will have a low signal- to-noise ratio (SNR), as such the model should also be less accurate). For example, it was found in *Drosophila* that the overlap between neural representations of odours could be described in terms of the Euclidean distance between neural activity vectors; similar odours were found to invoke similar response profiles (and thus shorter Euclidean distances) whereas dissimilar odours invoked dissimilar response profiles (and thus larger Euclidean distances) (Parnas et al. 2013). The derived Euclidean distances could be used to predict the discrimination

abilities of the animals within a strict two- alternative forced-choice paradigm. The study used behavioural chambers which were developed such that the left and right sides could present odours simultaneously, which allowed for choice behaviours to be evaluated based on what side animals preferred. When presented with an odour that fruit flies exhibit innate preferences for on one side and a dissimilar odour on the other, fruit flies show a clear bias for the side of the chamber which presents the preferred odour. Instead, if a similar odour is presented with the preferred odour, fruit flies will pick one side over another by chance, showing no preference for either side. As the Euclidean distance between the preferred odour and an alternative increases, fruit flies begin to show an increased bias for the side of the chamber which presents the preferred odour.

It is quite simple to control the similarity of odours, and honeybee choice behaviour in binary olfactory discrimination tasks has been studied quite profusely (Guerrieri et al. 2005). Previous experiments have shown that more similar odours induce a higher degree of generalisation in bees (or causes them to misidentify the 'correct' choice and thus make a higher proportion of errors). From these experiments, it can be concluded that task difficulty (or stimuli similarity) is highly influential on the choice behaviours of decision-makers and so it is a parameter which we factor into our model.

Classical decision-making models have historically struggled with the case of being presented with two alternatives that are very similar or the same perceptually. In the case of the DDM, the presentation of two equivalent stimuli results in the model integrating noise over time. The model thus remains in a deadlock and is unable to choose one option or the other. This is behaviourally implausible. There have been several methods proposed to overcome deadlock, for example, implementing a mechanism known as the 'urgency signal' such that the decision thresholds collapse over time (Cisek et al. 2009, Hanks et al. 2011). Whilst this does indeed break deadlock, a recent study found that models with static or fixed thresholds are better suited to

fit experimental data (Hawkins et al. 2015). We here make use of the Brown & Holmes model (Brown and Holmes 2001) as a deadlock breaker, which introduces accumulator biases and a saturation function on the inhibitory mechanism of the LCA. The dynamics of the model are such that, in the case of two equivalent alternatives, the model picks one or the other at random. The reason for this is described within the following section. We make use of the deadlock breaking parameters in our own model, as such it serves as an extension to the Brown & Holmes decision-maker.

Brown & Holmes

The Brown & Holmes model of decision-making (Brown and Holmes 2001) was developed as an extension to the original LCA model. As with most sequential sampling models, neural populations are accumulating noisy evidence over time, with each population integrating for one stimulus or the other (since this is a binary decision-making model). Like the original leaky competing accumulators, the neural populations are in competition with each other through laterally inhibitory connections and have a decay constant through which activation decreases (their definition of activation being 'a population-averaged analogy to membrane voltage' (Brown and Holmes 2001, p. 1). External stimuli are denoted as ρ_i in the original publication but are here adapted to I_i, such that they are consistent with the Bogacz et al., (2006) notation, which we also use for our own model definitions. The input stimuli are normalised such that when $I_i \neq 0, I_2 = 1 - I_1$. Two forms of 'priming biases' are introduced to the model which are not present in the original leaky competing accumulators. The first type of bias, denoted b_i, affects the neural populations independently. These biases may also be removed (as such $b_i = 0$). The second type, denoted as i_0, affects both populations. Lateral inhibition is modified such that it is implemented via a sigmoidal activation function:

$$f(x; g, b) = \frac{1}{1 + e^{-g(x-b)}} \tag{5.1}$$

where g, the 'activation gain', denotes the slope of the function, and an increase in this parameter results in a steeper curve. The 'activation bias', b, denotes the midpoint of the function, and thus when $x = b$ the slope is at a maximum. In the original publication, $b = 0.5$. To visualise the function and observe the changes in respect of activation gain g, the curve is plotted in Fig 5.2.

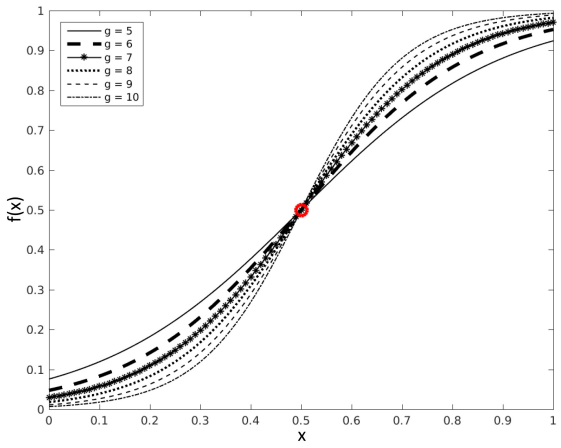

Figure 5.2: The Brown & Holmes activation function plotted for different values of g, the 'activation gain' or slope of the function. The red circle indicates the midpoint of the function curve where $x = b$.

Model Definitions

We now detail the definitions which describe the accumulator dynamics. The activation function $f(x; g, b)$ is from this point forward shortened to $f(x)$. As with other classical decision-making models, the model is described using differential equations:

$$\begin{aligned} dx_1 &= (I_1 + b_1 + h_0 - kx_1 - wf(x_2))dt \\ dx_2 &= (I_2 + b_2 + h_0 - kx_2 - wf(x_1))dt \end{aligned} \tag{5.2}$$

where I_i denotes the input stimuli, b_i the individual accumulator biases, h_0 the joint accumulator bias (originally i_0 in the Brown & Holmes model), k the decay constant (leak) and w (originally β) the strength of lateral inhibition.

To understand the deadlock breaking mechanisms of this model, the phase plane for equivalent values of I_i is illustrated in Fig 5.3.

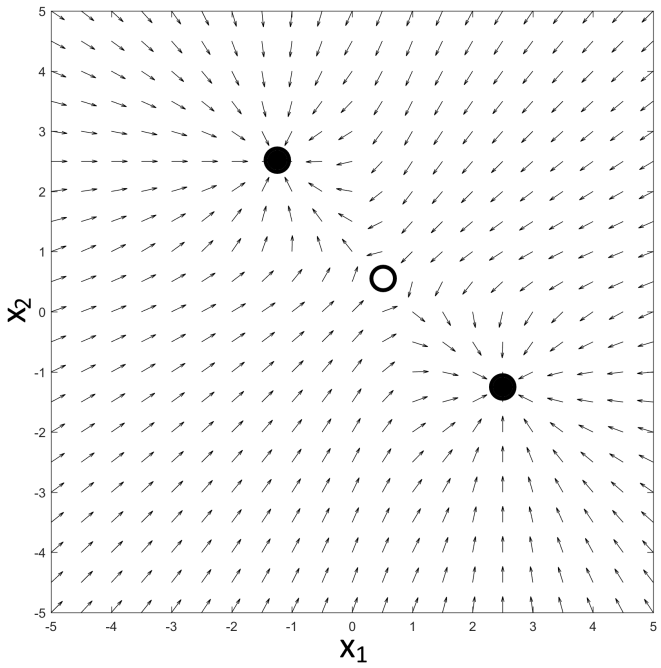

Figure 5.3: Phase plane for the Brown & Holmes model. Stable points denoted by closed circles, unstable points denoted by open circle. Accumulator activities converge at one stable point or the other. Parameters: $h_0 = 0.5, b = 0.5, g = 5, k = 0.2, w = 0.75, b_1 = 0, b_2 = 0, I_1 = 0, I_2 = 0$.

The phase plane serves as a visual representation of the fixed points in relation to the accumulator activities, as such we can see how the neural dynamics will change according to the state the system is in. Both the accumulators begin integration at the baseline activation, which is 0, and so from the origin they begin to increase. They will initially be drawn to an unstable fixed point which is close to the baseline. They are then repelled from this point, and the direction in which they go is determined by a small perturbation in the system. After this perturbation, the activities converge onto one of the stable points. The benefit of this is that, at these fixed points, one accumulator has a higher activation than the other, which is a crucial dynamic for decision-making. Exactly which stable point the accumulators will converge

at is random - this corresponds to the decision-maker selecting one alternative or the other without bias. This result is compatible with biological choice behaviour and thus solves the problem that other classical decision-making models faced when presented with two equivalent alternatives.

In order to see how these dynamics contribute to changes in neural activity over time, we simulated the model computationally in binary decision-making tasks. A single trial may last for any predetermined number of seconds, here the maximum has been set to 50 (it should be noted that this is a simulation of the *free-response paradigm*, where the decision-maker responds whenever the decision threshold has been crossed and as such it responds in its own time). We also set $dt = 0.01$, as in the (Bogacz et al. 2006) analysis. We here run simulations with the parameters according to the standard parameter set outlined in (Brown and Holmes 2001). Simulation results for a single choice task are shown in Fig 5.4. Additionally, since this model is an extension of the leaky competing accumulators, we also simulate the LCA model alongside the Brown & Holmes one. This is to further emphasise how the additional parameters are changing the decision process, such that the extended model is now able to make a choice at random.

Fig 5.4 makes it clear how the neural dynamics of the Brown & Holmes model aid in decision-making for two equivalent alternatives (where the difference in alternatives, Δv, is set to 0). Although the process of evidence accumulation is noisy in these decision-making models, here noise has been removed in order to emphasise the dynamics. Both of the leaky competing accumulators are stuck at the baseline level of activation and consequently the model never makes a decision. In contrast, the Brown & Holmes accumulators eventually diverge such that one of the accumulators is suppressed below the baseline whilst the other is accelerated beyond the decision threshold. Which accumulator reaches the threshold is random, and running multiple simulations will yield situations where the y_2 accumulator (where y_2 represents the activity level of the second accumulator, which is integrating evidence for the

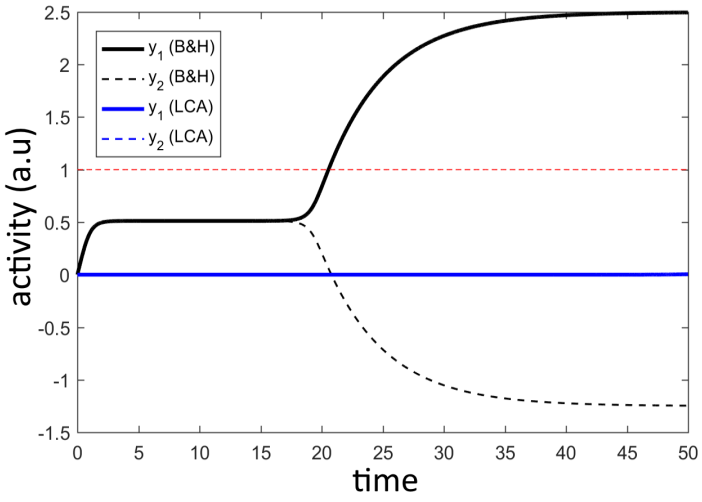

Figure 5.4: Accumulator activities of the Brown & Holmes model (black) compared with those of the original LCA model (blue). The decision threshold is denoted by the dashed red line. Note that noise has been removed (see main text). Parameters: $h_0 = 0.5, b = 0.5, g = 5, k = 0.2,\ w = 0.75, b_1 = 0, b_2 = 0, I_1 = 0, I_2 = 0$.

non-preferred alternative) will instead 'win' the trial.

It is important to note that these dynamics are dependent on the model parameters being used. In some cases, a difference in accumulator activations is never induced. This is demonstrated in Fig 5.5. Ideally, the difference in activities should be larger in order to aid in discrimination tasks; as such parameter sensitivity must be taken into consideration when running simulations.

From Fig 5.5 we can see that setting the joint bias (h_0) to 0.1 or 0.2 will result in the Brown & Holmes model behaving effectively like the leaky competing accumulators, where there is no difference between the activities of the two accumulators and as such the model is unable to make a decision. The best performances are obtained when $h_0 > 0.4$, and we use $h_0 = 0.5$ for our own simulations.

From this analysis we can conclude that the Brown & Holmes model is able to perform well in decision-making tasks where previous classical decision-making models failed. It builds upon the biologically plausible leaky compet-

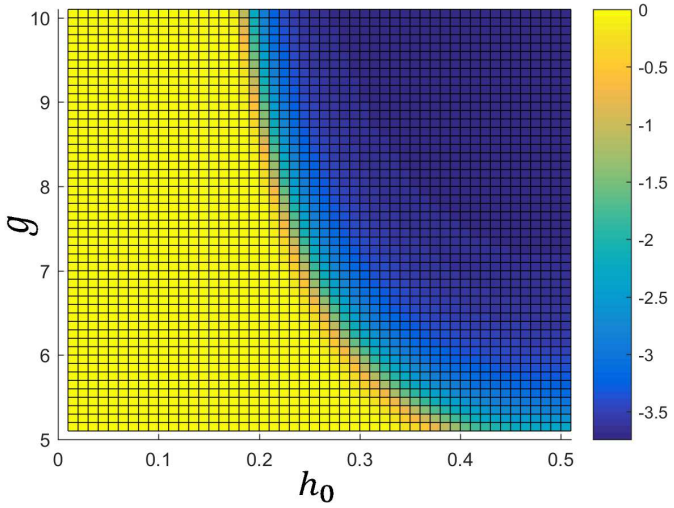

Figure 5.5: Examination of how the performance of the Brown & Holmes model changes according to parametric setup. Plotted is the difference in accumulator activities ($y_1 - y_2$) at the end of a trial, for systematically varied values of g (the slope of the activation function) and h_0 (the joint accumulator bias). Darker blue areas denote larger differences in accumulator activities and thus better performance. Yellow areas denote where the two accumulator activities are the same at the end of the trial, which is undesirable. Parameters are reproduced from the standard set.

ing accumulators and serves as a good approximation to the antennal lobe part of the decision-making circuit, where lateral inhibition is thought to be employed. We can think of these neural populations as the computational equivalent of the excitatory projection neurons of the antennal lobes, connected via lateral inhibition, the strength of which is mediated by the similarity of the odours presented. With this part of the decision-making circuit in place, we must now look to model the connections between the antennal lobes and the lateral protocerebrum. As mentioned previously, within the invertebrate olfactory system, ePN and iPN connections run in parallel, with ePN projections being excitatory and iPN projections being inhibitory, and pooled across glomeruli. The secondary part of the circuit can thus be imagined as a 'pooled-feedforward' model, which has been hypothesised to aid in discrimination by increasing the difference in the population activities (Parnas et al. 2013).

Weighted Lateral Inhibition

In the Brown & Holmes decision-maker, and indeed in other decision-making models, the strength of lateral inhibition between neural units is set to a constant value. However, there is evidence to suggest that inhibitory connections between antennal lobe neurons are dependent on glomerular response profiles, that is, a change of input brings about a change in inhibitory strength. As such, lateral inhibition should be mediated according to the stimuli that are presented.

In 1999, it was demonstrated in honeybees that usage of the GABA receptor antagonist picrotoxin impaired discrimination of similar odours but not dissimilar ones, suggesting that lateral inhibition is engaged to decorrelate similar odour response profiles but not dissimilar ones (Laurent et al. 1999). The authors noted that the neuronal responses to stimuli are often 'temporally structured' and suggested that the temporal dynamics of spike trains may be encoding information about stimuli. More recently, a study in *Drosophila* suggested that temporal patterns in the brain are dependent on the odours that have been presented (Wilson and Laurent 2005). The results demonstrated that GABAergic inhibition is employed to decorrelate the neural representations of odours, such that the differences between projection neuron activation levels are increased over time. That is, the similarity of neural representations decreased with GABA $_B$-mediated inhibition. This study also proposed that neural temporal patterns are a direct consequence of antennal lobe processing, as such patterns are not found in ORN responses to odourants.

Finally, research conducted by Linster and colleagues examined how temporal differences in neural responses may be taking place within the honeybee brain (Linster et al. 2005). Their study built upon older studies on both mammalian olfactory bulb and insect antennal lobe, where, as mentioned before, it has been shown that inhibitory networks are important for olfactory processing. Their work suggests that it is the response profiles of the glomeruli that determine the lateral connectivity between these neural structures.

Their computational model of the honeybee antennal lobes points to lateral inhibition being mediated by the response profiles of the glomeruli. Their network, wherein lateral inhibition is proportional to the similarity of the response profiles, was able to produce experimentally observed results (namely the output of PNs) more accurately than other networks. The study reiterates the point that odour representations in the brain are both spatially and temporally defined since representations are more dissimilar when leaving the antennal lobes than when entering them as ORN input.

Taken together, the results of these studies heavily suggest that lateral inhibition should not be constant between neural structures. It should be noted that, although the evidence for weighted lateral inhibition has been obtained from studies which have used different invertebrates, we here assume it can be generalised across groups (Galizia 2014). From these results we conclude that inhibitory strength should be dependent upon odour input, more specifically the similarity between odours (Δv in our model). As such, we modify the Brown & Holmes model to take the stimuli presented into account.

Pooled Feedforward Inhibition

The second 'layer' of the model has been designed to replicate the connectivity from the antennal lobe iPNs to the lateral protocerebrum. We thus employ a mechanism which has not been explored in sequential sampling models before, where input from the first layer of accumulators is summed and the total used as inhibition on the following layer of accumulators (Fig 5.6).

This inhibition is selective and it targets only one accumulator, representing the selective inhibition in the honeybee lateral protocerebrum (see Fig 5.9, in the model this layer is composed of the x_i and z populations. For the full definitions, see 'Materials & Methods'). Biologically, this denotes either an innate preference for one stimulus over another, or, a learned preference. The model can thus represent the state of a decision-maker post learning.

To see how this part of the model will aid in decision-making, as before

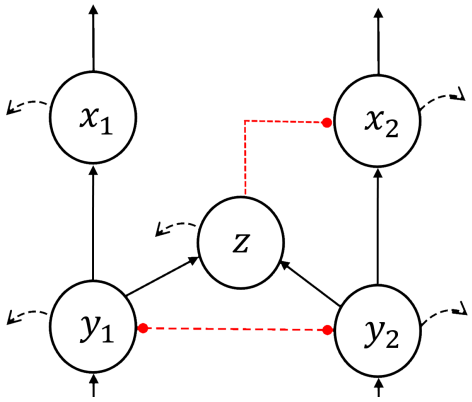

Figure 5.6: Schematic for pooled-feedforward inhibition part of the model.

with the Brown & Holmes model, we can analyse the phase plane. This is shown in Fig 5.7.

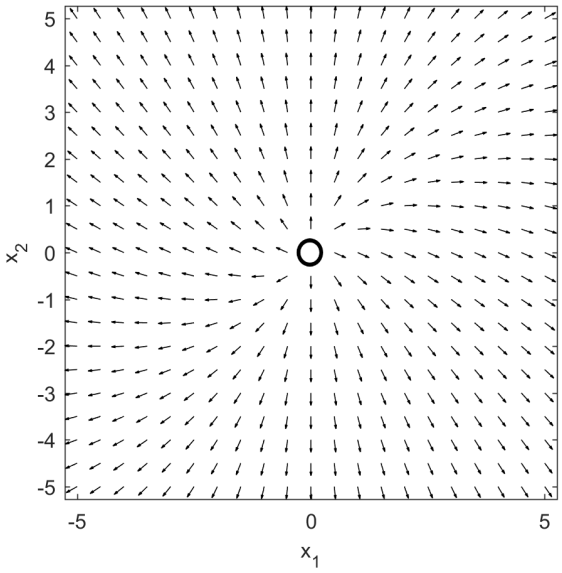

Figure 5.7: Phase plane for the pooled-feedforward layer of the model. Demonstrated is an unstable star node (unstable point denoted by open circle). Parameters: $I_1 = 0, I_2 = 0, h_0 = 0.5, w'' = 0.4, k'' = 0.1, g = 5, b = 0.5, k = 0.2$ $w = 0.75, \bar{k} = 0.2$.

