# Emotions, Motivation-Based Action Selection and Dynamic Environments

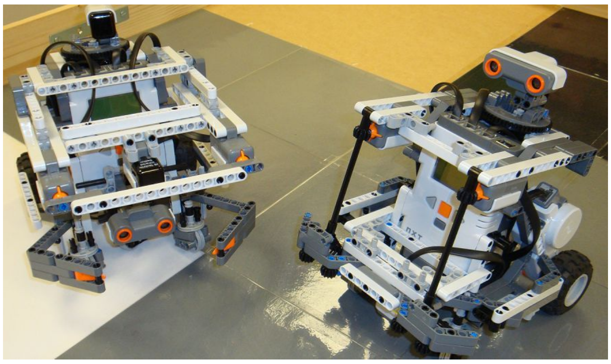

SUBMITTED TO THE
**UNIVERSITY OF HERTFORDSHIRE**
IN PARTIAL FULFILMENT OF THE REQUIREMENTS
OF THE DEGREE OF

*Doctor of Philosophy*

Claire O'Bryne

FACULTY OF ENGINEERING AND
INFORMATION SCIENCES
APRIL 2013

# A BOTTOM-UP INVESTIGATION OF BRAIN-BODY-ENVIRONMENT INTERACTIONS AND EMOTION-MODULATED ACTION SELECTION IN A PREDATOR-PREY SCENARIO

*For Mum and Dad...*

# Acknowledgements

Considering my own emotions in the time it has taken to write this thesis, I have sometimes felt just as much of a case study as my robots. This is perhaps especially true for fear and its effect on my own decision-making abilities. For example, while not always appreciating the irony, I *have* come to appreciate that (as with my robots) a little fear in the right place can be a “good thing”. Yet, unlike them, I have also had the consolation of emotions that are much more pleasant to experience. For helping to create some of those in terms of my research and its direction, I would like to say thank you first to my supervisors: Lola Cañamero, Iain Werry and Maria Schilstra. A special thank you has to go to Rene te Boekhorst, for his statistical advice and expertise, as well as general willingness to take time out of a very busy schedule to impart it. More generally, but no less sincerely, thanks to everyone, past and present, at UH for making the office such a welcoming place to be. This especially includes the ladies in the admin office who are always helpful when it comes to rules and forms. More personally though, thanks Mum, Dad, Ian, Lisa, Kim, Kerry and all who know me best! This includes Shirley, who helped make University a possibility for me. While I doubt any of you will ever actually read this fully, as you will each get bored long before the end (yeah, I know you will!) I hope you *will* manage to make it this far at least and know I really appreciate all you have done, and sacrificed, to help me along the way — whether it be designing and building an arena for my robots (thanks Dad!) flat-hunting or organising holidays (thanks Mum!). You have all helped keep me going, reminding me that, though I am lucky to have had the opportunity to “play with LEGO all day” there are also many other things in the world to be grateful for... like you. Finally, because I smile to think what you would make of me having written a thesis on LEGO Robots, but know you would have been all for it and I would like to think proud regardless... Thanks Grandma, Thanks Grandad!