We assume that the neural activities have settled into one stable point or another from the dynamics of the first layer (see Fig 5.3). At either of the stable points, the dynamics of the second layer will enhance the differences in accumulator activities, allowing one to increase whilst inhibiting the other. As

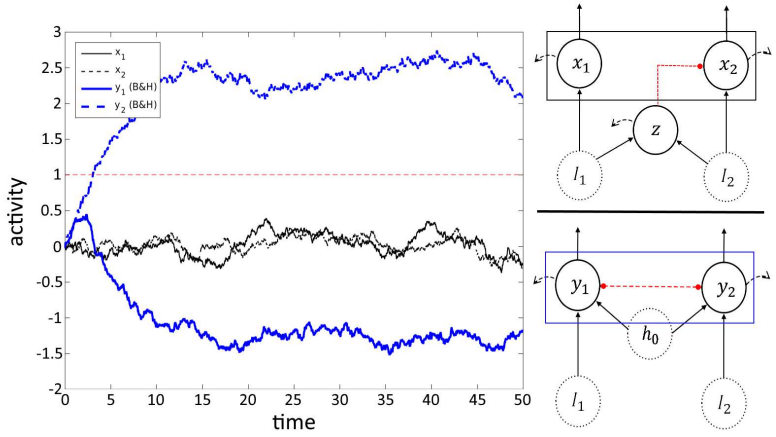

Figure 5.8: The pooled-feedforward inhibition model is unable to decorrelate when two equal alternatives are presented. Parameters: $I_1 = 0, I_2 = 0, h_0 = 0.5, w'' = 0.2, k'' = 0.3, g = 10, b = 0.5,\ k = 0.2, w = 0.75, \bar{k} = 0.2$.

such, the phase plane indicates that pooled-feedforward inhibition may indeed be acting as a decorrelator for neural signals, which agrees with the hypotheses put forward by previous research. This idea is tested in our simulations.

It is worth noting that this model, on its own, suffers fom the same problem as other classical decision-making models. With two equal alternatives, it is unable to separate the population activities and simply integrates noise over time. This is shown in Fig 5.8.

The figure shows the Brown & Holmes model (blue) and the pooled-feedforward inhibition model (black) implemented separately. In the presence of equal alternatives, the pooled-inhibition model hovers around the baseline level of activation and never reaches the decision threshold (red). However, since the model has been developed on the basis of honeybee neurobiological data (with the assumption that features of the honeybee olfactory system can be generalised to other invertebrates) it would make sense that such a model needs the first layer, where lateral inhibition is applied to introduce competition into the circuit and provide the initial decorrelation, in order to perform well. Because this model builds upon the dynamics of the first, it is not a problem that it is unable to decorrelate the signals of equal alternatives when implemented alone.

5.1 Materials and Methods

Our decision-making network is based on classical models of decision-making (Bogacz et al. 2006) and is composed of leaky accumulating units which are organised in a manner inspired by that of the invertebrate olfactory system (Fig. 5.9). It extends the Brown & Holmes model of decision-making (Brown and Holmes 2001) which makes use of priming biases as well as a neural activation function in order to invoke random choice behaviour upon the presentation of two alternatives equal in value. Taking inspiration from previous biological and computational studies, we alter lateral inhibition so that it is mediated according to the similarity of the alternatives presented.

We use our model of decision-making to simulate a T-maze olfactometer task. Here, two odours are presented to the network and it must discriminate between them and respond to the correct one (the one it would have been trained to prefer, as the network is in the post-learning state) within a given amount of time. Since we are simulating a two-alternative forced-choice task, the network is composed of two competing evidence accumulation channels; the first channel (denoted by integer 1 in the unit labels shown in Fig. 5.9) accumulates the preferred olfactory input and is thus the 'preferred' channel. In simulations, odours presented to the model invoke responses from the ORN populations and their mean firing rate is used for odour encoding. The odour which invokes a higher mean firing rate is denoted the preferred option, thus $I_1 > I_2$, and a greater difference in the means corresponds to a greater Euclidean distance within the neural representations (and therefore an easier task).

The pre-motor accumulator of the first channel, corresponding to neural units in the honeybee lateral protocerebrum, must reach the decision threshold in order for the network to respond to the preferred odour. At the end of each simulated trial the decision time is recorded, as well as whether or not the network made the correct decision.

We use this model to examine the impact of two inhibitory circuits to

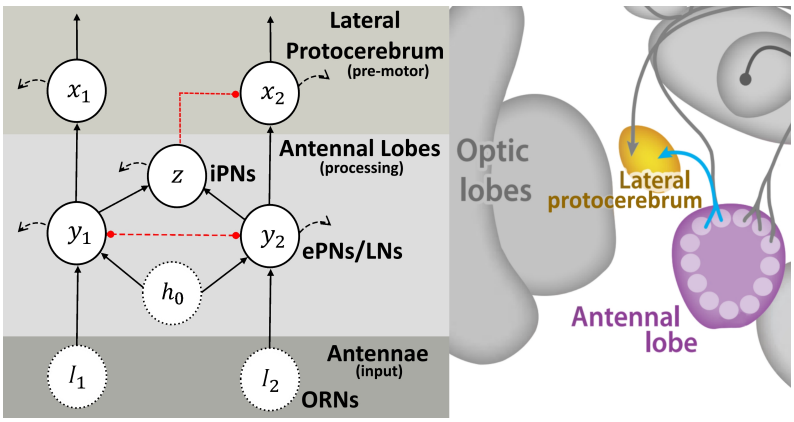

Figure 5.9: A biologically constrained network for binary decision-making. Right: frontal section of the honeybee brain, adapted from Perry and Barron 2013b with permission, with modelled brain regions highlighted in colour. A single olfactory tract, which bypasses the mushroom bodies and projects directly to the lateral protocerebrum, is shown in blue. Left: the network schematic. The decision network is composed of external sensory units which provide input to the system (dotted circles) and leaky accumulating units which integrate sensory information (solid circles). Neural pools are defined as follows: I_i populations denote ORNs, y_i denote antennal lobe neurons (both LNs and ePNs), z denotes the iPNs and x_i denote the pre-motor units in the lateral protocerebrum. Excitatory connections between units are denoted by solid black arrows, inhibitory connections by red filled circles, and neural leak by dashed arrows.

quantify their role in robust decision-making. Furthermore, we alter specific parameters to see how decision accuracy and time changes within a binary task. In particular, we vary the difficulty of the task (such that the presented odours are similar or dissimilar), the decision threshold (the level of which represents the internal state of the bee, discussed in the previous chapter), the strength of pooled inhibition arising from the iPN population and the plasticity of lateral inhibition within the antennal lobes.

Precisely what determines the level of the decision threshold, and causes some agents to respond more quickly than others, is currently unknown. Undoubtedly such a mechanism will be subject to individual differences, indeed, such was seen to be the case with bumblebees in a discrimination task using virtual flowers (Chittka et al. 2003). The level of training given to an agent prior to a task may also impact response time, however, it would be unreasonable to assume that agents completing the task will behave in the same

manner simply because they have received the same training. One factor that is often overlooked, perhaps due to the fact that it is difficult to control experimentally, is motivational state, and it is this factor that we wish to explore by means of the decision threshold within the model.

We now describe our network of accumulators for our model of decision-making, using ordinary differential equations to update the accumulator activity. The equations are defined as follows (*cf* Fig. 5.9)):

$$\begin{aligned} dy_1 &= (I_1 + h_0 - ky_1 - (\frac{w}{\Delta v})f(y_2))dt \\ dy_2 &= (I_2 + h_0 - ky_2 - (\frac{w}{\Delta v})f(y_1))dt \end{aligned} \tag{5.3}$$

where dy_i denote the activity of the antennal lobe accumulators, I_i denote the activity of the olfactory receptor neurons, h_0 represents a bias presented to both accumulators (otherwise known as the baseline activation), k denotes the strength of accumulator leak, w denotes the strength of lateral inhibition, Δv denotes the difference in chemical structure of the two presented odours, and $f(y_i)$ denote the neural saturation function used when implementing lateral inhibition (see Brown and Holmes 2001). Here, the strength of lateral inhibition is weighted according to the similarity of the two presented stimuli, Δv.

The group of inhibitory projection neurons are described as follows:

$$dz = (y_1 + y_2 - \tilde{k}z)dt \tag{5.4}$$

where $\tilde{k}$ denotes the leak of the iPN population. Finally, the pre-motor accumulators are described by the following:

$$\begin{aligned} dx_1 &= (y_1 - k'x_1)dt \\ dx_2 &= (y_2 - k'x_2 - w'z)dt \end{aligned} \tag{5.5}$$

where $\tilde{w}$ denotes the leak of the pre-motor accumulators and w' denotes the

strength of the inhibition from the iPN population.

The inputs of the model are normalised (Brown and Holmes 2001) and are determined as follows:

$$I_2 = \begin{cases} 1 - I_1, & \text{if } I_1 \neq 0 \\ 0, & \text{otherwise} \end{cases} \tag{5.6}$$

This model does not have the capacity to learn as the mushroom bodies have not been modelled here. Instead, we assume that the odour preference is a result of a training phase, thus the agent has been pre-trained and prefers one odour over the other as it is associated with a reward. Alternatively, it could be assumed that the odour preference is innate and evolutionarily determined. Innate odour preferences within a binary decision-making task have been explored quite recently with *Drosophila* (Parnas et al. 2013) and we can assume that this concept extends to honeybees. Here, the first odour (input from I_1) is considered the preferred stimulus and thus a correct decision corresponds to the first pre-motor accumulator, x_1, reaching the decision threshold first.

All results from the model are generated using what we refer to as the standard parameter set unless specifically stated otherwise. These parameters are in Table 5.1 and were derived from simulation results in order to produce good model performance in decision-making tasks.

Table 5.1: **The standard parameter set.**

	Value	**Function**
b	0.50	midpoint of saturation function
g	5.00	slope of saturation function
w	0.75	weight of lateral inhibition (static)
k	0.20	neural leak of antennal lobe units
$i0$	0.50	baseline neural activation; joint bias
w'	0.40	strength of pooled inhibition
k'	0.10	neural leak of pre-motor units
$\tilde{k}$	0.20	neural leak of iPN population
dt	0.01	time step used in Euler method
z	0.01 - 0.5	decision threshold
Δv	0.1 - 0.9	chemical difference / task difficulty

5.2 Results

We ran simulations of binary decision-making tasks using this model. In order to examine the full impact of the two inhibitory circuits (weighted lateral inhibition and pooled-feedforward inhibition), we first simulated the model with and without these circuits intact. We obtained an accuracy matrix for decision-making performances for the full model, under different task difficulties and levels of decision threshold (corresponding to the motivational state of the modelled animal). Finally, we then used the results obtained to infer how an animal's choice behaviours may differ according to task difficulty and motivational state, and we compared the theoretical data with our behavioural experiments performed in Chapter 4, as well as with another experiment performed by (Parnas et al. 2013) using *Drosophila*.

Inhibitory projection neurons enhance discrimination

We first investigated the role of pooled feedforward-inhibition in the decision-making process. In the model, a single accumulator (z) pools the antennal lobe activity and propagates the combined total forward as inhibition to a single channel in the lateral protocerebrum. This corresponds to the multiglomerular inhibitory projection neurons (iPNs) of the honeybee brain and functionally achieves selective attenuation of the non-preferred channel (see Fig 5.9, where only the non-preferred pre-motor accumulator is inhibited).

We thus compared the performance of two models: one with and one without the iPN population. We term the complete model 'WiPN' (with iPN) and the alternative model 'NiPN' (no iPN). We then ran decision-making tasks with these two models under the same conditions such that their neural dynamics could be compared. Analysis of the simulations showed that pooled inhibition enhanced the decorrelation of the neural signals and produced a greater difference in neural activity by the end of the trial. This is shown for several different simulations in (Fig. 5.10).

From the accumulator plot in Fig. 5.10, it is clear that both models are

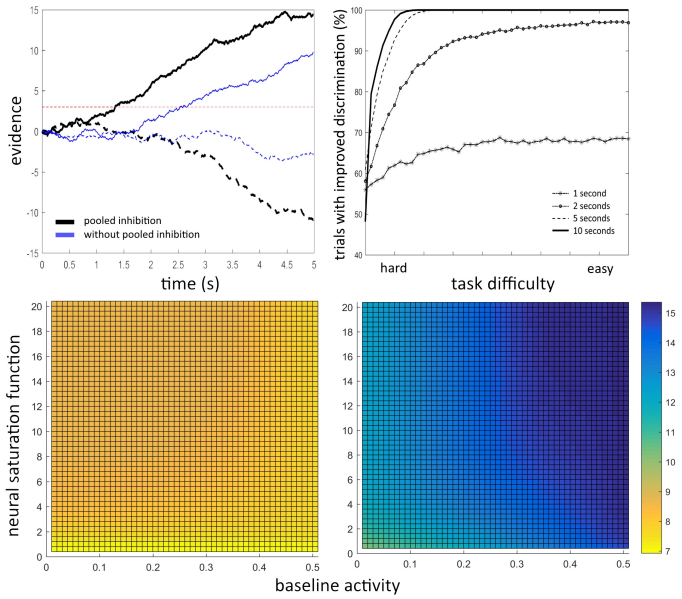

Figure 5.10: Pooled inhibition improves discrimination by decorrelating neural signals. Top-left: accumulator plot to show how the integrated evidence (neural activity) changes over the course of an easy discrimination trial for the pre-motor units. Solid lines denote preferred accumulator activity, dashed lines denote non-preferred accumulator activity. Black lines denote WiPN model, blue lines denote NiPN model. Top-right: the discrimination performances of the two models compared across ten thousand trials for tasks of varying difficulty and allowed integration time. Plotted is the percentage of trials where the WiPN model improves discrimination (thus, at 50 %, the WiPN model is performing as well as the NiPN model). Bottom: plots to show the absolute difference between the pre-motor units at the end of an easy decision-making trial (lasting for five seconds), for the WiPN model (right) and NiPN model (left), for varying baseline activation (h_0) and neural saturation (g) parameters. A higher absolute difference (blue areas) denotes enhanced decorrelation.

able to decorrelate the two neural signals. In this particular simulation, both models also make the correct choice, with the WiPN model selecting the preferred alternative slightly faster. However, by the end of the trial, the model with pooled inhibition was able to decorrelate its signals to a greater extent.

Further investigation showed that decorrelation was enhanced more consistently for easier tasks (performance plot in Fig. 5.10). Furthermore, enhancement was also dependent on the time allowed for the task: when the model was given more time to accumulate evidence the iPN population improved discrimination to a greater extent (Fig. 5.10). Although the performance curves for five and ten seconds are very similar, suggesting that five seconds is plenty of time for the pooled-feedforward model to decorrelate the signals, there is quite a large difference between one and two seconds. When allowed two seconds for evidence accumulation, and when given an easy task, the WiPN model successfully outperformed the NiPN model (in terms of decorrelation) around 95% of the time. For harder tasks, this percentage drops to around 60 - 70%. When five seconds are allowed, the difficulty of the task becomes less detrimental to performance.

We then performed a sensitivity analysis on the two models to see how decorrelation performance changed under varying parametric conditions. We varied two parameters of the original Brown & Holmes model, the slope of the neural saturation function (g) and the baseline activity or joint accumulator bias (h_0). The results are shown in the bottom row of Fig. 5.10. We found that, for all cases, the WiPN model was able to enhance decorrelation, with improvements being especially pronounced when h_0 was set to a higher value. As such, this model was able to outperform the NiPN model across parametric setups. Taken together, these results suggest how multiglomerular iPNs could act as an effective decorrelator of similar neural signals to improve discrimination.

We then investigated whether or not pooled inhibition would impact the overall decision accuracy and decision speed of the model. Since we hypothes-

ise that the motivational state of an animal will impact the speed-accuracy trade-off, we recognised that the level of the decision threshold would need to be taken into account, thus we ran simulations for tasks of varying difficulty with different decision thresholds. As with the results in Fig 5.10, we compared the results with the NiPN model to quantify the benefits of the iPN neurons.

The results are shown in Fig 5.11. The WiPN model tended to be slightly more accurate in decision-making, with particularly superior performances manifesting when the strength of pooled inhibition was increased beyond that of the standard parameter. However, pooled inhibition had little to no effect on the average decision speed of the WiPN model, even when the strength of inhibition was increased. The model took longer to make a decision in harder tasks than easier tasks, and this slowdown was enhanced by the implementation of a higher decision threshold. With a lower threshold, the model made decisions very quickly regardless of task difficulty. Taken together, these results may indicate that this mechanism will not impact the decision speed of an animal despite giving it a slight increase in accuracy. However, such a result also implies that the mechanism may be sub-optimal. As discussed before, drift-diffusion theory predicts that the speed-accuracy trade-off is unavoidable for optimal decision-makers.

Weighted lateral inhibition enhances discrimination for more difficult tasks

We next investigated the role of weighted lateral inhibition in the decision-making process. Within our full model, we varied the weight of the inhibitory connections according to the difficulty of the task such that a more difficult task resulted in stronger lateral inhibition between the antennal lobe populations. We compared this model to another which had static inhibitory connections. We refer to this as the SI model (static inhibition) and the model with weighted inhibition as the WI model. We ran comparisons akin to the

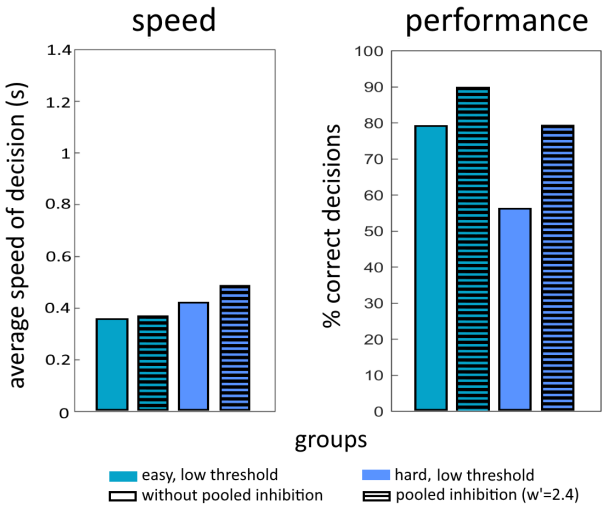

Figure 5.11: Pooled inhibition can improve decision accuracy without impacting decision speed. Left: average decision speeds for tasks of varying difficulty and threshold levels. Right: average decision accuracies for respective tasks. Averages in this figure were calculated from ten thousand trials. Low decision threshold: 0.1, higher decision threshold: 0.5.

ones for the feedforward-pooled inhibition part of the model, and the results are shown in Fig 5.12.

Analysis of the accumulator plot in Fig. 5.12 shows that, in contrast to the SI model, the WI model is able to effectively decorrelate its neural signals. It should be noted that this was a 'bad' simulation for the SI model, and that the performance shown here is not representative of the overall performance of the model. However, this emphasises the difference that weighted lateral inhibition can make.

Additional investigation showed that decorrelation was enhanced more consistently for harder tasks (performance plot in Fig. 5.12). As with the pooled-feedforward part of the model, enhancement was dependent on the time allowed for the model to integrate evidence. When the model was given more time, the introduction of weighted inhibition improved discrimination to a greater extent for more difficult tasks. Again, the performance curves for five and ten second trials are similar. The difference in performance between two and five seconds is here quite large, with the WI model in these cases being

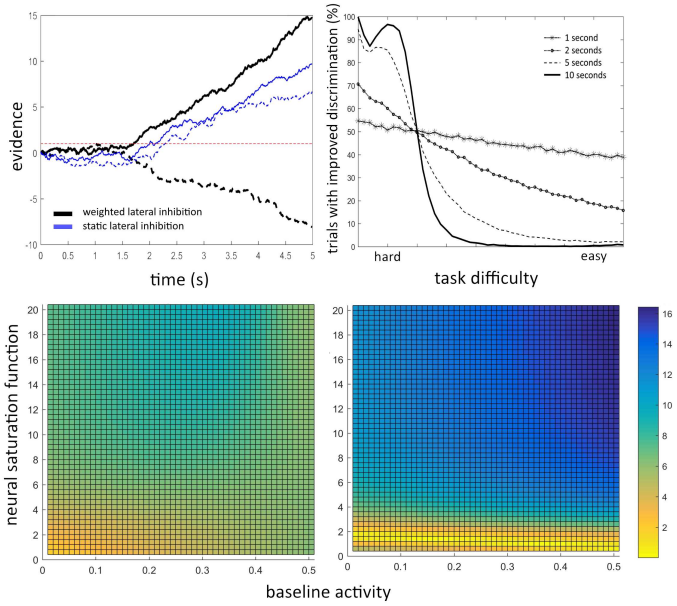

Figure 5.12: Weighted lateral inhibition gives rise to enhanced decorrelation. Top-left: accumulator plot for a single trial. Black lines denote the WI model, blue lines denote the SI model. Solid lines denote preferred accumulator activity, dashed lines denote non-preferred accumulator activity. Top-right: percentage of trials where the WI model improves discrimination. As before, a percentage of 50 % means that the WI model performs just as well as the SI model. Bottom: plots to show the absolute difference between the pre-motor units at the end of a harder decision-making trial, for the WI model (right) and SI model (left), for varying baseline activation (h_0) and neural saturation (g) parameters. A higher absolute difference (blue areas) denotes enhanced decorrelation.

able to outperform the SI model around 90 - 100 % of the time. For easy tasks, this performance drops, and the SI model instead is outperforming in terms of decorrelation. Interestingly, for pooled-feedforward inhibition, allowing five seconds or more for evidence accumulation results in the task difficulty having less impact on performance, however, for the lateral inhibition part of the model, the inverse is true. Here, when five seconds or more are allowed for the model to accumulate, it becomes increasingly sensitive to task difficulty, and easier tasks have a greater negative impact.

Finally, we then performed a sensitivity analysis on the two models to see how decorrelation performance changed under varying parametric conditions. As before, we varied two parameters of the original Brown & Holmes model, the slope of the neural saturation function (g) and the baseline activity or joint accumulator bias (h_0). The results are shown in the bottom row of Fig. 5.12. We found that, in most cases, the WI model was able to enhance discrimination. Improvements were especially pronounced for higher values of g and h_0, with the best performances being obtained for $h_0 = 0.5, g = 20$ (top right of matrix in bottom-right plot of Fig. 5.12). For the WI model, using smaller values of g reduced decorrelation performance. Regardless, this model was able to outperform the SI model across many parametric setups, and these results suggest that weighted lateral inhibition is particularly effective as a decorrelating mechanism for difficult decision-making tasks. However, since it is detrimental for easier tasks, this mechanism should perhaps be disengaged (neurobiological results suggest that this is the case, see brief discussion of weighted lateral inhibition above.)

A Model of Decision-Making can Predict Choice Behaviour

In this model, we used the level of the decision threshold to represent the motivational state of an animal, thus we make the assumption that internal state can directly impact the evidence accumulation process and change how much evidence is needed to initiate a response. This assumption was ex-

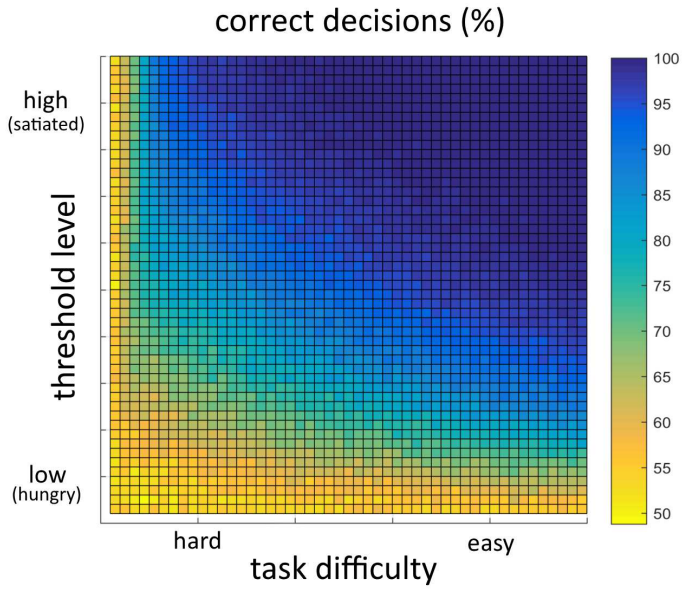

Figure 5.13: The full model, with both weighted lateral inhibition and pooled-feedforward inhibition, completed discrimination tasks of varying difficulties with different decision thresholds. Our underlying assumption is that a change in motivation causes an animal to alter its decision threshold, thus lower thresholds correspond to 'hungry' states and higher thresholds 'satiated' states. Dark blue areas denote high decision accuracy. Pixels highlighted with an asterisk signify data that have been used for comparison with data from the accompanying behavioural experiment.

plored in Chapter 4, where we used a well-known behavioural paradigm to test how honeybee choice behaviour differed in satiated and hungry bees. Before performing this experiment, we first used this model to make behavioural predictions. We ran decision-making trials where the decision threshold and task difficulty were systematically varied and performed ten thousand trials in each setup. We then recorded the average decision accuracy. The results are given in Fig 5.13.

The model results show that the accuracy of the decision-maker is dependent on task difficulty as well as motivational state. For task difficulty, the model predicts that the decision-maker will be more accurate in easier trials where the two odours presented are very dissimilar. For harder trials, the model suffers a reduction in accuracy. For two odours that are almost equivalent, the model picks the correct odour by chance. These results are in agreement with those of the study of (Parnas et al. 2013), which examined

Drosophila choice behaviour in a two-alternative forced-choice task.

Intuitively, the implementation of a low decision threshold caused a reduction in decision accuracy. In trials where the threshold was particularly low (corresponding to a very hungry animal), the model performance was reduced, even in tasks where two dissimilar odours were presented and should have been easy to discriminate between. As such the model predicts that hungry animals will make more mistakes than satiated animals when they are discriminating between alternatives, regardless of how easy the task is. The model performs the best when satiated and when discriminating between dissimilar alternatives (easier tasks), in some cases achieving 100% accuracy (thus choosing the correct option in each trial).

We here refer back to the behavioural data from Chapter 4. Although the results of this model cannot be directly applied to the experiments (as they were not true two-alternative forced-choice tasks), the model predictions can still be used to infer behavioural changes under the assumption that the process of evidence accumulation remains the same and that the underlying mechanisms still apply (noisy evidence accumulated over time, decision reached upon crossing the decision threshold, etc). In both our behavioural experiments, we saw that hungry honeybees were significantly less accurate than satiated bees. Furthermore, in our first experiment, this result was obtained from a choice task which used two dissimilar odours, as such they should have been easy to discriminate between. As such, these behavioural data are in agreement with the model.

5.3 Discussion

We developed a sequential sampling model which has connectivity based on honeybee neurobiology. Based on previous research in invertebrate olfactory systems, we here held the assumption that the principles of olfactory processing in honeybees can be generalised to other invertebrates. We replicated part of the olfactory system using a higher-level approach, focusing on pop-

ulations of neurons and how they interact with other brain centres. Using this model, we examined how two different inhibitory mechanisms may contribute to the discrimination abilities of an animal. We furthermore inferred how these two mechanisms might interact with other parameters which impact the evidence accumulation process, such as the decision threshold and task difficulty.

The results suggest that pooled-feedforward inhibition, a mechanism which is implemented in the honeybee brain via the multiglomerular iPNs, decorrelate neural representations of odours as they reach the pre-motor region of the brain, allowing for enhanced discrimination. This is achieved by means of targeted or selective inhibition, where the signal for the non-preferred odour is suppressed. We assume that this connection can be strengthened or weakened as a result of further learning. The results also suggest that the decision accuracy of the animal should improve due to the presence of these inhibitory neurons, since the accumulation channels integrating evidence for non-preferred options are being driven below the decision threshold.

Although selective inhibition could be a mechanism that arises due to learning or training, it could also be a consequence of evolutionary adaptation, resulting in the animal developing innate preferences for some odours over others. Previous research has suggested that this is the case. It has been shown, for example, that the aldehydes invoke responses (proboscis extensions) from bees more than other chemical groups (Guerrieri et al. 2005). Indeed, our preliminary data with honeybees show that untrained bees were more likely to respond to Hexanal, one of the aldehydes, than 1-Heptanol, on their first learning trial (around 20% of bees). Furthermore, honeybees have been shown to exhibit innate preferences for specific attributes of visual stimuli, for example, symmetrical or 'flower-like' patterns (Lehrer et al. 1995) and colour (Giurfa et al. 1995), as such it is reasonable to assume that they will have innate preferences for some odours over others.

We investigated whether lateral inhibition, a mechanism also implemen-

ted in the antennal lobes through the interactions of ePNs and LNs, would improve the decision-making process by being weighted according to odour similarity as opposed to being static. The results suggest that weighted lateral inhibition may have evolved to enhance neural decorrelation for harder discrimination tasks. For easier tasks, however, this mechanism eventually becomes detrimental and does not perform well as a decorrelator. As such, it may be the case that lateral inhibition is only 'switched on' for harder tasks, when decorrelation within the antennal lobes becomes crucial. Alternatively, there may be a minimum weight that inhibition cannot drop below.

Our sequential sampling model holds the assumptions that noisy sensory evidence is accumulated over time and that a response is triggered when the sampled evidence crosses the decision threshold. We further proposed that the decision threshold should be mediated by motivational state. As such, it is here predicted that the decision speed and accuracy of an animal within a choice task should be dependent on motivational state. The results of the model suggest that animals with low decision thresholds (corresponding to being hungry) should be more prone to error. The behavioural data presented in Chapter 4, which was obtained from real honeybees within a proboscis extension reflex paradigm task, agree with the model predictions (although it is important to note that they cannot be matched precisely, as the bees were not in a two-alternative forced-choice task). In our experiments, we found that hungry bees were indeed less accurate in their decision-making. This suggests that evidence accumulation is directly impacted by motivational state and that this state can thus alter decision accuracy and speed. Importantly, choice behaviours can differ between animals, even if they have been through the same training, if their motivational states are not controlled in some manner.

We found that the model results also agree with the behavioural data obtained by (Parnas et al. 2013). Their study, which used a T-maze to examine *Drosophila* choice behaviours within a two-alternative forced-choice task, demonstrated that odour similarity, or task difficulty, directly impacted the

way the flies behaved. More specifically, when the flies were presented with two dissimilar odours, one of which they innately preferred, they would exhibit a very clear bias for the preferred odour. Within the model, this corresponds to the decision-maker consistently choosing the correct odour over the incorrect odour. However, when the flies were presented with an odour that was very similar to the preferred odour, no preference was shown. This corresponds to the trials where the model chose the correct odour around 50% of the time, indicating no preference for either option.

It has also been suggested that fast and inaccurate decision-making may be a foraging strategy implemented by bees, such that they can maximise their nectar collection rate (Burns 2005). When a bee has depleted its food resources, mediation of the decision threshold may cause it to switch to implementing this strategy. In this state, waiting to accumulate more evidence is a time cost that the bee cannot afford, thus quick decision-making is perhaps the optimal strategy to implement. In the wild, honeybees (and other animals) face predatory threats and inaccurate decisions may increase the likelihood of an encounter, however, such an encounter will remain a possibility whilst starvation would be absolutely certain. In this situation, it can be inferred why speed should be favoured, and a change in foraging strategy is advantageous.

From these results it is suggested that more attention is given to evidence accumulation processes in decision-making and discrimination studies, and that motivational state should be controlled where possible. In many cases this may be too difficult, however, in tasks where animals are being trained to respond to food rewards, the level of satiation is relatively easy to control. Here, despite receiving the same training, honeybees in different motivational states responded differently during the testing phase: hungry bees struggled to discriminate in easy tasks which posed no problem for partially satiated bees. We recognise that other motivations, for example thirst and stress, are also likely to impact the decision-making process, however we do not model them here. Instead, we use the model to emphasise that decision-making is a

robust and dynamic process which can be influenced by internal states as well as external parameters (such as task difficulty).

Our model does not encompass every mechanism that has been implemented by the honeybee brain to aid in discrimination tasks. However, it has stressed the importance of two inhibitory mechanisms which are functionally different but achieve similar results: the decorrelation of neural signals, resulting in improved discrimination. The model has also suggested that these inhibitory mechanisms may interact with the motivational state of the bee to give rise to adaptive decision-making. From our network, we can infer how the brains of honeybees and other invertebrates may operate using these mechanisms.

Chapter 6

Future Work and Conclusions

This research has approached perceptual decision-making from the perspective of sequential sampling models, which assume that action selection rests upon the accumulation of evidence. Over the years, a tremendous amount of research has focused on the behaviours resulting from decision-making, and it is clear that it is a robust and dynamic process. The model that was developed here was based upon previous models of decision-making, the results of which emphasise the role of inhibitory networks within the invertebrate brain. Furthermore, the results of the biological experiment suggest that the crucial mechanism of the decision threshold is mediated by motivational state. Whilst the results are encouraging, it is important to keep in mind that the model that was developed here is no doubt an oversimplification of real decision-making mechanisms that ignores many of the aspects that will impact choice behaviour. Furthermore, the behavioural experiment used honeybees within laboratory settings wherein they were harnessed; as such, the conditions of the bees will have been affected. Nonetheless, these results may still give an insight into how certain neural circuits in the brain impact decision-making and how motivation can influence it. This research has presented several novel ideas that will hopefully be built upon in future projects.

The model of decision-making that was developed for this research attempted to replicate some of the olfactory pathways within the honeybee brain. It

is in no way, however, a complete picture of what is happening. A few of the limitations encountered with both the model and the biological experiment are here explored, and proposals for future work presented.

6.1 Level of Abstraction

When implementing biologically-plausible computational models, the question of how much detail to include within the model should always be considered. Some models, for example, will include tens of parameters for the sake of biological realism, whereas others will take a more 'black box' approach. The latter approach will be simpler, and often it is the case that such a model can still replicate real world data. More specifically for neural models, the complexity can range between modelling groups of neurons as a single unit, as was done for this research project, or individual neurons and their properties. Although it may seem intuitive to include as much detail as possible, this does not always produce more accurate results. Furthermore, more complex models require more computational power, and they may also introduce a drawback with regards to comprehensibility.

A typical neuron, at the most basic level, will be composed of a cell body (or soma), dendrites for input into the cell body and an axon to deliver outputs (Arbib 1995). Although there are many different types of neuron, computational models often simplify the matter and model neurons with as little complexity as possible, which is important for networks built up from hundreds or thousands of them. Of all the neuron models that exist, the point neurons are the simplest (for example, rate neurons which do not spike and integrate-and-fire neurons which do), although they are not used for many simulations as they are often assmed to be too simplistic to produce reliable results. On the other end of the spectrum is the Hodgkin-Huxley model (Hodgkin and Huxley 1952), inspired by data from the giant squid axon, which incorporates far more parameters which must be measured from biological experiments in order to replicate real world data.

With any of these models, what must be taken into account is how computationally expensive the model is and whether or not higher levels of detail will produce more accurate results or provide greater insight. For this research we decided to develop a decision-making model using a higher-level approach, which involved modelling neurons as populations. We thus assumed that each population could be representative of a region within the brain. For this kind of model, the details of the individual neurons must be ignored, which allows for a focus on the connectivity between the populations and how different regions might impact each other. Additionally, this kind of model assumes that all the neurons within a single population will function in an equivalent manner. This is not always the case. Furthermore, there was no spiking in the model, rather, activity was measured in terms of the mean firing rate of the populations. The end result was a model that was computationally inexpensive, which in turn meant that simulation results could be obtained quickly. Although the simulations presented in this work are not computationally intensive, searching the parameter space for model performance can quickly become a problem, certainly for searches comprising of three parameters (to show within a 3D performance plot, for example). If these searches are being performed for tasks wherein the decision-maker is given more time to make a decision (twenty or more seconds, for example), the time required for these simulations to finish will increase rapidly. As such, a model of this simplicity is beneficial.

It should also be noted that it is easier to start from a higher-level approach and add more biological detail to future versions of the model, rather than start with a more detailed model that is made increasingly abstract. Indeed, the former avenue has already been explored with quite a famous model of mammalian action selection. In 2001, Kevin Gurney and colleagues implemented a biologically-plausible population-level model of the basal ganglia, which is able to perform action selection (Gurney et al. 2001a, Gurney et al. 2001b). The model used leaky-integrator neurons and simulated neural inter-

action between regions of the mammalian basal ganglia. The model was able to reproduce some forms of action selection seen in animals with results that match neurobiological data, also incorporating dopaminergic modulation such that dopamine levels impacted choice performance. A few years later, this model of the mammalian basal ganglia was expanded upon (Humphries et al. 2006). This time, instead of using populations of neurons, the study implemented spiking neurons, a move that aimed to increase biological plausibility of the model. Here, the model was able to replicate some experimental data obtained from rats (the mean firing rates of certain brain regions) and indeed was able to reproduce the results of the original model. For example, setting dopamine to a low value resulted in the model not selecting any action, and higher levels of dopamine resulted in the model choosing more than one action. As such, it was shown that a higher-level model which focuses on populations of neurons can achieve the same results as one that incorporates individual spiking neurons. This model thus serves as a good example of one which was first developed to be a population model but was later successfully adapted to include a greater amount of biological detail.

6.2 Multi-Alternative Decision-Making

For this research we have focused on binary decision-making, however, choosing between only two alternatives is a simplification of real world situations. Within their natural environments, animals are far more likely to encounter multi-alternative choice tasks. Although binary decision-making tasks, such as the two-alternative forced-choice paradigm, can be well studied in laboratory settings, manipulated and adapted to suit different scenarios, they will not be as representative of real world decision-making as multi-alternative tasks. That said, it has been argued that binary decision-making is still worth studying: '*... it is representative of many problems faced by animals in their natural environments (e.g., whether to approach or avoid a novel stimulus)*' (Bogacz et al. 2006).

In light of this limitation, there are models which have been developed to build upon the original binary sequential sampling models and include more than two alternatives (Bogacz et al. 2007, Tsetsos et al. 2011). It is assumed that there are more than two integrators and, as with the two-alternative models, that each integrator is accumulating evidence for a single alternative. Tsetsos, Usher and McClelland develop multi-alternative extensions for the race model, LCA model and DDM in a recent publication (Tsetsos et al. 2011). As the model developed here is an extension of the LCA model, it should be possible to make a multi-alternative variant. It would be interesting to see how the model performed when asked to discriminate between more than two stimuli. A drawback to this approach, however, will be incorporating the added complexity into behavioural experiments (as well as an increase in computational power required for the model).

6.3 Honeybee Mushroom Bodies

The honeybee has a high capacity for learning and memory. This insect is required to navigate large distances in order to forage for food and return to the hive safely, detect rewarding flowers during foraging flights, communicate the whereabouts of available food sources to other hive mates, and contribute to collective decision-making. Furthermore, even though the honeybee does indeed have a brain the size of a grass seed, it has a rich behavioural repertoire (Menzel et al. 2001, Giurfa 2003a, Menzel and Giurfa 2006, Srinivasan 2010) Indeed, behavioural and neurobiological studies indicate that the honeybee is important to research in learning, memory and decision-making, despite its relatively simple nervous system. It has been shown to be able to learn abstract properties, such as orientation or colour, of stimuli (Van Hateren et al. 1990, Horridge et al. 1992, Horridge and Zhang 1995, Giurfa et al. 1996) and other concepts such as ‘sameness’ or ‘difference’ (Giurfa et al. 2001), learn two concepts simultaneously (Avarguès-Weber et al. 2012), navigate through complex mazes using visual cues (Zhang et al. 1999, Zhang et al. 2000) and

solve contextual problems (Collett and Kelber 1988, Zhang et al. 2006). Since decision-making can be influenced by learning, it is intuitive that it is beneficial to incorporate into a model. In particular, the learning of abstract properties of enables an agent to apply its knowledge to a decision-making task which introduces novel stimuli.

The model of decision-making that has been developed for this research is a higher level approach to modelling the honeybee brain. However, it is in no way a complete model and it does not incorporate all of the brain centres that are involved with decision-making and action selection, for example, the mushroom bodies and central complex. Furthermore, we here focus only on the olfactory centres. The model aimed to replicate some of the more well studied olfactory pathways, however, since there are no mushroom bodies, it cannot model them all. Another impact of this is that the model in its current state is unable to demonstrate any form of learning. The decision-maker is wired to exhibit a preference for one stimulus over another as a result of training, however, this preference cannot be changed. Since the model is making a choice within the proboscis extension reflex paradigm, a choice task that is traditionally used to study honeybee learning and memory, it would be desirable to incorporate aspects of learning into the model in order to replicate the learning curves that are derived from real PER experiments. It would be possible, for example, to adapt the model such that specific neural connections are strengthened in the presence of positive reinforcement or reward. Indeed, this has shown to be an effective method of implementing learning (Vasilaki et al. 2009).

During this project, work on such an extended model already began but currently remains in the theoretical stage. The more complete version includes the mushroom body kenyon cells and the extrinsic neurons which connect to other higher level brain centres, and has the capacity to learn to prefer one stimulus over another, or to reverse this preference. Since the theoretical foundations are in place, the computational implementation is the next

stage. This extended model would perhaps be the first thing to work on in future research, as it would give a more complete picture of the olfactory pathways within the brain. Indeed, since two of the more well-known pathways travel through the mushroom bodies, full decision-making behaviours cannot be modelled unless these brain regions are included. It would be interesting to examine the contribution of these pathways to choice behaviours, as this is still being studied.