# Abstract

In contrast to traditional approaches, where the focus is on developing or evolving artificial "brains" as the route to artificial intelligence (AI) more recent approaches have increasingly emphasised and modelled the role of "bodies" and "environments". In turn, this has further encouraged ideas regarding aspects of intelligence as being best thought of as distributed across agent brains, bodies and environments. That is, as system properties emerging from interactions of these components. Action selection is commonly recognised as one of the problems all agents, whether biological or artificial, must face: deciding at any given moment "what to do next". Researchers have generated many different action selection mechanisms as "solutions" to this problem. However, in the work of this thesis, we focus on one which takes its inspiration from biological ideas about the role and possible neural substrates of emotion. We use this to consider how models of brain-body-environment interactions might be more useful for the study of emotion, as well as action selection mechanisms. For, despite the many mechanisms proposed, the literature still lacks systematic ways to analyse their performance in combination with different physical and/or perceptual capabilities. That is, factors relating more directly to agent *embodiment*. In this thesis we have studied the performance of our selected architecture in a robotic predator-prey scenario known as the Hazardous Three Resource Problem. The predator-prey relationship is popular in artificial intelligence, both as an action selection problem and a situation which enables study of agent-agent interactions. Predators can act as catalysts for the evolution of prey agents in a "survival of the fittest" sense while, in their turn, prey agents are tests of predator ingenuity. For us, however, it is also a situation where emotion might naturally be assumed to have useful functions. To study action selection, emotion and brain-body-environment interactions in an artificial predator-prey relationship, we both advocate and adopt a bottom-up, animat approach. The animat approach to AI is one that emphasizes characteristics neglected by more traditional approaches. As such, it has embraced the study of robotic agents. One reason for this is the process of designing "real-world" agents forces us to consider practicalities simulations might not. What makes the use of robots particularly appealing for our work, however, is how it can give us a greater appreciation of more physical aspects of intelligence such as agent morphology and its integration with agent control mechanisms as well as *environmental dynamics*. Using LEGO robots, we show how the performance of our architecture varies in our chosen scenario with aspects of agent brain, body and environment. We argue our results complement existing research by contributing evidence from a real-world implementation, explicitly modelling ideas about action selection and emotion as distributed across, or best thought of as emerging from interactions between, agent brain, body and environment. In particular, this thesis shows how our selected architecture varies and benefits from further integration with aspects of agent "body". It also acts as an example of an alternative form for the bottom-up development of artificial emotion, demonstrating wider applications for creating more adaptive action selection mechanisms. Comparing the robotic predator-prey relationships we have created to ethological evidence and theories, we argue our architecture may also have specific potential for future research and applications — having already proven itself capable of emerging multiple functions and properties.

*“If we knew what it was we were doing, it would not be called research, would it?”*
**— Albert Einstein**

# Contents

# Chapter 1

# Introduction

*"We are not students of some subject matter, but students of problems. And problems may cut right across the borders of any subject matter or discipline."*
**— Karl Popper**

Within this thesis we[1] ask questions and explore problems which can be considered to "cut right across the borders" traditionally drawn between various research areas and scientific disciplines. An interdisciplinary endeavour, the initial motivation for the work of this thesis was to make a contribution to existing literature by drawing together ideas and evidence from a number of different research areas, concerning and incorporating ideas from the study of both biological and artificial agents. Generally speaking, in doing so we wanted to recognise and encourage others to think about how, in order to find solutions to problems in one discipline, we should consider and start to integrate our own work even further with the work of those in others. The interdisciplinary nature and implications of the questions we pose have the added advantage of making our work of potential interest and relevance both across and within disciplines. For example, by integrating ideas from fields where researchers in one have largely ignored the work of those outside their own. More specifically though, we aim here for our work to help lay the foundations for a new approach to the study and development of artificial agents for exploring aspects of intelligence. To this end, we have focused our efforts on helping to develop new tools and techniques to complement those currently used by researchers. Our starting point? The development of a model with which to explore our ideas about action selection, emotion and brain-body-environment interactions in a predator-prey relationship.

---

[1] For stylistic reasons, throughout this thesis "we" is used instead of "I". This is so as to be consistent with the general style of the papers submitted as a result of this author's research.

## 1.1 Cutting across borders

### 1.1.1 Artificial Intelligence

What is intelligence? How do we recognise it when we see it? Can we build artificial agents that are capable of the same level of "intelligence" as we claim for ourselves? If so, how and where should we focus our efforts so as to create such agents? These questions lie at the heart of artificial intelligence (AI) and overlap both intimately and intricately with many related areas and disciplines. Because of this, in the attempt to find answers, many researchers have naturally linked research from the study of biological organisms to the construction and design of artificial agents [1]. While some have constructed high-fidelity models of biological processes [2], others have focused more on useful abstractions of biological ideas [3]. For instance, extracting and abstracting ideas such as evolution to act as general design principles [4]. However, what is appealing about all such biologically-inspired research is that it not only gives us a source of inspiration for developing artificial intelligence, but may also help us find out more and clarify ideas about our own intelligence in return [5].

### 1.1.2 Artificial Life

In comparison with the field of AI, the questions at the heart of artificial life (AL) are somewhat broader, but nevertheless tend to follow the same general spirit of enquiry. What is life? How do we recognise it when we see it? Can we build artificial agents that "live" and/or can help us understand and develop principles about the processes of life? In terms of its position as a field of research, AL may be a relative latecomer compared to AI. However, one of its most exciting features in terms of philosophy is that it studies not only "life-as-we-know-it", but also "life-as-it-could-be" [6]. Perhaps one of its most notable strengths, further advocated in the work of this thesis, is that it has embraced bottom-up approaches to the study of artificial systems much more readily than other fields [7]. Such approaches try to answer these difficult questions by beginning with the details, before working up to the highest conceptual level. Though many AI researchers have also now started to accept this kind of approach, as evidenced by the introduction of the behaviour-based, animat approach to AI (more on this later) this is in direct contrast to the more traditional AI top-down approach, which tends to approach problems related to intelligence at the highest conceptual level, before working down to the details.

### 1.1.3 Action Selection

AL researchers tend to suggest that the path to AI is to model, whether in computers or in robots, the vital dynamics that support mental life [8]. Adopting an interdisciplinary approach, this thesis supports this view and likewise aims to contribute towards a better understanding of tasks an agent, whether biological or artificial, must deal with to be considered intelligent. It is commonly accepted that an “intelligent” agent must, by necessity, be capable of making its own decisions as to what course of action to take at any given time [9]. This relates strongly to notions of agent autonomy and adaptive behaviour, but also more specifically to the problem of action selection[1]. The latter can perhaps be defined most simply as a problem which all agents face of ‘what to do next’ [10; 11; 12]. As a research area that has its roots in the very core of AI research, many solutions to this problem have been proposed in the form of action selection mechanisms [13; 14; 15]. However, as the purpose of this chapter is essentially to provide an overview and brief introduction of the primary problems and ideas we are concerned with in our research, a more in-depth discussion of these in terms of the existing literature will be saved for later chapters.

### 1.1.4 Emotion

Counter-intuitively to early ideas regarding logical decisions as more “intelligent” ones and devoid of emotional influence, a relatively recent branch of the action selection literature has looked at how models of emotion might be used to develop more adaptive agents, especially in the context of dynamic environments (those which change over time) [16; 17]. Consequently, a particular subset of action selection mechanisms can now be classed as “emotion-based” [18]. That is to say, they incorporate and model ideas about the role of emotion in rational decision-making for humans. It is on this type of action selection mechanism that we focus our own attention. However, in the work of this thesis, in choosing to adopt a bottom-up approach, we further restrict our attentions to the study of one in particular. This we selected for several reasons, not least because it incorporates a hormone-like mechanism, used to modify inputs to a motivation-based architecture, which simulates *interoceptive modulation*[2]. This mechanism may be of special interest to those wanting to adopt a bottom-up approach to the study of emotion in that it relates to the work of researchers such as Fellous [19; 20], who suggest the biological roots (and neural substrate) of emotion are to be found as patterns of neuromodulation.

### 1.1.5 The Brain-Body-Environment (B-B-E) Relationship

Since the early days of AI, ideas about the relative importance (and potential of models) of agent brains, bodies and environments for the development of adaptive agents have also been revised: not only in terms of the biological, but also the artificial [21].

---

[1] Within other disciplines such as psychology, this can be thought of as akin to decision-making

[2] That is to say, it modulates the sensitivity to stimuli originating inside the body

Collectively, researchers to-date have focused on (modelling) many types and aspects of brain-body-environment relationships, both for producing and understanding adaptive behaviour. Researchers such as Clark [22] encourage us to rethink notions such as intelligence as properties of systems that emerge only when the balance between brain, body and environment is right. One aspect of this relationship, which we also focus on in our research, relates to ideas about the importance of the body. Initial research within AI generally focused on abstracting away the body of an agent completely, treating it as peripheral at most. However, the limitations of this kind of research have since led an increasing number of researchers to encourage us to model and think more about how our own (and artificial) intelligence might be shaped by and emerge from interactions between brain, body and environment [23].

### 1.1.6 Embodied Action Selection

With regards to action selection mechanisms and the importance of body, many action selection mechanisms have been proposed and successfully implemented in agent bodies with quite different capabilities: from "blind" agents, such as Miglino et al's [24] to agents with multi-modal abilities such as the robots Kismet [25] or Psikharpax [26]. Researchers such as Ziemke have even tried to co-evolve agent controllers with agent bodies [27]. However, within the literature to-date, it is still difficult to find systematic studies showing how the performance of a given emotion-based architecture will vary with specific aspects of agent embodiment per se. This thesis attempts to address this, by seeing how aspects of embodiment will impact our chosen emotion-based architecture's performance. In our work, we have adopted and advocate the use of the animat approach for this purpose, considering it ideal for conducting an exploratory study of aspects of the brain-body-environment relationship (the reasons for which we will demonstrate in subsequent chapters).