6.4 Computational and Experimental Comparisons

Much of the previous research in decision-making makes use of the two-alternative forced-choice task. Indeed, our sequential sampling model also aims to replicate behaviours for such a task. In this paradigm, both stimuli are presented to the animal at the same time. The biological experiment, which was devised to accompany the computational model, cannot strictly be an example of a two-alternative forced-choice task since the odours are presented sequentially as opposed to simultaneously. As such, the model predictions and behavioural data cannot be directly compared. This is quite a large drawback, however, it does produce two avenues for possible future research:

- Modification of the computational model so that the two odours are presented sequentially, as in the experimental paradigm
- Modification of the behavioural experiment so that the two odours are presented simultaneously. This involves a change in experimental apparatus, such that each of the odours are presented to one antennae of the bee only, and a choice is determined by movement of the head in one direction or another

Of these two options, the former is perhaps the easier avenue as it involves an adaptation of the model code, which will undoubtedly also invoke a change in the parametric setup. However, computational simulations are easy to run

and data can be obtained rather quickly. The latter option involves modification of the device used for odour delivery as well as how the bee is harnessed, since the traditional harnessing methodology for PER was devised to restrict movement of the head. The olfactory controller used for gathering the experimental data is not programmed to deliver two odours at once, as such this will also require modification. Changing the experimental equipment and gathering new behavioural data will perhaps be a slower process, however, it would be necessary in order to examine honeybee decision-making within a true two-alternative forced-choice task.

6.5 Decision Boundaries

We here suggested that an animal's decision boundary or threshold can be directly influenced by the motivational state of the animal, thus inducing observable behavioural changes in the animal's choice behaviour. This theory was supported by our experimental results. In our computational model, the threshold is a parameter which can be modified, however, the neural circuitry which actually sets the level of this threshold is yet to be identified. Furthermore, we do not know where in the invertebrate brain this circuitry might exist. In human decision-making, the subthalamic nucleus (STN) has been recently identified as a brain region which modulates the decision threshold (Herz et al. 2016) during perceptual decision-making. The STN is a brain region which has previously been identified to play a role in action selection, also featuring in the basal ganglia model by Gurney and colleagues mentioned above. This result is very encouraging, and since recent studies also bring to light the parallels between mammalian and invertebrate brain structures (Strausfeld and Hirth 2013), perhaps it can be assumed that an invertebrate equivalent can be found. The central complex and lateral protocerebrum are two higher level brain centres in invertebrates that are known to play a role in action selection but are currently rather understudied; it would be interesting to see whether or not a mechanism which mediates the decision threshold

exists in either of these regions.

6.6 Conclusions

This research has presented an examination of decision-making and action selection using both experimental and theoretical approaches. An underlying theme throughout this work has been the honeybee, an animal which has evolved to solve decision-making problems efficiently. Specifically, we have explored the role of inhibitory mechanisms and how they might benefit a decision-maker within two different models, one of which was inspired by honeybee neurobiology. Further, we have explored the role of motivation in honeybees, and quantified how an animal's satiation level may impact their choice behaviours. We have also attempted to describe the observed behavioural data using the classical drift-diffusion model of Ratcliff (Ratcliff 1978).

In Chapter 3, we presented an abstract model of action selection within a foraging context, based upon a previous model of behavioural switching developed by Houston and Sumida (Houston and Sumida 1985). The results showed that using lateral inhibition to couple the two competing motivations of the modelled animal improved action selection as the animal reduced its costly behavioural switching. We found that mediation of the inhibitory circuit provided further benefits to the animal and allowed it to decrease its foraging bout lengths consistently over time, a result which is more biologically plausible. However, in order to exhibit this behaviour, the animal began switching between alternatives more frequently and consequently incurred more costs, in some cases. This may indicate a trade-off between bout length reduction and behavioural switching.

In Chapter 4, we investigated the role of satiation in honeybee choice behaviours using the proboscis extension reflex (PER) paradigm and found that they can be described by a drift-diffusion process. In our two experiments, we found that there was a clear difference in the decision-making behaviours of hungry and satiated animals. In both experiments, we found that hungry

bees were more inaccurate than satiated bees. We furthermore analysed their reaction times and found that, for some odours, hungry bees were responding significantly faster than satiated bees. In our second experiment, we expanded on these results and analysed the behaviours of honeybees in more detail. We found that, in some cases, hungry bees would respond before stimulus presentation significantly more than satiated bees. We furthermore found that, in the event of an error, satiated bees would withdraw their proboscis significantly quicker than hungry bees. We found that all these differences in their behaviours are compatible with a drift-diffusion process, if we assume that the decision threshold is lowered according to level of satiation. A reduced threshold would make a decision-maker more inaccurate, quicker in responding and slower in correcting errors. Thus, from these results we suggest that differences in behaviours are caused by a mediation of the decision threshold, according to level of satiation. This is an additional theory to that put forward by Page and colleagues, which suggested that a change in response probability was caused by a change in sucrose sensitivity (Page Jr et al. 1998). If the sucrose sensitivity of an individual is defined by the decision threshold, an increase in sensitivity would reflect a lower threshold and thus a higher response probability.

Finally, in Chapter 5, we presented a novel computational model of decision-making, the connectivity of which was based on the honeybee brain and inspired by the previous decision-making models of Ratcliff (Ratcliff 1978), Usher and McClelland (Usher and McClelland 2001), and Brown and Holmes (Brown and Holmes 2001). We used this model to further investigate the role of inhibitory mechanisms in decision-making, and to see if this model could match the behavioural data from our experiments as well as others. We found that the mechanisms of lateral inhibition and pooled-feedforward inhibition, both of which have been identified as circuits within the invertebrate antennal lobes, were beneficial to aiding in discrimination tasks. Our results agree with previous theories that lateral inhibition should not be static, rather, it should

be robust and dependent on stimuli similarity. We found that weighted lateral inhibition was more beneficial to harder decision-making tasks, whilst pooled-feedforward inhibition was more beneficial to easier tasks. Our model also demonstrated that increasing the similarity of two stimuli presented within a two-alternative forced-choice task should make a decision-maker increasingly more inaccurate. We found that, for two stimuli that are almost equivalent, the decision-maker will exhibit no bias in picking one stimulus or the other. This result is in agreement with the results in the study of Parnas et al., (2013). Finally, we found that a reduced decision threshold caused the model to make more mistakes, even for easier tasks that should have been easy to solve. Although this result cannot be directly matched with our experimental study (since they use two different paradigms), the results indicate that a hungry animal should make decisions faster and be more inaccurate.

This research, which utilised both computational neuroscience and biology, has further emphasised the benefits of using the honeybee as a model for decision-making, and has also demonstrated the power of abstract and higher-level computational models.

Appendices

Appendix A

Statistical Tests

All statistical tests were carried out using R. For the decision accuracy data the test of given proportions was used. For all other data, normality was first tested using the Shapiro-Wilks normality test. If the data sets to be compared were both normal, a two sample t-test was performed, and if normality could not be assumed, a Wilcoxon-Mann-Whitney test was performed. The full results from the statistical analyses are presented here.

A.1 Experiment One

We first present the tests for the original experiment.

A.1.1 Decision Accuracy

```
prop.test(c(12,31),c(41,43), correct=FALSE, alternative = "less")
     2-sample test for equality of proportions without continuity
     correction
data:  c(12, 31) out of c(41, 43)
X-squared = 15.405, df = 1, p-value = 4.337e-05
alternative hypothesis:  less
95 percent confidence interval:
     -1.0000000 -0.2660136
sample estimates:
     prop 1 prop 2
0.2926829 0.7209302
```

Result: $p < .001$**; hungry bees significantly less accurate than satiated bees.**

A.1.2 Reaction Time Data: Shapiro-Wilks Normality Tests

Before presenting the statistical tests, we first determine whether or not the data sets are normal.

Rewarded odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.86824, p-value = 0.0004324
```

Result: This data set is normal.

Rewarded odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.83631, p-value = 0.0004027
```

Result: This data set is normal.

Unrewarded odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.88542, p-value = 0.2949
```

Result: This data set is not normal.

Unrewarded odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.82289, p-value = 0.01292
```

Result: This data set is normal.

A.1.3 Reaction Time Statistical Tests

Rewarded odour

```
t.test(..., paired=FALSE)
     Welch Two Sample t-test
t = 0.751, df = 60.038, p-value = 0.4556
alternative hypothesis:  true difference in means is not equal to
0
95 percent confidence interval:
     -0.4124763 0.9083943
sample estimates:
mean of x mean of y
```

```
     1.828649 1.580690
```

Result: For the rewarded odour, hungry bees were not significantly faster than satiated bees.

Unrewarded odour

```
wilcox.test(..., paired=FALSE, alternative="greater", correct=FALSE)
     Wilcoxon rank sum test
W = 59, p-value = 0.03964
```

Result: $p < .05$ for the unrewarded odour, hungry bees were significantly faster than satiated bees.

A.2 Experiment Two

A.2.1 Decision Accuracy (Overall)

```
prop.test(c(12,27),c(40,41), correct=FALSE, alternative = "less")
     2-sample test for equality of proportions without continuity
     correction
data:  c(12, 27) out of c(40, 41)
X-squared = 10.425, df = 1, p-value = 0.0006216
alternative hypothesis:  less
95 percent confidence interval:
     -1.0000000 -0.1881172
sample estimates:
     prop 1 prop 2
0.3000000 0.6585366
```

Result: $p < .001$; hungry bees significantly less accurate than satiated bees.

A.2.2 Decision Accuracy (100:0)

```
prop.test(c(36,32),c(40,41), correct=FALSE, alternative = "greater")
     2-sample test for equality of proportions without continuity
     correction
data:  c(36, 32) out of c(40, 41)
X-squared = 2.1464, df = 1, p-value = 0.07145
alternative hypothesis:  greater
95 percent confidence interval:
     -0.01237078 1.00000000
sample estimates:
     prop 1 prop 2
```

```
0.9000000 0.7804878
```

Result: $p > .05$; the proportion of hungry bees responding to this odour is statistically the same as the proportion of satiated bees responding.

A.2.3 Decision Accuracy (70:30)

```
prop.test(c(16,8),c(40,41), correct=FALSE, alternative = "greater")
     2-sample test for equality of proportions without continuity
     correction
data:  c(16, 8) out of c(40, 41)
X-squared = 4.076, df = 1, p-value = 0.02175
alternative hypothesis:  greater
95 percent confidence interval:
     0.04179288 1.00000000
sample estimates:
     prop 1 prop 2
0.400000 0.195122
```

Result: $p < .05$; significantly more hungry bees responded to this odour than satiated bees.

A.2.4 Decision Accuracy (50:50)

```
prop.test(c(12,5),c(40,41), correct=FALSE, alternative = "greater")
     2-sample test for equality of proportions without continuity
     correction
data:  c(12, 5) out of c(40, 41)
X-squared = 3.8706, df = 1, p-value = 0.02457
alternative hypothesis:  greater
95 percent confidence interval:
     0.03220594 1.00000000
sample estimates:
     prop 1 prop 2
0.3000000 0.1219512
```

Result: $p < .05$; significantly more hungry bees responded to this odour than satiated bees.

A.2.5 Decision Accuracy (30:70)

```
prop.test(c(5,2),c(40,41), correct=FALSE, alternative = "greater")
     2-sample test for equality of proportions without continuity
     correction
```

```
data:  c(5, 2) out of c(40, 41)
X-squared = 1.4898, df = 1, p-value = 0.1111
alternative hypothesis:  greater
95 percent confidence interval:
     -0.02605422 1.00000000
sample estimates:
     prop 1 prop 2
0.12500000 0.04878049
```

Result: $p > .05$; the proportion of hungry bees responding to this odour is statistically the same as the proportion of satiated bees responding.

A.2.6 Decision Accuracy (0:100)

```
prop.test(c(16,4),c(40,41), correct=FALSE, alternative = "greater")
     2-sample test for equality of proportions without continuity
     correction
data:  c(16, 4) out of c(40, 41)
X-squared = 9.9597, df = 1, p-value = 8e-04
alternative hypothesis:  greater
95 percent confidence interval:
     0.1539698 1.0000000
sample estimates:
     prop 1 prop 2
0.40000000 0.09756098
```

Result: $p < .001$; significantly more hungry bees responded to this odour than satiated bees.

A.2.7 Reaction Time Data: Shapiro-Wilks Normality Tests

Before presenting the statistical tests, we first determine whether or not the data sets are normal.

Rewarded odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.90861, p-value = 0.0119
```

Result: This data set is normal.

Rewarded odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.91056, p-value = 0.007696
```

Result: This data set is normal.

70:30 compound odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.96126, p-value = 0.822
```

Result: This data set is not normal.

70:30 compound odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.67751, p-value = 0.0002182
```

Result: This data set is normal.

50:50 compound odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.85499, p-value = 0.2108
```

Result: This data set is not normal.

50:50 compound odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.84844, p-value = 0.05565
```

Result: This data set is not normal.

30:70 compound odour (satiated bees):

Sample size too small to perform test.

30:70 compound odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.7936, p-value = 0.07179
```

Result: This data set is not normal.

Punished odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.76024, p-value = 0.04793
```

Result: This data set is normal.

Punished odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
```

```
W = 0.75913, p-value = 0.00461
```

Result: This data set is normal.

A.2.8 Reaction Time Statistical Tests

Rewarded odour (100:0)

```
t.test(..., paired=FALSE, alternative="greater")
     Welch Two Sample t-test
t = 2.2001, df = 53.013, p-value = 0.01609
alternative hypothesis:  true difference in means is greater than
0
95 percent confidence interval:
     0.2270901 Inf
sample estimates:
mean of x mean of y
     2.779355 1.829429
```

Result: $p < .05$ for the rewarded odour, hungry bees were significantly faster than satiated bees.

70:30 compound odour

```
wilcox.test(..., paired=FALSE, alternative="greater", correct=FALSE)
     Wilcoxon rank sum test
W = 74.5, p-value = 0.1032
```

Result: $p > .05$ for the 70:30 compound odour, hungry bees were not significantly faster than satiated bees.

50:50 compound odour

```
wilcox.test(..., paired=FALSE, alternative="greater", correct=FALSE)
     Wilcoxon rank sum test
W = 25.5, p-value = 0.4756
```

Result: $p > .05$ for the 50:50 compound odour, hungry bees were not significantly faster than satiated bees.

30:70 compound odour

```
wilcox.test(..., paired=FALSE, alternative="greater", correct=FALSE)
     Wilcoxon rank sum test
W = 4, p-value = 0.7143
```

Result: $p > .05$ for the 30:70 compound odour, hungry bees were not significantly faster than satiated bees.

Punished odour (0:100)

```
t.test(rt13, rt14, paired=FALSE, alternative="greater")
     Welch Two Sample t-test
t = 0.43671, df = 7.3036, p-value = 0.3375
alternative hypothesis:  true difference in means is greater than
0
95 percent confidence interval:
     -1.821249 Inf
sample estimates:
mean of x mean of y
     2.71 2.16
```

Result: $p > .05$ for the punished odour, hungry bees were not significantly faster than satiated bees.

A.2.9 Number of Bouts: Shapiro-Wilks Normality Tests

Rewarded odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.79285, p-value = 3.816e-05
```

Result: This data set is normal.

Rewarded odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.8455, p-value = 0.0001502
```

Result: This data set is normal.

70:30 compound odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.6412, p-value = 0.0004791
```

Result: This data set is normal.

70:30 compound odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
```

```
W = 0.83177, p-value = 0.007431
```

Result: This data set is normal.

50:50 compound odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.73872, p-value = 0.02332
```

Result: This data set is normal.

50:50 compound odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.79416, p-value = 0.007882
```

Result: This data set is normal.

30:70 compound odour (satiated bees):
Sample size too small to perform test.

30:70 compound odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.8494, p-value = 0.2242
```

Result: This data set is not normal.

Punished odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.80788, p-value = 0.003473
```

Result: This data set is normal.

Punished odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.90769, p-value = 0.01128
```

Result: This data set is normal.

A.2.10 Total Number of Bouts: Statistical Tests

Rewarded odour (100:0)

```
t.test(..., paired=FALSE)
     Welch Two Sample t-test
t = -0.0046458, df = 60.198, p-value = 0.9963
alternative hypothesis:  true difference in means is greater than
0
95 percent confidence interval:
     -0.7733468 0.7697626
sample estimates:
mean of x mean of y
     2.387097 2.388889
```

Result: $p > .05$ for the rewarded odour, hungry bees did not show significantly more proboscis extensions than satiated bees.

70:30 compound odour

```
t.test(..., paired=FALSE)
     Welch Two Sample t-test
t = -2.6881, df = 20.477, p-value = 0.01396
alternative hypothesis:  true difference in means is greater than
0
95 percent confidence interval:
     -1.9966898 -0.2533102
sample estimates:
mean of x mean of y
     1.375 2.500
```

Result: $p < .05$ for the 70:30 compound odour, hungry bees showed significantly more proboscis extensions than satiated bees.

50:50 compound odour

```
t.test(..., paired=FALSE)
     Welch Two Sample t-test
t = -0.30484, df = 6.8666, p-value = 0.7695
alternative hypothesis:  true difference in means is greater than
0
95 percent confidence interval:
     -2.875908 2.221363
sample estimates:
mean of x mean of y
```

```
2.400000 2.727273
```

Result: $p > .05$ for the 50:50 compound odour, hungry bees did not show significantly more proboscis extensions than satiated bees.

30:70 compound odour

```
wilcox.test(..., paired=FALSE, correct=FALSE)
     Wilcoxon rank sum test
W = 2, p-value = 0.1949
```

Result: $p > .05$ for the 30:70 compound odour, hungry bees did not show significantly more proboscis extensions than satiated bees.

Punished odour (0:100)

```
wilcox.test(..., paired=FALSE, correct=FALSE)
     Wilcoxon rank sum test
W = 16, p-value = 0.1231
```

Result: $p > .05$ for the punished odour, hungry bees did not show significantly more proboscis extensions than satiated bees.

A.2.11 Total Time Responding: Shapiro-Wilks Normality Tests

Rewarded odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.90769, p-value = 0.01128
```

Result: This data set is normal.

Rewarded odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.92864, p-value = 0.02277
```

Result: This data set is normal.

70:30 compound odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
```

```
W = 0.67974, p-value = 0.001326
```

Result: This data set is normal.

70:30 compound odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.92242, p-value = 0.1845
```

Result: This data set is not normal.

50:50 compound odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.82294, p-value = 0.123
```

Result: This data set is not normal.

50:50 compound odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.8821, p-value = 0.1106
```

Result: This data set is not normal.

30:70 compound odour (satiated bees):

Sample size too small to perform test.

30:70 compound odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.87007, p-value = 0.2667
```

Result: This data set is not normal.

Punished odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.94213, p-value = 0.6673
```

Result: This data set is not normal.

Punished odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
```

```
W = 0.88526, p-value = 0.0469
```

Result: This data set is normal.

A.2.12 Total Time Responding to Odour: Statistical Tests

Rewarded odour (100:0)

```
t.test(..., paired=FALSE)
     Welch Two Sample t-test, alternative="less"
t = -1.6043, df = 64.367, p-value = 0.05677
alternative hypothesis:  true difference in means is greater than
0
95 percent confidence interval:
     -Inf 0.0654051
sample estimates:
mean of x mean of y
     5.539355 7.164722
```

Result: $p > .05$ for the rewarded odour, hungry bees did not spend significantly more time responding than satiated bees.

70:30 compound odour

```
wilcox.test(..., paired=FALSE, correct=FALSE, alternative="less")
     Wilcoxon rank sum test
W = 34, p-value = 0.03511
```

Result: $p < .05$ for the 70:30 compound odour, hungry bees spent significantly more time responding than satiated bees.

50:50 compound odour

```
wilcox.test(..., paired=FALSE, correct=FALSE, alternative="less")
     Wilcoxon rank sum test
W = 21, p-value = 0.2548
```

Result: $p < .05$ for the 70:30 compound odour, hungry bees did not spend significantly more time responding than satiated bees.

30:70 compound odour

```
wilcox.test(..., paired=FALSE, correct=FALSE, alternative="less")
```

```
     Wilcoxon rank sum test
W = 2, p-value = 0.1905
```

Result: $p < .05$ for the 30:70 compound odour, hungry bees did not spend significantly more time responding than satiated bees.

Punished odour (0:100)

```
wilcox.test(..., paired=FALSE, correct=FALSE)
     Wilcoxon rank sum test
W = 11, p-value = 0.02497
```

Result: $p < .05$ for the punished odour, hungry bees spent significantly more time responding than satiated bees.

A.2.13 Total Time Responding Before Odour Onset: Shapiro-Wilks Normality Tests

Rewarded odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.32314, p-value = 7.331e-11
```

Result: This data set is normal.