### 1.1.7 Dynamic Environments

Other aspects of the brain-body-environment relationship, which we also consider in our model, relate more directly to agent environment. Again, in tune with the interdisciplinary and biologically-inspired nature of our research, we focus on a specific type: that of predator and prey. One reason for this is that it is an environment which both encapsulates the problem of action selection and is one in which emotion is thought to have a special role[1]. But another is that it is a scenario where we can take an even wider perspective than is common. That is, to look at and compare affect-modulated action selection (as performed by our chosen mechanism) not only in terms of brain-

---

[1] Explanations of emotions often suggest the primary emotion of fear having a role here in the fight-or-flight response of the prey [28]

body-environment interactions, but agent-agent interactions as well. The predator-prey relationship can be used as a good metaphor for many environments, but particularly those with hazardous aspects. That is, dynamic environments where risk assessment and risk-taking has to be both considered and dealt with; requiring agents to be adaptive.

### 1.1.8 Our Research

This thesis argues that the adaptive value (including the emergent functions) of an emotion-based architecture will be affected by the bodily and perceptual abilities of both the agent itself and other agents present in its environment. Moreover, by studying such systems and manipulating these abilities we might therefore create, or learn how to create, more adaptive agents from simple, reactive ones. This may be particularly true for resource-limited agents making real-time decisions, such as robots, performing tasks involving trade-offs and opportunity costs. The aim of this thesis is to better understand these elements by considering them, in greater detail, in one body of (systematic) research. To this end a bottom-up[1], animat approach is both advocated and followed here to ask and answer the question: "How will changes in the perceptual abilities of predator and prey agents interact to affect the dynamics of their relationship, especially in terms of the adaptive value of an emotion-based architecture for action selection?"

In the research outlined in this thesis, we model brain-body-environment interactions in an artificial predator-prey relationship, developing a methodology and running experiments to enable us to explore both how interactions between components of our model (brain, body, environment and agent) may be affected by the addition of an emotion-like component, as well as how such an emotion-like component might fit into such a relationship or "mediate" these components (brain, body, environment and agent). Our approach therefore allows us to provide a much more qualitative, detailed picture of emotion-based action selection at work in a particular scenario. One of our principal aims here is to highlight or provide markers for points of interest that might be useful to other researchers. Thus, we mainly aim for our results to provide a "proof of concept", seeing how we can connect the success of our emotion-based architecture to aspects of the brain-body-environment relationship.

### 1.1.9 A Goldilocks Approach

One way to describe our own approach is by the Goldilocks principle. For, we do not aim to optimise our action selection mechanism as such, but to hopefully discover — or

---

[1] More on this term in the next chapter.

at least learn something useful about — those conditions[1] under which emotion-based action selection mechanisms become "just right" for artificial agents. Our research thus coincidentally looks at how well the Goldilocks principle explains the success (or not) of emotion-based architectures. One recent suggestion following this is that more adaptive agents will be developed when we stop abstracting away the body, embracing it as being potentially as important for intelligent or adaptive behaviour as agent "brains". In this respect, we study intelligence and emotion as properties of systems, emerging from the interactions of brains, bodies and environments.

### 1.1.10 Hazardous Three Resource Problem

To study the predator-prey relationship further, we have designed and created an implementation of a specific predator-prey robotic scenario (the Hazardous Three Resource Problem; a model designed for the study of action selection [29]) for use as a research platform. LEGO NXT robots, an off-the-shelf robotic platform, were used to create prey and predator agents for several reasons. One of these is that, within this thesis, we are primarily interested in action-selection for real-world systems, with real-world constraints being imposed on them. This includes, for example, the laws of physics, noisy sensors and real-time decision-making.

Varying aspects of both our developed robotic agents' embodiment, such as perceptual distance and perceptual field, we have conducted experiments to examine and study the effect observed on each agent's performance, as well as their relationship, when our prey is implemented with an emotion-based architecture. Drawing on our early findings and observations, in later experiments we have then examined how the emotion-like component of our chosen architecture might be extended across the brain-body-environment divide (and in what way this will be to our agent's advantage). To do this, in our experiments we have looked at how we might manipulate the connections between our emotion-like mechanism and the brain-body-environment components represented in our model.

In this way, we have used our implementation to look at how various factors, represented in changes to our basic implementation of the H3RP, may affect the adaptive value of our chosen emotion-based architecture. From these experiments, we aim to see how our architecture might be improved or manipulated so as to take advantage of brain-body-environment interactions and generate more adaptive behaviour. Our more recent findings, outlined in later chapters, have led us to think about how such changes might generate adaptive behaviours when physical abilities are given some kind of "cost".

---

[1] Brain-Body-Environment combinations

### 1.1.11 Levels of Analysis

One of the advantages of both adopting the animat approach and studying predator-prey relationships is the way it allows us to look at our developed model using multiple levels of analysis. With regards to our own analysis of the data we have collected during our experiments, we have adopted numerous measures in the process of our research, so as to enable us to consider how different measures define "adaptive value" and/or recognise the adaptive behaviour of our agents differently. For instance, during our experiments, we have looked at how ideas about state transitions (as used by ethologists to study animal behaviour) may provide an alternative means for both capturing and describing changes in our agents' behaviour. While most of our data has been collected internally from our robots, in later experiments we have also considered results obtained from observational analysis. This has both enabled us to make comparisons between different measures, as well as to think about introducing additional measures of our own. In particular, within this thesis we put forward our own, novel plots for use in analysis and exploration of agent behaviour — referred to here as "brain-body-environment maps". Mapping elements of the brain-body-environment relationship represented in our model, and going by the extra detail these maps have enabled us to capture as a result, we believe these maps hold great promise as a tool to aid future research.

## 1.2 Overview

This section aims to give a brief overview of the different chapters composing this thesis:

### 1.2.1 Chapter 2

This chapter introduces the animat approach to AI and AL as a method for developing and studying artificial agents. We believe this to be one that is particularly well suited to those with our research interests (as already identified in this chapter). Thus, a general outline of this approach is given, before being used to provide the context for our thesis, setting the scene in terms of the existing literature. In this way, we are not just adopting this approach ourselves, but further using it to develop a new framework for categorising the literature related to our research question by relating elements of an animat to the concepts we are interested in studying together. Here, this is achieved using the concepts of animat brains, bodies and environments to help us both categorise and bring together several different threads of research. This includes work related to action selection and aspects of the brain-body-environment relationship. We argue the animat approach is a good choice of approach for studying such ideas further.

The chapter then concludes with discussion of our particular environment — that of predator and prey — outlining in more detail some of the advantages of studying the predator-prey relationship and focusing on it in terms of action selection and emotion.

### 1.2.2 Chapter 3

Here we describe our model, providing an outline and history of our chosen scenario (the H3RP) and the architecture studied and used throughout the rest of this thesis as a model of action selection. This is a scenario which has been previously studied so a brief review of the findings from previous research is given. We review some of the literature with regards to emotion-based architectures and biologically-inspired, robotic implementations, comparing our chosen architecture to existing emotion-based ones using several criteria. Through these criteria both the aims and work of this thesis are further distinguished from those of existing researchers (highlighting the novel aspects at its core). The focus of this chapter then switches to the consideration of how adaptive value has and, with regard to the performance of our implemented agents, might be measured. The idea of using other measures of performance is also introduced in light of the re-conceptualisation of action selection as activity cycles.

### 1.2.3 Chapter 4

This chapter details the technical details of our novel implementation of the Hazardous Three Resource Problem, developed as a research platform for this thesis. Focusing on providing a description of our basic agents (which take the form of a LEGO NXT prey and LEGO NXT predator robot) the different components, including “brain” “body” and “environment”, of these animats are both described and illustrated. This chapter thereby provides an overview of both the software and hardware used to conduct the experiments reported in later chapters. It also provides specific details of the implemented hormone-like mechanism, used to simulate interoceptive modulation in our prey. Finally, we define the measures used in our experiments to study the performance of our robots.

### 1.2.4 Chapters 5-7

The results of experiments performed using our implementation of the H3RP are then reported. Initial experiments looked at the impact of prey perceptual distance on performance of the architecture used for a prey animat. Moving on from and extending these experiments, we then looked at how prey *and* predator perceptual distance might interact to affect performance in the H3RP. Modifications were made to the underlying hormone-like mechanism to see whether a sigmoidal decay rate might provide more interesting results. The implication here is that the predator, as a “brain within a body” as well as an aspect of the prey's environment, will also determine the success of our chosen architecture.