Rewarded odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.51805, p-value = 1.025e-09
```

Result: This data set is normal.

70:30 compound odour (satiated bees):
Cannot perform test as all values are the same.

70:30 compound odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.66715, p-value = 7.5e-05
```

Result: This data set is normal.

50:50 compound odour (satiated bees):

```
shapiro.test(...)
```

```
     Shapiro-Wilk normality test
W = 0.58678, p-value = 0.0004144
```

Result: This data set is normal.

50:50 compound odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.71299, p-value = 0.0006876
```

Result: This data set is normal.

30:70 compound odour (satiated bees):

Sample size too small to perform test.

30:70 compound odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.62978, p-value = 0.001241
```

Result: This data set is normal.

Punished odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.87446, p-value = 0.03184
```

Result: This data set is normal.

Punished odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.9109, p-value = 0.01364
```

Result: This data set is normal.

A.2.14 Total Time Responding Before Odour Onset: Statistical Tests

Rewarded odour (100:0)

```
t.test(..., paired=FALSE, alternative="less")
     Welch Two Sample t-test, alternative="less"
t = -1.9206, df = 44.326, p-value = 0.03061
alternative hypothesis:  true difference in means is greater than
```

```
0
95 percent confidence interval:
     -Inf -0.08297658
sample estimates:
mean of x mean of y
     0.1896774 0.8519444
```

Result: $p < .05$ for the rewarded odour, hungry bees spent significantly more time responding before the odour onset than satiated bees.

70:30 compound odour

```
wilcox.test(..., paired=FALSE, correct=FALSE, alternative="less")
     Wilcoxon rank sum test
W = 44, p-value = 0.0423
```

Result: $p < .05$ for the 70:30 compound odour, hungry bees spent significantly more time responding before the odour onset than satiated bees.

50:50 compound odour

```
t.test(..., paired=FALSE, alternative="less")
     Welch Two Sample t-test, alternative="less"
t = 0.16746, df = 5.4361, p-value = 0.5635
alternative hypothesis:  true difference in means is greater than
0
95 percent confidence interval:
     -Inf 3.54847
sample estimates:
mean of x mean of y
     1.624000 1.347273
```

Result: $p > .05$ for the 50:50 compound odour, hungry bees did not spend significantly more time responding before the odour onset than satiated bees.

30:70 compound odour

```
wilcox.test(..., paired=FALSE, correct=FALSE, alternative="less")
     Wilcoxon rank sum test
```

```
W = 44, p-value = 0.0423
```

Result: $p > .05$ for the 30:70 compound odour, hungry bees did not spend significantly more time responding before the odour onset than satiated bees.

Punished odour (0:100)

```
t.test(..., paired=FALSE, alternative="less")
     Welch Two Sample t-test, alternative="less"
t = -1.9156, df = 9.18, p-value = 0.04352
alternative hypothesis:  true difference in means is greater than
0
95 percent confidence interval:
     -Inf -0.1032791
sample estimates:
mean of x mean of y
     0.870 3.155
```

Result: $p < .05$ for the punished odour, hungry bees spent significantly more time responding before the odour onset than satiated bees.

A.2.15 Total Time Responding After Odour Onset: Shapiro-Wilks Normality Tests

Rewarded odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.9109, p-value = 0.01364
```

Result: This data set is normal.

Rewarded odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.96957, p-value = 0.414
```

Result: This data set is not normal.

70:30 compound odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
```

```
W = 0.67974, p-value = 0.001326
```

Result: This data set is normal.

70:30 compound odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.92954, p-value = 0.2398
```

Result: This data set is not normal.

50:50 compound odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.90261, p-value = 0.4245
```

Result: This data set is not normal.

50:50 compound odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.91518, p-value = 0.2805
```

Result: This data set is not normal.

30:70 compound odour (satiated bees):

Sample size too small to perform test.

30:70 compound odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.89867, p-value = 0.4245
```

Result: This data set is not normal.

Punished odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.89867, p-value = 0.4245
```

Result: This data set is not normal.

Punished odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
```

```
W = 0.96026, p-value = 0.6665
```

Result: This data set is not normal.

A.2.16 Total Time Responding After Odour Onset: Statistical Tests

Rewarded odour (100:0)

```
wilcox.test(..., paired=FALSE, correct=FALSE, alternative="less")
     Wilcoxon rank sum test
W = 461.5, p-value = 0.1125
```

Result: $p > .05$ for the rewarded odour, hungry bees did not spend significantly more time responding after the odour onset than satiated bees.

70:30 compound odour

```
wilcox.test(..., paired=FALSE, correct=FALSE, alternative="less")
     Wilcoxon rank sum test
W = 37, p-value = 0.05282
```

Result: $p > .05$ for the 70:30 compound odour, hungry bees did not spend significantly more time responding after the odour onset than satiated bees.

50:50 compound odour

```
wilcox.test(..., paired=FALSE, correct=FALSE, alternative="less")
     Wilcoxon rank sum test
W = 19, p-value = 0.1676
```

Result: $p > .05$ for the 50:50 compound odour, hungry bees did not spend significantly more time responding after the odour onset than satiated bees.

30:70 compound odour

```
wilcox.test(..., paired=FALSE, correct=FALSE, alternative="less")
     Wilcoxon rank sum test
```

```
W = 2, p-value = 0.1905
```

Result: $p > .05$ for the 30:70 compound odour, hungry bees did not spend significantly more time responding after the odour onset than satiated bees.

Punished odour (0:100)

```
wilcox.test(..., paired=FALSE, correct=FALSE, alternative="less")
     Wilcoxon rank sum test
W = 14.5, p-value = 0.04904
```

Result: $p < .05$ for the punished odour, hungry bees spent significantly more time responding after the odour onset than satiated bees.

A.2.17 Number of Bouts Above the Plane: Shapiro-Wilks Normality Tests

Rewarded odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.72593, p-value = 2.872e-06
```

Result: This data set is normal.

Rewarded odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.7913, p-value = 1.099e-05
```

Result: This data set is normal.

70:30 compound odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.6412, p-value = 0.0004791
```

Result: This data set is normal.

70:30 compound odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
```

```
W = 0.81401, p-value = 0.004205
```

Result: This data set is normal.

50:50 compound odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.55218, p-value = 0.000131
```

Result: This data set is normal.

50:50 compound odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.34499, p-value = 2.243e-08
```

Result: This data set is normal.

30:70 compound odour (satiated bees):
Sample size too small to perform test.

30:70 compound odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.89867, p-value = 0.4245
```

Result: This data set is not normal.

Punished odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.68403, p-value = 0.00647
```

Result: This data set is normal.

Punished odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.45449, p-value = 9.106e-07
```

Result: This data set is normal.

A.2.18 Number of Bouts Above Plane of Head: Statistical Tests

Rewarded odour (100:0)

```
t.test(..., paired=FALSE, alternative="less")
     Welch Two Sample t-test, alternative="less"
t = -0.96385, df = 64.564, p-value = 0.1694
alternative hypothesis:  true difference in means is greater than
0
95 percent confidence interval:
     -Inf 0.1441808
sample estimates:
mean of x mean of y
     0.5806452 0.7777778
```

Result: $p > .05$ for the rewarded odour, hungry bees did not show significantly more proboscis extensions above the plane of the head than satiated bees.

70:30 compound odour

```
t.test(..., paired=FALSE, alternative="less")
     Welch Two Sample t-test, alternative="less"
t = -1.9496, df = 18.837, p-value = 0.03314
alternative hypothesis:  true difference in means is greater than
0
95 percent confidence interval:
     -Inf -0.05634282
sample estimates:
mean of x mean of y
     0.375 0.875
```

Result: $p < .05$ for the 70:30 compound odour, hungry bees showed significantly more proboscis extensions above the plane of the head than satiated bees.

50:50 compound odour

```
t.test(..., paired=FALSE, alternative="less")
     Welch Two Sample t-test, alternative="less"
t = 2.7116, df = 23.885, p-value = 0.9939
alternative hypothesis:  true difference in means is greater than
0
95 percent confidence interval:
```

```
     -Inf 1.130629
sample estimates:
mean of x mean of y
     0.8750000 0.1818182
```

Result: $p > .05$ for the 50:50 compound odour, hungry bees did not show significantly more proboscis extensions above the plane of the head than satiated bees.

30:70 compound odour

```
wilcox.test(..., paired=FALSE, correct=FALSE, alternative="less")
     Wilcoxon rank sum test
W = 3, p-value = 0.1636
```

Result: $p > .05$ for the 30:70 compound odour, hungry bees did not show significantly more proboscis extensions above the plane of the head than satiated bees.

Punished odour (0:100)

```
wilcox.test(..., paired=FALSE, correct=FALSE, alternative="less")
     Wilcoxon rank sum test
W = 24, p-value = 0.139
```

Result: $p > .05$ for the punished odour, hungry bees did not show significantly more proboscis extensions above the plane of the head than satiated bees.

A.2.19 Number of Bouts At the Plane: Shapiro-Wilks Normality Tests

Rewarded odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.86153, p-value = 0.0009008
```

Result: This data set is normal.

Rewarded odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
```

```
W = 0.89081, p-value = 0.001932
```

Result: This data set is normal.

70:30 compound odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.84891, p-value = 0.09288
```

Result: This data set is not normal.

70:30 compound odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.81087, p-value = 0.003811
```

Result: This data set is normal.

50:50 compound odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.85991, p-value = 0.2279
```

Result: This data set is not normal.

50:50 compound odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.76023, p-value = 0.002828
```

Result: This data set is normal.

30:70 compound odour (satiated bees):

Sample size too small to perform test.

30:70 compound odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.96086, p-value = 0.814
```

Result: This data set is not normal.

Punished odour (satiated bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
```

```
W = 0.8494, p-value = 0.2242
```

Result: This data set is not normal.

Punished odour (hungry bees):

```
shapiro.test(...)
     Shapiro-Wilk normality test
W = 0.77026, p-value = 0.001126
```

Result: This data set is normal.

A.2.20 Number of Bouts At Plane of Head: Statistical Tests

Rewarded odour (100:0)

```
t.test(..., paired=FALSE, alternative="less")
     Welch Two Sample t-test, alternative="less"
t = 0.53941, df = 60.517, p-value = 0.7042
alternative hypothesis:  true difference in means is greater than
0
95 percent confidence interval:
     -Inf 0.8002629
sample estimates:
mean of x mean of y
     1.806452 1.611111
```

Result: $p > .05$ for the rewarded odour, hungry bees did not show significantly more proboscis extensions at the plane of the head than satiated bees.

70:30 compound odour

```
wilcox.test(..., paired=FALSE, correct=FALSE, alternative="less")
     Wilcoxon rank sum test
W = 54, p-value = 0.2551
```

Result: $p > .05$ for the 70:30 compound odour, hungry bees did not show significantly more proboscis extensions at the plane of the head than satiated bees.

50:50 compound odour

```
wilcox.test(..., paired=FALSE, correct=FALSE, alternative="less")
     Wilcoxon rank sum test
```

```
W = 22.5, p-value = 0.2756
```

Result: $p > .05$ for the 50:50 compound odour, hungry bees did not show significantly more proboscis extensions at the plane of the head than satiated bees.

30:70 compound odour

```
wilcox.test(..., paired=FALSE, correct=FALSE, alternative="less")
     Wilcoxon rank sum test
W = 3, p-value = 0.2085
```

Result: $p > .05$ for the 30:70 compound odour, hungry bees did not show significantly more proboscis extensions at the plane of the head than satiated bees.

Punished odour (0:100)

```
wilcox.test(..., paired=FALSE, correct=FALSE, alternative="less")
     Wilcoxon rank sum test
W = 19, p-value = 0.1045
```

Result: $p > .05$ for the punished odour, hungry bees did not show significantly more proboscis extensions at the plane of the head than satiated bees.

Appendix B

PER Training Data Sheets

B.1 First Experiment

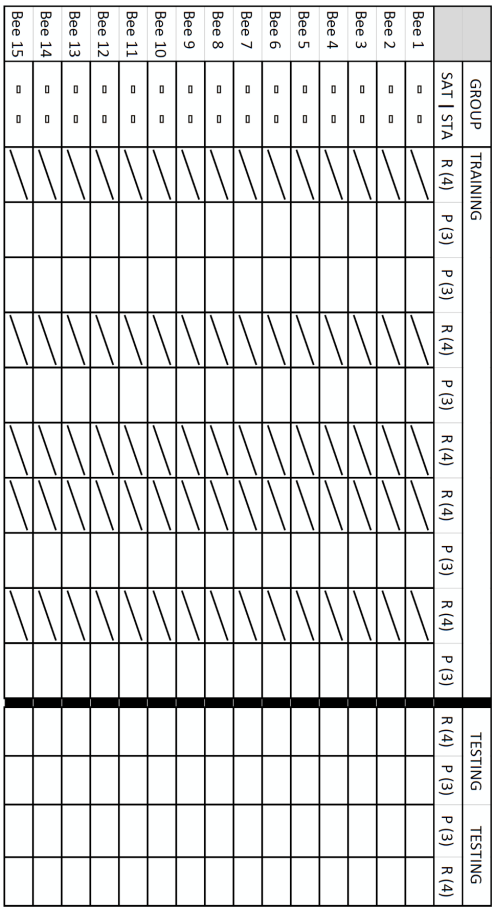

Figure B.1: Data sheet used for the first experiment.

B.2 Second Experiment

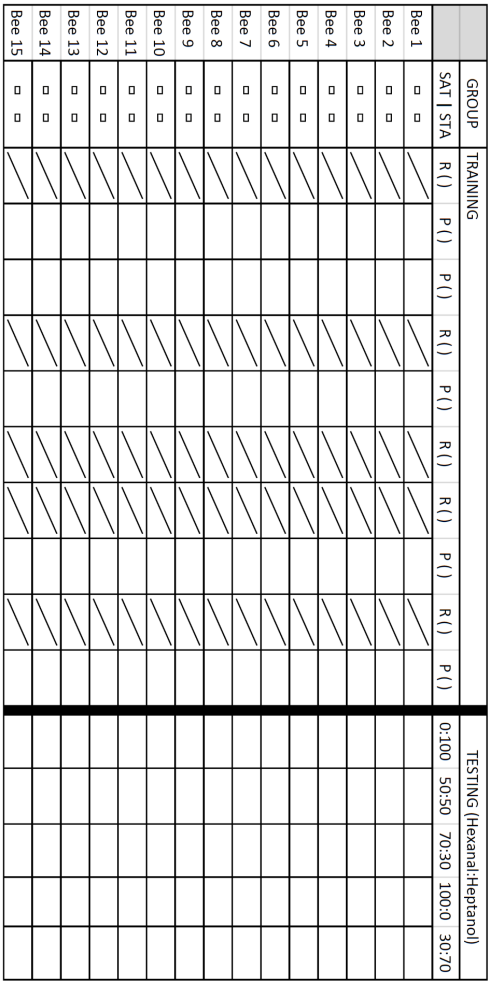

Figure B.2: Data sheet used for the second experiment.

Bibliography

Abel, R., Rybak, J., and Menzel, R. Structure and response patterns of olfactory interneurons in the honeybee, Apis mellifera. *Journal of Comparative Neurology*, 437(3):363–383, 2001.

Akalal, D.-B. G., Wilson, C. F., Zong, L., Tanaka, N. K., Ito, K., and Davis, R. L. Roles for Drosophila mushroom body neurons in olfactory learning and memory. *Learning & Memory*, 13(5):659–668, 2006.

Albantakis, L. and Deco, G. The encoding of alternatives in multiple-choice decision-making. *BMC Neuroscience*, 10(Suppl 1):P166, 2009.

Arbib, M. Introducing the Neuron. In *Handbook of Brain Theory and Neural Networks*, pages 4–11. MIT Press, 1995.

Arevian, A. C., Kapoor, V., and Urban, N. N. Activity-dependent gating of lateral inhibition in the mouse olfactory bulb. *Nature Neuroscience*, 11(1): 80–87, 2008.

Arnott, G. and Elwood, R. W. Information gathering and decision making about resource value in animal contests. *Animal Behaviour*, 76(3):529–542, 2008.

Avarguès-Weber, A. and Mota, T. Advances and limitations of visual conditioning protocols in harnessed bees. *Journal of Physiology-Paris*, 2016.

Avarguès-Weber, A., de Brito Sanchez, M. G., Giurfa, M., and Dyer, A. G.

Aversive reinforcement improves visual discrimination learning in free-flying honeybees. *PLoS One*, 5(10):e15370, 2010.

Avarguès-Weber, A., Dyer, A. G., Combe, M., and Giurfa, M. Simultaneous mastering of two abstract concepts by the miniature brain of bees. *Proceedings of the National Academy of Sciences*, 109(19):7481–7486, 2012.

Avila, C., Tyndale, E., and Kuzirian, A. Feeding behavior and growth of Hermissenda crassicornis (Mollusca: Nudibranchia) in the laboratory. *Marine & Freshwater Behaviour & Phy*, 31(1):1–19, 1998.

Azevedo, F. A., Carvalho, L. R., Grinberg, L. T., Farfel, J. M., Ferretti, R. E., Leite, R. E., Lent, R., Herculano-Houzel, S., et al. Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain. *Journal of Comparative Neurology*, 513(5):532–541, 2009.

Balamurali, G., Somanathan, H., and De Ibarra, N. H. Motion cues improve the performance of harnessed bees in a colour learning task. *Journal of Comparative Physiology A*, 201(5):505–511, 2015.

Balling, A., Technau, G., and Heisenberg, M. Are the structural changes in adult Drosophila mushroom bodies memory traces? Studies on biochemical learning mutants. *Journal of Neurogenetics*, 4(1):65–73, 1987.

Bargmann, C. I. Beyond the connectome: how neuromodulators shape neural circuits. *Bioessays*, 34(6):458–465, 2012.

Barron, A. B., Gurney, K. N., Meah, L. F., Vasilaki, E., and Marshall, J. A. Decision-making and action selection in insects: inspiration from vertebrate-based theories. *Frontiers in Behavioral Neuroscience*, 9, 2015.

Barth, M. and Heisenberg, M. Vision affects mushroom bodies and central complex in Drosophila melanogaster. *Learning & Memory*, 4(2):219–229, 1997.

Bausenwein, B., Müller, N. R., and Heisenberg, M. Behavior-dependent activity labeling in the central complex of Drosophila during controlled visual stimulation. *Journal of Comparative Neurology*, 340(2):255–268, 1994.

Beggs, K. T., Hamilton, I. S., Kurshan, P. T., Mustard, J. A., and Mercer, A. R. Characterization of a d2-like dopamine receptor (am dop3) in honey bee, Apis mellifera. *Insect Biochemistry and Molecular Biology*, 35(8):873–882, 2005.

Bell, W. J. Searching behavior patterns in insects. *Annual review of Entomology*, 35(1):447–467, 1990.

Bitterman, M., Menzel, R., Fietz, A., and Schäfer, S. Classical conditioning of proboscis extension in honeybees (Apis mellifera). *Journal of Comparative Psychology*, 97(2):107, 1983.

Blakemore, C. and Tobin, E. A. Lateral inhibition between orientation detectors in the cat's visual cortex. *Experimental Brain Research*, 15(4):439–440, 1972.

Bogacz, R. Optimal decision making theories. *Handbook of Reward and Decision Making*, pages 375–397, 2009.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J. D. The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, 113(4):700, 2006.

Bogacz, R., Usher, M., Zhang, J., and McClelland, J. L. Extending a biologically inspired model of choice: multi-alternatives, nonlinearity and value-based multidimensional choice. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1485):1655–1670, 2007.

Bogacz, R., Hu, P. T., Holmes, P. J., and Cohen, J. D. Do humans produce the speed–accuracy trade-off that maximizes reward rate? *The Quarterly Journal of Experimental Psychology*, 63(5):863–891, 2010a.