Considering the "sensitivity" of our emotion-like mechanism, we then go on in our experiments to look at how changing the parameters of the emotion-like component of our action selection mechanism itself (with regards to the "level" or "gradience" of the fear it simulates) might affect prey performance. Focusing on both different levels as well as the design and use of a gradient response, the implication here is that modifying the emotion-like component of our mechanism to take into account other sensory information will also lead to differences in prey performance, as measured in environments using different types of predator agents. Finally, we move on to examine the factors of prey perceptual field and speed (trying to integrate the emotion-like component more closely with aspects of prey body).

### 1.2.5 Chapter 8

This chapter draws together and compares the results obtained across all of our experiments during the study of our developed system (H3RP). Here, we hope to extract further insights from our work, drawing on our research experience so as to provide a summary of our conclusions in terms of our initial research objectives. Reflections are given on the potential implications of the results obtained and reported in previous chapters, while consideration is given to the possibilities for future research.

## 1.3 Contributions

The main contributions of this thesis can be summarised as follows (these will be expanded on in greater detail in subsequent chapters):

### 1.3.1 For Emotion-Based Action Selection

- Development of an implementation of the Hazardous Three Resource Problem (H3RP) — a predator-prey environment — in which to study and develop architectures for action selection, as well as to study the influence of different aspects of an agent's brain-body-environment relationship such as embodiment.[1] With regards to our research question, we advocate this as an ideal platform for our research; in which we can study together all the elements we are interested in, as encapsulated in our research question

---

[1] In initial experiments the focus was on perceptual distance as a specific aspect of embodiment, for which we divided perception into proximal and distal types (combinations of which made further subproblems/versions of the scenario). In latter experiments perceptual field and physical ability (speed) were also studied as aspects of embodiment

- A more systematic (step-by-step) study of the predator-prey type relationship than has yet been conducted with regards to action selection and affect. That is, looking at how an emotion-based action selection mechanism, when placed in a prey agent, not only varies with individual changes to a single prey... but can interact with and causes changes in a single predator
- An analysis of the adaptive value, in terms of costs and benefits, of both emotions and decisions in the predator-prey relationship. To do this, we take inspiration from ethological research (the study of animal behaviour)
- Consideration and investigation of emotion-like action selection mechanisms as adaptive mechanisms, in addition to (even beyond) as mechanisms for adaptation. In other words, we consider how emotion-based architectures themselves can adapt to changes over time, in addition to being used to adapt an agent

### 1.3.2 For the Analysis and Design of Adaptive Systems

- A comparison and evaluation of measures of adaptive value (both quantitative and qualitative) that might be adopted (to evaluate performance of an agent)
- An analysis of the action selection problem in terms of the brain-body-environment relationship
- Consideration and integration of internal and observational metrics beyond the work of other researchers
- Demonstration of how researchers might manipulate or adjust parameters so as to better “fine-tune” mechanisms like our own and increase their value for adaptive action selection in similar contexts (of the predator-prey scenario)

## 1.4 Summary

To summarise then, the primary problems or areas that characterize this thesis include:

- The Animat Approach: as one which we consider ideal for studying all the problems and areas we are interested in, advocating as it does robotic implementation (and simulations)
- Action Selection: in the form of action selection mechanisms
- Artificial Emotion: specifically, in the form of emotion-based action selection mechanisms

- The Brain-Body-Environment Relationship: through consideration of the importance of individual components of this relationship, including aspects such as embodiment, as well as the interactions arising between them
- The Predator-Prey Relationship: among other aspects, as a type of environment that is *dynamic*, but also as a very interesting relationship between two agents
- The Goldilocks Principle: as a way of describing the approach we take (and the philosophy which sees intelligence as a system property) to studying our emotion-based action selection mechanism

### 1.4.1 Studying Emotion for Action Selection in a Predator-Prey Relationship

This thesis focuses primarily on the study of action selection in conjunction with the modelling of emotion in artificial prey agents (robots in particular). That is, how we might design artificial agents that will decide at any given moment “what to do next”, but especially focusing on the promise of biologically-inspired emotion-based architectures for designing and creating ever more adaptive agents that will perform adaptively in dynamic environments. In this case we argue our choice of studying the predator-prey relationship allows further insights to be obtained as to how an emotion-based architecture might impact such a relationship between artificial agents, as well as to see how we might usefully measure agent behaviour in such contexts (the predator-prey one being a context that is very interesting, for many reasons, in its own right).

### 1.4.2 System Properties generated by B-B-E Interactions

However, perhaps one of the most distinctive features of the research outlined in this thesis, and which therefore separates it most from the existing research, is that in our experiments we adopt a bottom-up approach to model and study action selection and emotion as emergent properties of systems, generated by interactions of brain, body and environment.

# Chapter 2

# Animat Brains, Bodies and Environments

*"Intelligence is determined by the dynamics of the interaction with the world."*

**— Rodney Brooks**

This chapter begins with a brief summary of the animat approach, outlining some of the specific contributions it aims to make (and is making) to both Artificial Intelligence (AI) and Artificial Life (AL). Moving on, we then use this to provide a context for our work. That is to say, to set the scene in terms of existing literature. In this thesis, we choose not only to adopt the animat approach as part of our research methodology, but to use it to develop a framework for categorising literature related to our research questions. In relating elements of this approach to those concepts we are interested in studying together, we focus on that research which is particularly relevant to the study and understanding of intelligence with regards to adaptive behaviour. This naturally includes work relating to biological, as well as artificial, agents. To do this, we use the concepts of animat brains, bodies and environments to help both categorise and bring together several different, in some cases previously separated, strands of research.

Examples are given of some of the most relevant work carried out to date, relating to our own research interests. With regards to the problem of action selection (how an agent decides "what to do next") and the brain-body-environment relationship, the focus here is initially placed on work which has manipulated the components (brain, body and environment) and interactions of the latter so as to create more autonomous and adaptive systems for solving the former. Recent ideas about the importance of aspects of embodiment for the development and study of adaptive agents, as well as the possible importance or influence of the environment, are therefore included. We then discuss research which has involved incorporating mechanisms mimicking one of the suggested roles of emotions in humans (with regards to decision-making). Specifically, we draw together a branch of the more recent research on action selection to focus on the promise of emotion-based action selection mechanisms, as well as questions about the role of the body, for generating adaptive behaviour in artificial agents.

A particularly interesting action selection scenario from the perspective of the role of emotion in decision-making is the case of predator-prey interactions, which involves threats to the survival of agents in a *dynamic* environment. This chapter concludes with a more in-depth discussion of the predator-prey relationship and its study. As a crucial and defining feature of the environment we also choose to focus on in our research, we identify some of the main advantages of studying this relationship, focusing on it as an action selection problem that is ideal for developing models of artificial emotion. Our discussion thereby enables consideration of the possible usefulness of the predator-prey environment in helping animat researchers to further investigate affect-modulated action selection in the context of brain-body-environment interactions. Again, we consider the animat approach the ideal choice of approach for exploring our ideas in this context.

## 2.1 The Animat Approach

The animat approach is a behaviour-based approach to Artificial Intelligence (AI) for the design, creation and study of autonomous and adaptive agents [30]. The behaviour-based approach to AI is one (the advent of which is most often attributed to Brooks' [31] early research with robotic agents and demonstrated in his proposal of the subsumption architecture [32]) that, in essence, sees intelligence as composed of a large number of modular elements that are simple to design and which can be developed in layers. As opposed to the more *deliberative* intelligence traditionally studied in AI, it is thus more associated with and focused on producing more *reactive* intelligence. At the time it was introduced, this approach was in many respects revolutionary, overcoming limitations of the more traditional approaches. For example, by demonstrating greater robustness and practical applications for real-time systems such as robots [33; 34].

The animat approach is also numbered among those developed as alternatives to earlier approaches when the initial 1950s promise of symbolic reasoning started to fade. As such, its history dates back as far as the mid-80s, when this new kind of bottom-up type of approach was first adopted and advocated by researchers including not only Brooks [13; 31; 33; 34], but others such as Wilson (cited in [35]) Meyer and Guillot [36]. As a paradigm, the animat approach takes a synthetic approach to the study and creation of autonomous agents, as well as the creation of artificial intelligence, which more than a decade later Dean [30] described as an "attempt to comprehend the capacity of animals for autonomous generation of adaptive, intelligent behaviour in complex, changing environments".