Bogacz, R., Wagenmakers, E.-J., Forstmann, B. U., and Nieuwenhuis, S. The neural basis of the speed–accuracy tradeoff. *Trends in Neurosciences*, 33(1): 10–16, 2010b.

Bono, M. d. and Villu Maricq, A. Neuronal substrates of complex behaviors in C. elegans. *Annual Review of Neuroscience*, 28:451–501, 2005.

Bose, T., Reina, A., and Marshall, J. A. Collective decision-making. *Current Opinion in Behavioral Sciences*, 16:30–34, 2017.

Brandt, R., Rohlfing, T., Rybak, J., Krofczik, S., Maye, A., Westerhoff, M., Hege, H.-C., and Menzel, R. Three-dimensional average-shape atlas of the honeybee brain and its applications. *Journal of Comparative Neurology*, 492 (1):1–19, 2005.

Britten, K. H., Shadlen, M. N., Newsome, W. T., and Movshon, J. A. The analysis of visual motion: a comparison of neuronal and psychophysical performance. *The Journal of Neuroscience*, 12(12):4745–4765, 1992.

Brown, E. and Holmes, P. Modeling a simple choice task: stochastic dynamics of mutually inhibitory neural groups. *Stochastics and Dynamics*, 1(02):159–191, 2001.

Brown, S. D. and Heathcote, A. The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, 57(3): 153–178, 2008.

Burkitt, A. N. A review of the integrate-and-fire neuron model: I. homogeneous synaptic input. *Biological Cybernetics*, 95(1):1–19, 2006.

Burns, J. G. Impulsive bees forage better: the advantage of quick, sometimes inaccurate foraging decisions. *Animal Behaviour*, 70(6):e1–e5, 2005.

Burns, J. G. and Dyer, A. G. Diversity of speed-accuracy strategies benefits social insects. *Current Biology*, 18(20):R953–R954, 2008.

Busemeyer, J. R. and Johnson, J. G. Computational models of decision making. *Blackwell Handbook of Judgment and Decision Making*, pages 133–154, 2004.

Busemeyer, J. R. and Townsend, J. T. Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100(3):432, 1993.

Camazine, S. and Sneyd, J. A model of collective nectar source selection by honey bees: self-organization through simple rules. *Journal of Theoretical Biology*, 149(4):547–571, 1991.

Chittka, L., Thomson, J. D., and Waser, N. M. Flower constancy, insect psychology, and plant evolution. *Naturwissenschaften*, 86(8):361–377, 1999.

Chittka, L., Dyer, A. G., Bock, F., and Dornhaus, A. Psychophysics: bees trade off foraging speed for accuracy. *Nature*, 424(6947):388–388, 2003.

Chittka, L., Skorupski, P., and Raine, N. E. Speed–accuracy tradeoffs in animal decision making. *Trends in Ecology & Evolution*, 24(7):400–407, 2009.

Cisek, P., Puskas, G. A., and El-Murr, S. Decisions in changing conditions: the urgency-gating model. *The Journal of Neuroscience*, 29(37):11560–11571, 2009.

Cleland, T. A. and Linster, C. Computation in the olfactory system. *Chemical Senses*, 30(9):801–813, 2005.

Cohen-Gilbert, J., Killgore, W., White, C., Schwab, Z., Crowley, D., Covell, M., Sneider, J., and Silveri, M. Differential influence of safe versus threatening facial expressions on decision-making during an inhibitory control task in adolescence and adulthood. *Developmental Science*, 17(2):212–223, 2014.

Collett, T. and Kelber, A. The retrieval of visuo-spatial memories by honeybees. *Journal of Comparative Physiology A*, 163(1):145–150, 1988.

Conradt, L. and Roper, T. J. Group decision-making in animals. *Nature*, 421 (6919):155–158, 2003.

Conradt, L. and Roper, T. J. Consensus decision making in animals. *Trends in Ecology & Evolution*, 20(8):449–456, 2005.

Cope, A. J., Sabo, C., Gurney, K., Vasilaki, E., and Marshall, J. A. A model for an angular velocity-tuned motion detector accounting for deviations in the corridor-centering response of the bee. *PLoS Computational Biology*, 12 (5):e1004887, 2016.

Couvillon, M. The dance legacy of Karl von Frisch. *Insectes Sociaux*, 59(3): 297–306, 2012.

Couzin, I. D., Ioannou, C. C., Demirel, G., Gross, T., Torney, C. J., Hartnett, A., Conradt, L., Levin, S. A., and Leonard, N. E. Uninformed individuals promote democratic consensus in animal groups. *Science*, 334(6062):1578–1580, 2011.

Dayan, P. and Abbott, L. F. *Theoretical Neuroscience*, volume 806. Cambridge, MA: MIT Press, 2001.

de Brito Sanchez, M. G., Chen, C., Li, J., Liu, F., Gauthier, M., and Giurfa, M. Behavioral studies on tarsal gustation in honeybees: sucrose responsiveness and sucrose-mediated olfactory conditioning. *Journal of Comparative Physiology A*, 194(10):861–869, 2008.

De Ibarra, N. H., Vorobyev, M., and Menzel, R. Mechanisms, functions and ecology of colour vision in the honeybee. *Journal of Comparative Physiology A*, 200(6):411–433, 2014.

De Martino, B., Fleming, S. M., Garrett, N., and Dolan, R. J. Confidence in value-based choice. *Nature Neuroscience*, 16(1):105–110, 2013.

de Vries, S. E. and Clandinin, T. R. Loom-sensitive neurons link computation

to action in the *Drosophila* visual system. *Current Biology*, 22(5):353–362, 2012.

Deisig, N., Lachnit, H., Sandoz, J.-C., Lober, K., and Giurfa, M. A modified version of the unique cue theory accounts for olfactory compound processing in honeybees. *Learning & Memory*, 10(3):199–208, 2003.

Ditterich, J., Mazurek, M. E., and Shadlen, M. N. Microstimulation of visual cortex affects the speed of perceptual decisions. *Nature Neuroscience*, 6(8): 891–898, 2003.

Ditterich, J. et al. A comparison between mechanisms of multi-alternative perceptual decision making: ability to explain human behavior, predictions for neurophysiology, and relationship with decision theory. *Frontiers in Neuroscience*, 4, 2010.

Dobrin, S. E. and Fahrbach, S. E. Visual associative learning in restrained honey bees with intact antennae. *PLoS One*, 7(6):e37666, 2012.

Doya, K., Ishii, S., Pouget, A., and Rao, R. P. N. *Bayesian brain: Probabilistic approaches to neural coding*. Cambridge, MA: MIT Press, 2007.

Drezner-Levy, T., Smith, B. H., and Shafir, S. The effect of foraging specialization on various learning tasks in the honey bee (Apis mellifera). *Behavioral Ecology and Sociobiology*, 64(1):135, 2009.

Dujardin, F. Mémoire sur le système nerveux des insectes. *Annales des Sciences Naturelles, Zoologie.*, 14(1850):547–560, 1850.

Dukas, R. and Clark, C. W. Sustained vigilance and animal performance. *Animal Behaviour*, 49(5):1259–1267, 1995.

Durst, C., Eichmüller, S., and Menzel, R. Development and experience lead to increased volume of subcompartments of the honeybee mushroom body. *Behavioral and Neural Biology*, 62(3):259–263, 1994.

Dussutour, A. and Simpson, S. J. Communal nutrition in ants. *Current Biology*, 19(9):740–744, 2009.

Dyer, A. and Chittka, L. Biological significance of distinguishing between similar colours in spectrally variable illumination: bumblebees (Bombus terrestris) as a case study. *Journal of Comparative Physiology A*, 190(2): 105–114, 2004a.

Dyer, A. G. and Chittka, L. Bumblebees (Bombus terrestris) sacrifice foraging speed to solve difficult colour discrimination tasks. *Journal of Comparative Physiology A*, 190(9):759–763, 2004b.

Dyer, A. G., Spaethe, J., and Prack, S. Comparative psychophysics of bumblebee and honeybee colour discrimination and object detection. *Journal of Comparative Physiology A*, 194(7):617–627, 2008.

Erber, J. Neural correlates of learning in the honeybee. *Trends in Neurosciences*, 4:270–273, 1981.

Erber, J. Tactile antennal learning in the honey bee. In *Honeybee Neurobiology and Behavior*, pages 439–455. Springer, 2012.

Esch, H. Beiträge zum problem der entfernungsweisung in den schwänzeltänzen der honigbiene. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, 48(5):534–546, 1964.

Faber, T., Joerges, J., and Menzel, R. Associative learning modifies neural representations of odors in the insect brain. *Nature Neuroscience*, 2(1): 74–78, 1999.

Fechner, G. Elements of psychophysics. Vol. I. New York: Holt, Rinehart and Winston (Original work published 1860), 1966.

Felsenberg, J., Gehring, K. B., Antemann, V., and Eisenhardt, D. Behavioural pharmacology in classical conditioning of the proboscis extension response

in honeybees (Apis mellifera). *JoVE (Journal of Visualized Experiments)*, (47):e2282–e2282, 2011.

Feng, S., Holmes, P., Rorie, A., and Newsome, W. T. Can monkeys choose optimally when faced with noisy stimuli and unequal rewards? *PLoS Computational Biology*, 5(2):e1000284, 2009.

Fernandez, P. C., Locatelli, F. F., Person-Rennell, N., Deleo, G., and Smith, B. H. Associative conditioning tunes transient dynamics of early olfactory processing. *The Journal of Neuroscience*, 29(33):10191–10202, 2009.

FitzHugh, R. Impulses and physiological states in theoretical models of nerve membrane. *Biophysical Journal*, 1(6):445–466, 1961.

Franceschini, N., Ruffier, F., and Serres, J. A bio-inspired flying robot sheds light on insect piloting abilities. *Current Biology*, 17(4):329–335, 2007.

Franks, N. R., Dornhaus, A., Fitzsimmons, J. P., and Stevens, M. Speed versus accuracy in collective decision making. *Proceedings of the Royal Society of London B: Biological Sciences*, 270(1532):2457–2463, 2003.

Free, J. The flower constancy of honeybees. *Journal of Animal Ecology*, 32 (1):119–131, 1963.

Free, J. The flower constancy of bumblebees. *The Journal of Animal Ecology*, pages 395–402, 1970.

Friedrich, A., Thomas, U., and Müller, U. Learning at different satiation levels reveals parallel functions for the cAMP–protein kinase A cascade in formation of long-term memory. *Journal of Neuroscience*, 24(18):4460–4468, 2004.

Frohlich, B., Riederer, M., and Tautz, J. Comb-wax discrimination by honeybees tested with the proboscis extension reflex. *Journal of Experimental Biology*, 203(10):1581–1587, 2000.

Frost, E. H., Shutler, D., and Hillier, N. K. The proboscis extension reflex to evaluate learning and memory in honeybees (Apis mellifera): some caveats. *Naturwissenschaften*, 99(9):677–686, 2012.

Galizia, C. G. Olfactory coding in the insect brain: data and conjectures. *European Journal of Neuroscience*, 39(11):1784–1795, 2014.

Galizia, G., Eisenhardt, D., and Giurfa, M. Honeybee neurobiology and behavior: a tribute to randolf menzel. Springer Science & Business Media, 2012.

Gauthier, M. and Grünewald, B. Neurotransmitter systems in the honey bee brain: functions in learning and memory. In *Honeybee Neurobiology and Behavior*, pages 155–169. Springer, 2012.

Getting, P. A. and Dekin, M. S. Tritonia swimming. In *Model Neural Networks and Behavior*, pages 3–20. Springer, 1985.

Giurfa, M. Cognitive neuroethology: dissecting non-elemental learning in a honeybee brain. *Current Opinion in Neurobiology*, 13(6):726–735, 2003a.

Giurfa, M. Cognitive neuroethology: dissecting non-elemental learning in a honeybee brain. *Current Opinion in Neurobiology*, 13(6):726–735, 2003b.

Giurfa, M. Conditioning procedure and color discrimination in the honeybee Apis mellifera. *Naturwissenschaften*, 91(5):228–231, 2004.

Giurfa, M. Behavioral and neural analysis of associative learning in the honeybee: a taste from the magic well. *Journal of Comparative Physiology A*, 193(8):801–824, 2007.

Giurfa, M. Cognition with few neurons: higher-order learning in insects. *Trends in Neurosciences*, 36(5):285–294, 2013.

Giurfa, M. and Malun, D. Associative mechanosensory conditioning of the proboscis extension reflex in honeybees. *Learning & Memory*, 11(3):294–302, 2004.

Giurfa, M. and Sandoz, J.-C. Invertebrate learning and memory: fifty years of olfactory conditioning of the proboscis extension response in honeybees. *Learning & Memory*, 19(2):54–66, 2012.

Giurfa, M., Nunez, J., Chittka, L., and Menzel, R. Colour preferences of flower-naive honeybees. *Journal of Comparative Physiology A*, 177(3):247–259, 1995.

Giurfa, M., Eichmann, B., and Menzel, R. Symmetry perception in an insect. *Nature*, 382(6590):458–461, 1996.

Giurfa, M., Zhang, S., Jenett, A., Menzel, R., and Srinivasan, M. V. The concepts of 'sameness' and 'difference' in an insect. *Nature*, 410(6831):930–933, 2001.

Giurfa, M., Schubert, M., Reisenman, C., Gerber, B., and Lachnit, H. The effect of cumulative experience on the use of elemental and configural visual discrimination strategies in honeybees. *Behavioural Brain Research*, 145(1): 161–169, 2003.

Godin, J.-G. J. and Crossman, S. L. Hunger-dependent predator inspection and foraging behaviours in the threespine stickleback (Gasterosteus aculeatus) under predation risk. *Behavioral Ecology and Sociobiology*, 34 (5):359–366, 1994.

Gold, J. I. and Shadlen, M. N. Neural computations that underlie decisions about sensory stimuli. *Trends in Cognitive Sciences*, 5(1):10–16, 2001.

Gold, J. I. and Shadlen, M. N. The neural basis of decision making. *Annual Review of Neuroscience*, 30:535–574, 2007.

Goldstein, E. B. and Brockmole, J. *Sensation and Perception*. Cengage Learning, 2016.

Gould, J. L. The biology of learning. *Annual Review of Psychology*, 37(1): 163–192, 1986.

Gould, J. L. Honey bee cognition. *Cognition*, 37(1):83–103, 1990.

Grant, V. The flower constancy of bees. *The Botanical Review*, 16(7):379–398, 1950.

Grünewald, B. Physiological properties and response modulations of mushroom body feedback neurons during olfactory learning in the honeybee, Apis mellifera. *Journal of Comparative Physiology A*, 185(6):565–576, 1999.

Guerrieri, F., Schubert, M., Sandoz, J.-C., and Giurfa, M. Perceptual and neural olfactory similarity in honeybees. *PLoS Biology*, 3(4):e60, 2005.

Guerrieri, F. J. and d'Ettorre, P. Associative learning in ants: conditioning of the maxilla-labium extension response in Camponotus aethiops. *Journal of Insect Physiology*, 56(1):88–92, 2010.

Gupta, N. and Stopfer, M. Functional analysis of a higher olfactory center, the lateral horn. *The Journal of Neuroscience*, 32(24):8138–8148, 2012.

Gurney, K., Prescott, T. J., and Redgrave, P. A computational model of action selection in the basal ganglia. I. a new functional anatomy. *Biological Cybernetics*, 84(6):401–410, 2001a.

Gurney, K., Prescott, T. J., and Redgrave, P. A computational model of action selection in the basal ganglia. II. analysis and simulation of behaviour. *Biological Cybernetics*, 84(6):411–423, 2001b.

Hall, J. C. The mating of a fly. *Science*, 264(5166):1702–1714, 1994.

Hammer, M. An identified neuron mediates the unconditioned stimulus in associative olfactory learning in honeybees. *Nature*, 366:59, 1993a.

Hammer, M. An identified neuron mediates the unconditioned stimulus in associative olfactory learning in honeybees. *Nature*, 366(6450):59, 1993b.

Hammer, M. and Menzel, R. Learning and memory in the honeybee. *The Journal of Neuroscience*, 15(3):1617–1630, 1995.

Hammer, M. and Menzel, R. Multiple sites of associative odor learning as revealed by local brain microinjections of octopamine in honeybees. *Learning & Memory*, 5(1):146–156, 1998.

Hampton, A. N., Bossaerts, P., and O'doherty, J. P. The role of the ventromedial prefrontal cortex in abstract state-based inference during decision making in humans. *The Journal of Neuroscience*, 26(32):8360–8367, 2006.

Hanesch, U., Fischbach, K.-F., and Heisenberg, M. Neuronal architecture of the central complex in Drosophila melanogaster. *Cell and Tissue Research*, 257(2):343–366, 1989.

Hanks, T. D., Mazurek, M. E., Kiani, R., Hopp, E., and Shadlen, M. N. Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *Journal of Neuroscience*, 31(17):6339–6352, 2011.

Hansson, B. S. *Insect olfaction*. Springer Science & Business Media, 1999.

Harris-Warrick, R. M. and Marder, E. Modulation of neural networks for behavior. *Annual Review of Neuroscience*, 14(1):39–57, 1991.

Harrison, R. R. and Koch, C. A robust analog VLSI motion sensor based on the visual system of the fly. *Autonomous Robots*, 7(3):211–224, 1999.

Hartline, H. K., Wagner, H. G., and Ratliff, F. Inhibition in the eye of Limulus. *The Journal of General Physiology*, 39(5):651–673, 1956.

Hawkins, G. E., Forstmann, B. U., Wagenmakers, E.-J., Ratcliff, R., and Brown, S. D. Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. *The Journal of Neuroscience*, 35(6): 2476–2484, 2015.

Hazlett, B., Rubenstein, D., and Rittschof, D. Starvation, energy reserves, and aggression in the crayfish Orconectes virilis (Hagen, 1870)(Decapoda, Cambaridae). *Crustaceana*, pages 11–16, 1975.

Heimbeck, G., Bugnon, V., Gendre, N., Keller, A., and Stocker, R. F. A central neural circuit for experience-independent olfactory and courtship behavior in Drosophila melanogaster. *Proceedings of the National Academy of Sciences*, 98(26):15336–15341, 2001.

Heisenberg, M. What do the mushroom bodies do for the insect brain? An introduction. *Learning & Memory*, 5(1):1–10, 1998.

Heisenberg, M., Borst, A., Wagner, S., and Byers, D. Drosophila mushroom body mutants are deficient in olfactory learning: Research papers. *Journal of Neurogenetics*, 2(1):1–30, 1985.

Heitz, R. P. The speed-accuracy tradeoff: history, physiology, methodology, and behavior. *Frontiers in Neuroscience*, 8:150, 2014.

Heitz, R. P. and Schall, J. D. Neural mechanisms of speed-accuracy tradeoff. *Neuron*, 76(3):616–628, 2012.

Helm, B. R., Slater, G., Rajamohan, A., Yocum, G. D., Greenlee, K. J., and Bowsher, J. H. The geometric framework for nutrition reveals interactions between protein and carbohydrate during larval growth in honey bees. *Biology Open*, pages bio–022582, 2017.

Henmon, V. A. C. The relation of the time of a judgment to its accuracy. *Psychological Review*, 18(3):186, 1911.

Hensch, T. K., Fagiolini, M., Mataga, N., Stryker, M. P., Baekkeskov, S., and Kash, S. F. Local gaba circuit control of experience-dependent plasticity in developing visual cortex. *Science*, 282(5393):1504–1508, 1998.

Hergarden, A. C., Tayler, T. D., and Anderson, D. J. Allatostatin-a neurons inhibit feeding behavior in adult Drosophila. *Proceedings of the National Academy of Sciences*, 109(10):3967–3972, 2012.

Herz, D. M., Zavala, B. A., Bogacz, R., and Brown, P. Neural correlates of

decision thresholds in the human subthalamic nucleus. *Current Biology*, 26 (7):916–920, 2016.

Hodgkin, A. L. and Huxley, A. F. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4):500, 1952.

Homberg, U. Interneurones of the central complex in the bee brain (Apis mellifera, L.). *Journal of Insect Physiology*, 31(3):251–264, 1985.

Hori, S., Takeuchi, H., Arikawa, K., Kinoshita, M., Ichikawa, N., Sasaki, M., and Kubo, T. Associative visual learning, color discrimination, and chromatic adaptation in the harnessed honeybee Apis mellifera L. *Journal of Comparative Physiology A*, 192(7):691–700, 2006.

Hori, S., Takeuchi, H., and Kubo, T. Associative learning and discrimination of motion cues in the harnessed honeybee Apis mellifera L. *Journal of Comparative Physiology A*, 193(8):825–833, 2007.

Horridge, G. and Zhang, S. Pattern vision in honeybees (Apis mellifera): Flower-like patterns with no predominant orientation. *Journal of Insect Physiology*, 41(8):681–688, 1995.

Horridge, G., Zhang, S., and Lehrer, M. Bees can combine range and visual angle to estimate absolute size. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 337(1279):49–57, 1992.

Hourcade, B., Perisse, E., Devaud, J.-M., and Sandoz, J.-C. Long-term memory shapes the primary olfactory center of an insect brain. *Learning & Memory*, 16(10):607–615, 2009.

Houston, A. and Sumida, B. A positive feedback model for switching between two activities. *Animal Behaviour*, 33(1):315–325, 1985.

Houston, A. I., Higginson, A. D., and McNamara, J. M. Optimal foraging

for multiple nutrients in an unpredictable environment. *Ecology Letters*, 14 (11):1101–1107, 2011.

Huang, Z.-Y. and Otis, G. W. Inspection and feeding of larvae by worker honey bees (Hymenoptera: Apidae): effect of starvation and food quantity. *Journal of Insect Behavior*, 4(3):305–317, 1991.

Huk, A. C. and Shadlen, M. N. Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *The Journal of Neuroscience*, 25(45):10420–10436, 2005.

Humphries, M. D., Stewart, R. D., and Gurney, K. N. A physiologically plausible model of action selection and oscillatory activity in the basal ganglia. *The Journal of Neuroscience*, 26(50):12921–12942, 2006.

Ings, T. C. and Chittka, L. Speed-accuracy tradeoffs and false alarms in bee responses to cryptic predators. *Current Biology*, 18(19):1520–1524, 2008.

Ito, K., Suzuki, K., Estes, P., Ramaswami, M., Yamamoto, D., and Strausfeld, N. J. The organization of extrinsic neurons and their implications in the functional roles of the mushroom bodies in Drosophila melanogaster meigen. *Learning & Memory*, 5(1):52–77, 1998.

Ito, K., Shinomiya, K., Ito, M., Armstrong, J. D., Boyan, G., Hartenstein, V., Harzsch, S., Heisenberg, M., Homberg, U., Jenett, A., et al. A systematic nomenclature for the insect brain. *Neuron*, 81(4):755–765, 2014.

Izhikevich, E. M. Simple model of spiking neurons. *IEEE Transactions on Neural Networks*, 14(6):1569–1572, 2003.

Jost, J. T., Nam, H. H., Amodio, D. M., and Van Bavel, J. J. Political neuroscience: The beginning of a beautiful friendship. *Political Psychology*, 35(S1):3–42, 2014.

Kanai, R., Feilden, T., Firth, C., and Rees, G. Political orientations are

correlated with brain structure in young adults. *Current Biology*, 21(8): 677–680, 2011.

Kenyon, F. The brain of the bee. a preliminary contribution to the morphology of the nervous system of the Arthropoda. *Journal of Comparative Neurology*, 6(3):133–210, 1896.

Khoury, D. S., Myerscough, M. R., and Barron, A. B. A quantitative model of honey bee colony population dynamics. *PloS One*, 6(4):e18491, 2011.

Kien, J. The initiation and maintenance of walking in the locust: an alternative to the command concept. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 219(1215):137–174, 1983.

Kim, J.-N. and Shadlen, M. N. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nature Neuroscience*, 2(2):176–185, 1999.

Kirschfeld, K. and Lutz, B. Lateral inhibition in the com pound eye of the fly, Musca. *Zeitschrift für Naturforschung C*, 29(1-2):95–97, 1974.

Kirschner, S., Kleineidam, C. J., Zube, C., Rybak, J., Grünewald, B., and Rössler, W. Dual olfactory pathway in the honeybee, Apis mellifera. *Journal of Comparative Neurology*, 499(6):933–952, 2006.

Krajbich, I. and Rangel, A. Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences*, 108(33):13852–13857, 2011.

Krajbich, I., Armel, C., and Rangel, A. Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience*, 13(10): 1292–1298, 2010.

Krajbich, I., Lu, D., Camerer, C., and Rangel, A. The attentional drift-diffusion model extends to simple purchasing decisions. *Frontiers in Psychology*, 3, 2012.

Kurshan, P. T., Hamilton, I. S., Mustard, J. A., and Mercer, A. R. Developmental changes in expression patterns of two dopamine receptor genes in mushroom bodies of the honeybee, Apis mellifera. *Journal of Comparative Neurology*, 466(1):91–103, 2003.

Kuwabara, M. Bildung des bedingten reflexes von pavlovs typus bei der honigbiene, Apis mellifica (mit 1 textabbildung). *Journal of the Faculty of Science, Hokkaido University Series VI. Zoology*, 13(1-4):458–464, 1957.

Laming, D. R. J. Information theory of choice-reaction times. Oxford, England: Academic Press, 1968.

Latty, T. and Beekman, M. Irrational decision-making in an amoeboid organism: transitivity and context-dependent preferences. *Proceedings of the Royal Society of London B: Biological Sciences*, 278(1703):307–312, 2011.

Laughlin, S. B. and Sejnowski, T. J. Communication in neuronal networks. *Science*, 301(5641):1870–1874, 2003.

Laurent, G., MacLeod, K., Stopfer, M., and Wehr, M. Dynamic representation of odours by oscillating neural assemblies. In *Proceedings of the 10th International Symposium on Insect-Plant Relationships*, pages 7–18. Springer, 1999.

Lehrer, M., Horridge, G., Zhang, S., and Gadagkar, R. 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, 1995.

Lepora, N. F., Sullivan, J. C., Mitchinson, B., Pearson, M., Gurney, K., and Prescott, T. J. Brain-inspired bayesian perception for biomimetic robot

touch. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 5111–5116. IEEE, 2012.

Levy, W. B. and Baxter, R. A. Energy efficient neural codes. *Neural Computation*, 8(3):531–543, 1996.

Liang, L., Li, Y., Potter, C. J., Yizhar, O., Deisseroth, K., Tsien, R. W., and Luo, L. Gabaergic projection neurons route selective olfactory inputs to specific higher-order neurons. *Neuron*, 79(5):917–931, 2013.

Lindemann, J. P., Weiss, H., Möller, R., and Egelhaaf, M. Saccadic flight strategy facilitates collision avoidance: closed-loop performance of a cyberfly. *Biological Cybernetics*, 98(3):213–227, 2008.

Link, S. and Heath, R. A sequential theory of psychological discrimination. *Psychometrika*, 40(1):77–105, 1975.

Linster, C., Sachse, S., and Galizia, C. G. Computational modeling suggests that response properties rather than spatial position determine connectivity between olfactory glomeruli. *Journal of Neurophysiology*, 93(6):3410–3417, 2005.

Locatelli, F. F., Fernandez, P. C., Villareal, F., Muezzinoglu, K., Huerta, R., Galizia, C. G., and Smith, B. H. Nonassociative plasticity alters competitive interactions among mixture components in early olfactory processing. *European Journal of Neuroscience*, 37(1):63–79, 2013.

Loewenstein, G. Out of control: Visceral influences on behavior. *Organizational Behavior and Human Decision Processes*, 65(3):272–292, 1996.

Lorenz, E. N. Deterministic nonperiodic flow. *Journal of the Atmospheric Sciences*, 20(2):130–141, 1963.

Lorenz, K. and Kickert, R. W. The foundations of ethology. Springer, 1981.

Lorenz, K., Leyhausen, P., and Tonkin, B. *Motivation of human and animal behavior: an ethological view*. Van Nostrand Reinhold Company, 1973.

Mares, S., Ash, L., and Gronenberg, W. Brain allometry in bumblebee and honey bee workers. *Brain, Behavior and Evolution*, 66(1):50–61, 2005.

Marshall, J. A., Bogacz, R., Dornhaus, A., Planqué, R., Kovacs, T., and Franks, N. R. On optimal decision-making in brains and social insect colonies. *Journal of the Royal Society Interface*, pages rsif–2008, 2009.

Marshall, J. A., Bogacz, R., and Gilchrist, I. D. Consistent implementation of decisions in the brain. *PloS One*, 7(9):e43443, 2012.

Marshall, J. A., Favreau-Peigné, A., Fromhage, L., Mcnamara, J. M., Meah, L. F., and Houston, A. I. Cross inhibition improves activity selection when switching incurs time costs. *Current Zoology*, 61(2):242–250, 2015.

Martin, J.-R., Ernst, R., and Heisenberg, M. Mushroom bodies suppress locomotor activity in Drosophila melanogaster. *Learning & Memory*, 5(1): 179–191, 1998.

Martin, J.-R., Raabe, T., and Heisenberg, M. Central complex substructures are required for the maintenance of locomotor activity in Drosophila melanogaster. *Journal of Comparative Physiology A*, 185(3):277–288, 1999.

Masterman, R., Smith, B., and Spivak, M. Brood odor discrimination abilities in hygienic honey bees (Apis mellifera L.) using proboscis extension reflex conditioning. *Journal of Insect Behavior*, 13(1):87–101, 2000.

Mazurek, M. E., Roitman, J. D., Ditterich, J., and Shadlen, M. N. A role for neural integrators in perceptual decision making. *Cerebral Cortex*, 13(11): 1257–1269, 2003.

McConnell, J. V. Comparative physiology: learning in invertebrates. *Annual Review of Physiology*, 28(1):107–136, 1966.

McFarland, D. Decision making in animals. *Nature*, 1977.

McGuire, S. E., Le, P. T., and Davis, R. L. The role of Drosophila mushroom body signaling in olfactory memory. *Science*, 293(5533):1330–1333, 2001.

McMillen, T. and Holmes, P. The dynamics of choice among multiple alternatives. *Journal of Mathematical Psychology*, 50(1):30–57, 2006.

Menzel, R. Neurobiology of learning and memory: the honeybee as a model system. *Naturwissenschaften*, 70(10):504–511, 1983.

Menzel, R. The honeybee as a model for understanding the basis of cognition. *Nature Reviews Neuroscience*, 13(11):758–768, 2012.

Menzel, R. The insect mushroom body, an experience-dependent recoding device. *Journal of Physiology-Paris*, 2014.

Menzel, R. and Erber, J. Learning and memory in bees. *Scientific American*, 239(1):102–111, 1978.

Menzel, R. and Giurfa, M. Cognitive architecture of a mini-brain: the honeybee. *Trends in Cognitive Sciences*, 5(2):62–71, 2001.

Menzel, R. and Giurfa, M. Dimensions of cognition in an insect, the honeybee. *Behavioral and Cognitive Neuroscience Reviews*, 5(1):24–40, 2006.

Menzel, R. and Muller, U. Learning and memory in honeybees: from behavior to neural substrates. *Annual Review of Neuroscience*, 19(1):379–404, 1996.

Menzel, R., Giurfa, M., Gerber, B., and Hellstern, F. Cognition in insects: the honeybee as a study case. *Brain Evolution and Cognition*, pages 333–366, 2001.

Meyer, A. and Galizia, C. G. Elemental and configural olfactory coding by antennal lobe neurons of the honeybee (Apis mellifera). *Journal of Comparative Physiology A*, 198(2):159–171, 2012.

Meyer-Rochow, V. Electrophysiology and histology of the eye of the bumblebee Bombus hortorum (L.)(Hymenoptera: Apidae). *Journal of the Royal Society of New Zealand*, 11(2):123–153, 1981.

Milosavljevic, M., Malmaud, J., Huth, A., Koch, C., and Rangel, A. The drift diffusion model can account for value-based choice response times under high and low time pressure. *Judgment and Decision Making*, 5(6):437–449, 2010.

Mizunami, M., Iwasaki, M., Okada, R., and Nishikawa, M. Topography of four classes of kenyon cells in the mushroom bodies of the cockroach. *Journal of Comparative Neurology*, 399(2):162–175, 1998a.

Mizunami, M., Weibrecht, J. M., and Strausfeld, N. J. Mushroom bodies of the cockroach: their participation in place memory. *Journal of Comparative Neurology*, 402(4):520–537, 1998b.

Mobbs, P. The brain of the honeybee Apis mellifera. I. the connections and spatial organization of the mushroom bodies. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, pages 309–354, 1982.

Morawetz, L. and Spaethe, J. Visual attention in a complex search task differs between honeybees and bumblebees. *The Journal of Experimental Biology*, 215(14):2515–2523, 2012.

Mori, K., Nagao, H., and Yoshihara, Y. The olfactory bulb: coding and processing of odor molecule information. *Science*, 286(5440):711–715, 1999.

Mulder, M., Van Maanen, L., and Forstmann, B. Perceptual decision neurosciences–a model-based review. *Neuroscience*, 277:872–884, 2014.

Müller, D., Abel, R., Brandt, R., Zöckler, M., and Menzel, R. Differential parallel processing of olfactory information in the honeybee, Apis mellifera L. *Journal of Comparative Physiology A*, 188(5):359–370, 2002.

Mustard, J. A., Pham, P. M., and Smith, B. H. Modulation of motor behavior by dopamine and the d1-like dopamine receptor amdop2 in the honey bee. *Journal of Insect Physiology*, 56(4):422–430, 2010.

Nawrot, M. P. Dynamics of sensory processing in the dual olfactory pathway of the honeybee. *Apidologie*, 43(3):269–291, 2012.

Neckameyer, W. S. Dopamine and mushroom bodies in Drosophila: experience-dependent and independent aspects of sexual behavior. *Learning & Memory*, 5(1):157–165, 1998.

Neyman, J. and Pearson, E. S. On the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 231:289–337, 1933.

Nieh, J. C. The stop signal of honey bees: reconsidering its message. *Behavioral Ecology and Sociobiology*, 33(1):51–56, 1993.

Niggebrügge, C., Leboulle, G., Menzel, R., Komischke, B., and de Ibarra, N. H. Fast learning but coarse discrimination of colours in restrained honeybees. *Journal of Experimental Biology*, 212(9):1344–1350, 2009.

Niwa, M. and Ditterich, J. Perceptual decisions between multiple directions of visual motion. *The Journal of Neuroscience*, 28(17):4435–4445, 2008.

O'Dell, K., Armstrong, J. D., Yao Yang, M., and Kaiser, K. Functional dissection of the Drosophila mushroom bodies by selective feminization ofagenetically defined subcompartments. *Neuron*, 15(1):55–61, 1995.

Okada, R., Ikeda, J., and Mizunami, M. Sensory responses and movement-related activities in extrinsic neurons of the cockroach mushroom bodies. *Journal of Comparative Physiology A*, 185(2):115–129, 1999.

Olsen, S. R. and Wilson, R. I. Lateral presynaptic inhibition mediates gain control in an olfactory circuit. *Nature*, 452(7190):956, 2008.

Olsen, S. R., Bhandawat, V., and Wilson, R. I. Divisive normalization in olfactory population codes. *Neuron*, 66(2):287–299, 2010.

Otto, D. Untersuchungen zur zentralnervösen kontrolle der lauterzeugung von grillen. *Zeitschrift für Vergleichende Physiologie*, 74(3):227–271, 1971.

Page Jr, R., Erber, J., and Fondrk, M. The effect of genotype on response thresholds to sucrose and foraging behavior of honey bees (Apis mellifera L.). *Journal of Comparative Physiology A*, 182(4):489–500, 1998.

Pais, D., Hogan, P. M., Schlegel, T., Franks, N. R., Leonard, N. E., and Marshall, J. A. A mechanism for value-sensitive decision-making. *PLoS One*, 8(9):e73216, 2013.

Pankiw, T., Waddington, K. D., and Page, R. E. Modulation of sucrose response thresholds in honey bees (Apis mellifera L.): influence of genotype, feeding, and foraging experience. *Journal of Comparative Physiology A*, 187 (4):293–301, 2001.

Paoli, P. P., Donley, D., Stabler, D., Saseendranath, A., Nicolson, S. W., Simpson, S. J., and Wright, G. A. Nutritional balance of essential amino acids and carbohydrates of the adult worker honeybee depends on age. *Amino Acids*, 46(6):1449–1458, 2014.

Parnas, M., Lin, A. C., Huetteroth, W., and Miesenböck, G. Odor discrimination in Drosophila: from neural population codes to behavior. *Neuron*, 79 (5):932–944, 2013.

Passino, K. M. and Seeley, T. D. Modeling and analysis of nest-site selection by honeybee swarms: the speed and accuracy trade-off. *Behavioral Ecology and Sociobiology*, 59(3):427–442, 2006.

Pastor, K. A. and Seeley, T. D. The brief piping signal of the honey bee: begging call or stop signal? *Ethology*, 111(8):775–784, 2005.

Pavlov, I. P. Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex. Oxford, England: Oxford University Press., 1927.

Peele, P., Ditzen, M., Menzel, R., and Galizia, C. G. Appetitive odor learning does not change olfactory coding in a subpopulation of honeybee antennal lobe neurons. *Journal of Comparative Physiology A*, 192(10):1083–1103, 2006.

Perry, C. J. and Barron, A. B. Honey bees selectively avoid difficult choices. *Proceedings of the National Academy of Sciences*, 110(47):19155–19159, 2013a.

Perry, C. J. and Barron, A. B. Neural mechanisms of reward in insects. *Annual review of Entomology*, 58:543–562, 2013b.

Pettersson, L. B. and Brönmark, C. Trading off safety against food: state dependent habitat choice and foraging in crucian carp. *Oecologia*, 95(3): 353–357, 1993.

Pirrone, A., Stafford, T., and Marshall, J. A. When natural selection should optimize speed-accuracy trade-offs. *Frontiers in Neuroscience*, 8, 2014.

Plath, J. A. and Barron, A. B. Current progress in understanding the functions of the insect central complex. *Current Opinion in Insect Science*, 12:11–18, 2015.

Pool, A.-H., Kvello, P., Mann, K., Cheung, S. K., Gordon, M. D., Wang, L., and Scott, K. Four GABAergic interneurons impose feeding restraint in Drosophila. *Neuron*, 83(1):164–177, 2014.

Purcell, B. A. and Palmeri, T. J. Relating accumulator model parameters and neural dynamics. *Journal of Mathematical Psychology*, 76:156–171, 2017.

Ratcliff, R. A theory of memory retrieval. *Psychological Review*, 85(2):59, 1978.

Ratcliff, R. Continuous versus discrete information processing: Modeling accumulation of partial information. 1988.

Ratcliff, R. and McKoon, G. The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, 20(4):873–922, 2008.

Ratcliff, R. and Rouder, J. N. Modeling response times for two-choice decisions. *Psychological Science*, 9(5):347–356, 1998.

Ratcliff, R. and Rouder, J. N. A diffusion model account of masking in two-choice letter identification. *Journal of Experimental Psychology: Human Perception and Performance*, 26(1):127, 2000.

Ratcliff, R. and Smith, P. L. A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, 111(2):333, 2004.

Ratcliff, R., Van Zandt, T., and McKoon, G. Connectionist and diffusion models of reaction time. *Psychological Review*, 106(2):261–300, 1999.

Ratcliff, R., Gomez, P., and McKoon, G. A diffusion model account of the lexical decision task. *Psychological Review*, 111(1):159, 2004.

Ratcliff, R., Hasegawa, Y. T., Hasegawa, R. P., Smith, P. L., and Segraves, M. A. Dual diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task. *Journal of Neurophysiology*, 97(2):1756, 2007.

Ratcliff, R., Smith, P. L., Brown, S. D., and McKoon, G. Diffusion decision model: current issues and history. *Trends in Cognitive Sciences*, 20(4): 260–281, 2016.

Reeve, J. A grand theory of motivation: Why not? *Motivation and Emotion*, 40(1):31–35, 2016.

Reina, A., Valentini, G., Fernández-Oto, C., Dorigo, M., and Trianni, V. A design pattern for decentralised decision making. *PloS One*, 10(10): e0140950, 2015.

Renn, S. C., Armstrong, J. D., Yang, M., Wang, Z., An, X., Kaiser, K., and Taghert, P. H. Genetic analysis of the Drosophila ellipsoid body neuropil:

organization and development of the central complex. *Journal of Neurobiology*, 41(2):189–207, 1999.

Rescorla, R. A. Evidence for "unique stimulus" account of configural conditioning. *Journal of Comparative and Physiological Psychology*, 85(2):331, 1973.

Rieke, F. *Spikes: exploring the neural code*. Cambridge, MA: MIT press, 1999.

Riveros, A. J. and Gronenberg, W. Decision-making and associative color learning in harnessed bumblebees (Bombus impatiens). *Animal Cognition*, 15(6):1183–1193, 2012.

Rodrigues, D. Both associative learning and speed–accuracy trade-off occur in the southern monarch butterfly when visiting flowers. *Ethology Ecology & Evolution*, 28(1):30–41, 2016.

Rodrigues, M. A., Martins, N. E., Balancé, L. F., Broom, L. N., Dias, A. J., Fernandes, A. S. D., Rodrigues, F., Sucena, É., and Mirth, C. K. Drosophila melanogaster larvae make nutritional choices that minimize developmental time. *Journal of Insect Physiology*, 81:69–80, 2015.

Roitman, J. D. and Shadlen, M. N. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *The Journal of Neuroscience*, 22(21):9475–9489, 2002.

Rolls, E. T., Grabenhorst, F., and Deco, G. Choice, difficulty, and confidence in the brain. *Neuroimage*, 53(2):694–706, 2010.

Roubieu, F. L., Serres, J., Franceschini, N., Ruffier, F., and Viollet, S. A fully-autonomous hovercraft inspired by bees: wall following and speed control in straight and tapered corridors. In *Robotics and Biomimetics (ROBIO), 2012 IEEE International Conference on*, pages 1311–1318. IEEE, 2012.

Roussel, E., Sandoz, J.-C., and Giurfa, M. Searching for learning-dependent changes in the antennal lobe: simultaneous recording of neural activity and

aversive olfactory learning in honeybees. *Frontiers in Behavioral Neuroscience*, 4, 2010.

Rowe, J. B., Hughes, L., and Nimmo-Smith, I. Action selection: a race model for selected and non-selected actions distinguishes the contribution of premotor and prefrontal areas. *Neuroimage*, 51(2):888–896, 2010.

Ruffier, F., Viollet, S., Amic, S., and Franceschini, N. Bio-inspired optical flow circuits for the visual guidance of micro air vehicles. In *Circuits and Systems, 2003. ISCAS'03. Proceedings of the 2003 International Symposium on*, volume 3, pages III–846. IEEE, 2003.

Rybak, J. The digital honey bee brain atlas. In *Honeybee Neurobiology and Behavior*, pages 125–140. Springer, 2012.

Rybak, J., Kuß, A., Lamecker, H., Zachow, S., Hege, H.-C., Lienhard, M., Singer, J., Neubert, K., and Menzel, R. The digital bee brain: integrating and managing neurons in a common 3D reference system. *Frontiers in Systems Neuroscience*, 4, 2010.

Sandoz, J., Galizia, C. G., and Menzel, R. Side-specific olfactory conditioning leads to more specific odor representation between sides but not within sides in the honeybee antennal lobes. *Neuroscience*, 120(4):1137–1148, 2003.

Sandoz, J.-C. Olfaction in honey bees: from molecules to behavior. In *Honeybee Neurobiology and Behavior*, pages 235–252. Springer, 2012.

Scheiner, R. Responsiveness to sucrose and habituation of the proboscis extension response in honey bees. *Journal of Comparative Physiology A*, 190 (9):727–733, 2004.

Scheiner, R., Erber, J., and Page Jr, R. Tactile learning and the individual evaluation of the reward in honey bees (Apis mellifera L.). *Journal of Comparative Physiology A*, 185(1):1–10, 1999.

Scheiner, R., Barnert, M., and Erber, J. Variation in water and sucrose responsiveness during the foraging season affects proboscis extension learning in honey bees. *Apidologie*, 34(1):67–72, 2003.

Scheiner, R., Page, R. E., and Erber, J. Sucrose responsiveness and behavioral plasticity in honey bees (Apis mellifera). *Apidologie*, 35(2):133–142, 2004.

Schmuker, M. and Schneider, G. Processing and classification of chemical data inspired by insect olfaction. *Proceedings of the National Academy of Sciences*, 104(51):20285–20289, 2007.

Schoofs, A., Hückesfeld, S., Schlegel, P., Miroschnikow, A., Peters, M., Zeymer, M., Spieβ, R., Chiang, A.-S., and Pankratz, M. J. Selection of motor programs for suppressing food intake and inducing locomotion in the Drosophila brain. *PLoS Biology*, 12(6):e1001893, 2014.

Schröter, U. and Menzel, R. A new ascending sensory tract to the calyces of the honeybee mushroom body, the subesophageal-calycal tract. *Journal of Comparative Neurology*, 465(2):168–178, 2003.

Seeley, T. D. *Honeybee democracy*. Princeton University Press, 2010.

Seeley, T. D. and Buhrman, S. C. Group decision making in swarms of honey bees. *Behavioral Ecology and Sociobiology*, 45(1):19–31, 1999.

Seeley, T. D., Camazine, S., and Sneyd, J. Collective decision-making in honey bees: how colonies choose among nectar sources. *Behavioral Ecology and Sociobiology*, 28(4):277–290, 1991.

Seeley, T. D., Visscher, P. K., and Passino, K. M. Group decision making in honey bee swarms. *American Scientist*, 94(3):220, 2006.

Seeley, T. D., Visscher, P. K., Schlegel, T., Hogan, P. M., Franks, N. R., and Marshall, J. A. Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science*, 335(6064):108–111, 2012.

Shadlen, M. N. and Kiani, R. Decision making as a window on cognition. *Neuron*, 80(3):791–806, 2013.

Shadlen, M. N. and Newsome, W. T. Motion perception: seeing and deciding. *Proceedings of the National Academy of sciences*, 93(2):628–633, 1996.

Shadlen, M. N. and Newsome, W. T. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *The Journal of Neuroscience*, 18(10):3870–3896, 1998.

Shadlen, M. N. and Newsome, W. T. Neural basis of a perceptual decision in the parietal cortex (area lip) of the rhesus monkey. *Journal of Neurophysiology*, 86(4):1916–1936, 2001.

Shafir, E. and Tversky, A. *Decision making.* MIT Press, 1995.

Shih, C.-T., Sporns, O., Yuan, S.-L., Su, T.-S., Lin, Y.-J., Chuang, C.-C., Wang, T.-Y., Lo, C.-C., Greenspan, R. J., and Chiang, A.-S. Connectomics-based analysis of information flow in the Drosophila brain. *Current Biology*, 25(10):1249–1258, 2015.

Sibly, R. How incentive and deficit determine feeding tendency. *Animal Behaviour*, 23:437–446, 1975.

Silbering, A. F., Okada, R., Ito, K., and Galizia, C. G. Olfactory information processing in the Drosophila antennal lobe: anything goes? *The Journal of Neuroscience*, 28(49):13075–13087, 2008.

Simpson, S. and Raubenheimer, D. A multi-level analysis of feeding behaviour: the geometry of nutritional decisions. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 342(1302):381–402, 1993.

Simpson, S. J. and Raubenheimer, D. The nature of nutrition: a unifying framework from animal adaptation to human obesity. Princeton University Press, 2012.

Simpson, S. J., Sibly, R. M., Lee, K. P., Behmer, S. T., and Raubenheimer, D. Optimal foraging when regulating intake of multiple nutrients. *Animal Behaviour*, 68(6):1299–1311, 2004.

Smith, B. H. and Menzel, R. The use of electromyogram recordings to quantify odourant discrimination in the honey bee, Apis mellifera. *Journal of Insect Physiology*, 35(5):369–375, 1989.

Smith, D., Wessnitzer, J., and Webb, B. A model of associative learning in the mushroom body. *Biological Cybernetics*, 99(2):89–103, 2008.

Smith, P. L. Stochastic dynamic models of response time and accuracy: A foundational primer. *Journal of Mathematical Psychology*, 44(3):408–463, 2000.

Smith, P. L. and Ratcliff, R. Psychology and neurobiology of simple decisions. *Trends in Neurosciences*, 27(3):161–168, 2004.

Smith, P. L. and Vickers, D. The accumulator model of two-choice discrimination. *Journal of Mathematical Psychology*, 32(2):135–168, 1988.

Spieser, L., Servant, M., Hasbroucq, T., and Burle, B. Beyond decision! Motor contribution to speed–accuracy trade-off in decision-making. *Psychonomic Bulletin & Review*, pages 1–7, 2016.

Sporns, O., Tononi, G., and Kötter, R. The human connectome: a structural description of the human brain. *PLoS Computational Biology*, 1(4):e42, 2005.

Srinivasan, M., Zhang, S., Lehrer, M., and Collett, T. Honeybee navigation en route to the goal: visual flight control and odometry. *Journal of Experimental Biology*, 199(1):237–244, 1996.

Srinivasan, M., Zhang, S., and Gadakar, R. Context-dependent learning in honeybees. In *Proceedings of the 26th Göttingen Neurobiology Conference*, page 521, 1998.

Srinivasan, M. V. Honey bees as a model for vision, perception, and cognition. *Annual Review of Entomology*, 55:267–284, 2010.

Srinivasan, M. V. Honeybees as a model for the study of visually guided flight, navigation, and biologically inspired robotics. *Physiological Reviews*, 91(2): 413–460, 2011.

Stabler, D., Paoli, P. P., Nicolson, S. W., and Wright, G. A. Nutrient balancing of the adult worker bumblebee (Bombus terrestris) depends on the dietary source of essential amino acids. *Journal of Experimental Biology*, 218(5): 793–802, 2015.

Starns, J. J. and Ratcliff, R. The effects of aging on the speed–accuracy compromise: Boundary optimality in the diffusion model. *Psychology and Aging*, 25(2):377, 2010.

Stone, M. Models for choice-reaction time. *Psychometrika*, 25(3):251–260, 1960.

Strausfeld, N. and Campos-Ortega, J. Vision in insects: pathways possibly underlying neural adaptation and lateral inhibition. *Science*, 195(4281): 894–897, 1977.

Strausfeld, N. J. Atlas of an insect brain. Springer-Verlag Berlin, Heidelberg, New York, 1976.

Strausfeld, N. J. Organization of the honey bee mushroom body: representation of the calyx within the vertical and gamma lobes. *Journal of Comparative Neurology*, 450(1):4–33, 2002.

Strausfeld, N. J. and Hirth, F. Deep homology of arthropod central complex and vertebrate basal ganglia. *Science*, 340(6129):157–161, 2013.

Strausfeld, N. J., Hansen, L., Li, Y., Gomez, R. S., and Ito, K. Evolution, discovery, and interpretations of arthropod mushroom bodies. *Learning & Memory*, 5(1):11–37, 1998.

Strauss, R. The central complex and the genetic dissection of locomotor behaviour. *Current Opinion in Neurobiology*, 12(6):633–638, 2002.

Strauss, R. and Heisenberg, M. A higher control center of locomotor behavior in the Drosophila brain. *The Journal of Neuroscience*, 13(5):1852–1861, 1993.

Su, C.-Y., Menuz, K., Reisert, J., and Carlson, J. R. Non-synaptic inhibition between grouped neurons in an olfactory circuit. *Nature*, 492(7427):66–71, 2012.

Susswein, A. J., Weiss, K. R., and Kupfermann, I. The effects of food arousal on the latency of biting in Aplysia. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, 123(1):31–41, 1978.

Szyszka, P., Galkin, A., and Menzel, R. Associative and non-associative plasticity in kenyon cells of the honeybee mushroom body. *Frontiers in Systems Neuroscience*, 2, 2008.

Takeda, K. Classical conditioned response in the honey bee. *Journal of Insect Physiology*, 6(3):168–179, 1961.

Tammero, L. F. and Dickinson, M. H. Collision-avoidance and landing responses are mediated by separate pathways in the fruit fly, Drosophila melanogaster. *Journal of Experimental Biology*, 205(18):2785–2798, 2002.

Tanaka, N. K., Endo, K., and Ito, K. Organization of antennal lobe-associated neurons in adult Drosophila melanogaster brain. *Journal of Comparative Neurology*, 520(18):4067–4130, 2012.

Tang, S. and Guo, A. Choice behavior of Drosophila facing contradictory visual cues. *Science*, 294(5546):1543–1547, 2001.

Teodorescu, A. R. and Usher, M. Disentangling decision models: From independence to competition. *Psychological Review*, 120(1):1, 2013.

Tereshko, V. and Loengarov, A. Collective decision making in honey-bee foraging dynamics. *Computing and Information Systems*, 9(3):1, 2005.

Thura, D., Beauregard-Racine, J., Fradet, C.-W., and Cisek, P. Decision making by urgency gating: theory and experimental support. *Journal of Neurophysiology*, 108(11):2912–2930, 2012.

Tinbergen, N. On the orientation of the digger wasp Philanthus triangulum Fabr. I. *The Animal in its World*, 1:103–127, 1972.

Travers, S. E. and Sih, A. The influence of starvation and predators on the mating behavior of a semiaquatic insect. *Ecology*, 72(6):2123–2136, 1991.

Trimmer, P. C., Houston, A. I., Marshall, J. A., Mendl, M. T., Paul, E. S., and McNamara, J. M. Decision-making under uncertainty: biases and bayesians. *Animal Cognition*, 14(4):465–476, 2011.

Tsetsos, K., Usher, M., and McClelland, J. L. Testing multi-alternative decision models with non-stationary evidence. *Frontiers in Neuroscience*, 5, 2011.

Tsetsos, K., Chater, N., and Usher, M. Salience driven value integration explains decision biases and preference reversal. *Proceedings of the National Academy of Sciences*, 109(24):9659–9664, 2012.

Uhlenbeck, G. E. and Ornstein, L. S. On the theory of the brownian motion. *Physical Review*, 36(5):823, 1930.

Urban, N. N. Lateral inhibition in the olfactory bulb and in olfaction. *Physiology & Behavior*, 77(4):607–612, 2002.

Usher, M. and McClelland, J. L. The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Review*, 108(3):550, 2001.

Usher, M. and McClelland, J. L. Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological Review*, 111(3):757, 2004.

Van den Bos, R., Harteveld, M., and Stoop, H. Stress and decision-making in humans: performance is related to cortisol reactivity, albeit differently in men and women. *Psychoneuroendocrinology*, 34(10):1449–1458, 2009.

van den Heuvel, M. P., Bullmore, E. T., and Sporns, O. Comparative connectomics. *Trends in Cognitive Sciences*, 20(5):345–361, 2016.

Van Hateren, J., Srinivasan, M., and Wait, P. Pattern recognition in bees: orientation discrimination. *Journal of Comparative Physiology A*, 167(5): 649–654, 1990.

Van Zandt, T., Colonius, H., and Proctor, R. W. A comparison of two response time models applied to perceptual matching. *Psychonomic Bulletin & Review*, 7(2):208–256, 2000.

Vasilaki, E., Frémaux, N., Urbanczik, R., Senn, W., and Gerstner, W. Spike-based reinforcement learning in continuous state and action space: when policy gradient methods fail. *PLoS Computational Biology*, 5(12):e1000586, 2009.

Vickers, D. Evidence for an accumulator model of psychophysical discrimination. *Ergonomics*, 13(1):37–58, 1970.

Visscher, P. K. Group decision making in nest-site selection among social insects. *Annual Review of Entomology*, 52:255–275, 2007.

Von Frisch, K. Dialects in the language of the bees. *Scientific American*, 207 (2):78–89, 1962.

Von Frisch, K. Die tänze der Bienen. In *Tanzsprache und Orientierung der Bienen*, pages 3–330. Springer, 1965.

Von Frisch, K. The dance language and orientation of bees. Harvard University Press, 1967.

Vowles, D. The structure and connexions of the corpora pedunculata in bees and ants. *Quarterly Journal of Microscopical Science*, 3(34):239–255, 1955.

Wadepuhl, M. Control of grasshopper singing behavior by the brain: Responses to electrical stimulation1. *Zeitschrift für Tierpsychologie*, 63(2-3): 173–200, 1983.

Wadepuhl, M. and Huber, F. Elicitation of singing and courtship movements by electrical stimulation of the brain of the grasshopper. *Naturwissenschaften*, 66(6):320–322, 1979.

Wald, A., Wolfowitz, J., et al. Optimum character of the sequential probability ratio test. *The Annals of Mathematical Statistics*, 19(3):326–339, 1948.

Walker-Smith, G. J., Gale, A. G., and Findlay, J. M. Eye movement strategies involved in face perception. *Perception*, 42(11):1120–1133, 2013.

Wang, M.-Y., Ings, T. C., Proulx, M. J., and Chittka, L. Can bees simultaneously engage in adaptive foraging behaviour and attend to cryptic predators? *Animal Behaviour*, 86(4):859–866, 2013.

Wang, X.-J. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, 36(5):955–968, 2002.

Wang, X.-J. Decision making in recurrent neuronal circuits. *Neuron*, 60(2): 215–234, 2008.

Waser, N. M. Flower constancy: definition, cause, and measurement. *American Naturalist*, pages 593–603, 1986.

Wehner, R. Desert ant navigation: how miniature brains solve complex tasks. *Journal of Comparative Physiology A*, 189(8):579–588, 2003.

Weiss, K., Koch, U., Koester, J., Rosen, S., and Kupfermann, I. The role of arousal in modulating feeding behavior of Aplysia: neural and behavioral studies. *The Neural Basis of Feeding and Reward*, pages 25–57, 1982.

White, J. G., Southgate, E., Thomson, J. N., and Brenner, S. The structure of the nervous system of the nematode Caenorhabditis elegans: the mind

of a worm. *Philosophical Transactions of the Royal Society London B*, 314: 1–340, 1986.

Williamson, C. E. The predatory behavior of Mesocyclops edax: Predator preferences, prey defenses, and starvation-induced changes. *Limnology and Oceanography*, 25(5):903–909, 1980.

Wilson, R. I. and Laurent, G. Role of GABAergic inhibition in shaping odor-evoked spatiotemporal patterns in the Drosophila antennal lobe. *The Journal of Neuroscience*, 25(40):9069–9079, 2005.

Wright, G. A., Carlton, M., and Smith, B. H. A honeybee's ability to learn, recognize, and discriminate odors depends upon odor sampling time and concentration. *Behavioral Neuroscience*, 123(1):36, 2009.

Wright, R. A. Motivation theory essentials: Understanding motives and their conversion into effortful goal pursuit. *Motivation and Emotion*, 40(1):16–21, 2016.

Young, J. and Armstrong, J. Structure of the adult central complex in Drosophila: organization of distinct neuronal subsets. *Journal of Comparative Neurology*, 518(9):1500–1524, 2010.

Zeleny, M. Multiple criteria decision making: Eight concepts of optimality. *Human Systems Management*, 17(2):97–107, 1998.

Zeleny, M. and Cochrane, J. L. *Multiple criteria decision making*, volume 25. McGraw-Hill New York, 1982.

Zettler, F. and Järvilehto, M. Lateral inhibition in an insect eye. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, 76(3):233–244, 1972.

Zhang, J. The effects of evidence bounds on decision-making: theoretical and empirical developments. *Frontiers in Psychology*, 3, 2012.

Zhang, K., Guo, J. Z., Peng, Y., Xi, W., and Guo, A. Dopamine-mushroom body circuit regulates saliency-based decision-making in Drosophila. *Science*, 316(5833):1901–1904, 2007.

Zhang, S. and Srinivasan, M. Prior experience enhances pattern discrimination in insect vision. *Nature*, 1994.

Zhang, S., Bartsch, K., and Srinivasan, M. Maze learning by honeybees. *Neurobiology of Learning and Memory*, 66(3):267–282, 1996.

Zhang, S., Lehrer, M., and Srinivasan, M. Honeybee memory: navigation by associative grouping and recall of visual stimuli. *Neurobiology of Learning and Memory*, 72(3):180–201, 1999.

Zhang, S., Mizutani, A., and Srinivasan, M. V. Maze navigation by honeybees: learning path regularity. *Learning & Memory*, 7(6):363–374, 2000.

Zhang, S., Schwarz, S., Pahl, M., Zhu, H., and Tautz, J. Honeybee memory: a honeybee knows what to do and when. *Journal of Experimental Biology*, 209(22):4420–4428, 2006.

Zhang, S., Si, A., and Pahl, M. Decision making in honeybees. *Frontiers in Neuroscience*, 2012a.

Zhang, S., Si, A., and Pahl, M. Visually guided decision making in foraging honeybees. *Frontiers in Neuroscience*, 6, 2012b.

Zwaka, H., Münch, D., Manz, G., Menzel, R., and Rybak, J. The circuitry of olfactory projection neurons in the brain of the honeybee, Apis mellifera. *Frontiers in Neuroanatomy*, 10, 2016.