In essence, this approach argues that the best route to approaching the complexities of human intelligence, in terms of both simulating and understanding it, is to build and study (which will hopefully lead to the understanding of) simple but complete creatures first. Key arguments or views of the animat approach, as nicely outlined by Dean [30] include:

- Adaptive behaviour is best understood by focusing on the interaction between an agent and its environment, hence an interest in embodied agents situated in natural environments
- Specific abilities or behaviours are more natural units of analysis and design than general, information-processing functions and world models
- High-level behaviours will emerge as systems composed of simple behavioural competences become more complex

In this way, animat research begins with the design of agents with low-level abilities, with the aim of later moving up towards higher, cognitive functions. Yet Dean also reflects on some of the arguments of its detractors, whose doubts include:

- Whether animat research will contribute to a general theory of situated agents and environments
- Whether animats can demonstrate the validity of the assumption that high-level competence will emerge naturally and in reasonable time when complexity is scaled up
- Whether the ultimate explanatory value of research with animats will be enough for connecting to information-processing accounts of higher, cognitive functions

The animat approach naturally encourages the use of biologically-inspired control mechanisms for creating autonomous and/or adaptive simulated animals or robots [37]. Such animats, and the models which use them, can act as tools for the study of adaptive behaviour, while also providing a means for considering intelligence, adaptation and perception at higher levels. From a broader perspective, this approach thereby attempts to better address some of the main criticisms directed towards more traditional approaches, including those of situatedness and embodiment [38]. In particular, it does this by considering the importance of agent-environment interactions [39].

Perhaps somewhat conversely though, the approach also resonates strongly with the idea that, in order to understand and reproduce human-level intelligence (however we might define this) we should first try to understand and reproduce the probable roots of this intelligence. Because its main assumption is that we can develop more complex animats by building on more simple ones, this places initial emphasis on the basic adaptive capacities of animals that have to deal with the full complexity of autonomous interaction with their environment. For example, to concentrate first on studying simpler abilities like navigating, seeking food and avoiding dangers [40].

Both in analysis and design, the animat approach takes its inspiration from the work of ethologists, psychologists, neurobiologists and evolutionary biologists, as well as connectionists [30]. For AI and robotic researchers, the primary research goal is often the production of autonomous and adaptive behaviour. However, it is also hoped that an increased understanding of the possible mechanisms able to generate adaptive behaviour will lead to the discovery of tools and concepts which might further aid natural scientists in their attempts to understand biological systems [41].

Researchers looking to study adaptive behaviour and autonomous agents in general have included Steels [42; 43; 44], Maes [45] and Meyer [40; 46], who each identified some of the advantages of this approach over the classical, more top-down approaches, yet also offer criticisms of the newer approach. For example, identifying various open questions still to be addressed, whilst also trying to more adequately attempt to define what it means for an agent to be both autonomous and adaptive (if for no other reason than to use as a yardstick for assessing how close animats have come to being both so far).

More than ten years later, while a search of the literature for the work of those researchers who have since adopted and used the term “animat” alone to describe their approach and/or robots may only generate a relatively small selection of research papers, it is nevertheless possible to identify many examples of what can be considered the animat approach/animats under other headings. Such examples include, but are not limited to, many found scattered through the literature of artificial ethology [2], embodied AI [31; 47] and bio-robotics [48]. In this way, although the term itself is less often used, in as much as it has perhaps not caught on enough for the widespread use of the term animat to be applied whenever the approach is adopted, the general approach may actually be more frequently identified within the literature.[1]

## 2.2 Animat Brain, Body and Environment

One open question in particular in the study of autonomous and adaptive agents, for animat research and the field of AI in general was identified by Meyer [40], who suggests that the dynamics of interactions between agent and environment are still not

---

[1] Having said this, a simple search at the time of writing this returned more than 12,000 results for the term “animat”. Out of interest, when these results are seen by decade we see several thousand results that do not relate to the use of this term in this context at all. Yet between 1990 and 2000 there are 180 results for the term “animat” and “action selection” combined. Between 2001 and 2010 there are a further 327 results, 33 of which are found to have been published in 2010 (14 in 2009 - one of which was written as a result of the work in this thesis.).

well understood. At the time, Maes [45] identified Beer [49], Kiss [50] and Steels [51], among others, as numbered among those having started to try approaching the problem using a dynamical systems perspective but, said Meyer, the field is “far from being able to prove in general what the emergent behaviour is of a distributed network of competence modules”. Though to-date there might still be much about this we do not understand, the animat approach nevertheless provides a means of exploring such interactions more thoroughly. Moreover, it gives us a different perspective than more top-down approaches, allowing us to experiment more with ideas about the evolution of our own intelligence.

### 2.2.1 Action Selection Mechanisms as Animat Brains

As an open question itself in AI, action selection is a problem of considerable interest to animat researchers. The problem of action selection is that of making a decision as to what behaviour to execute next so as to carry out conflicting tasks and guarantee survival in a given environment and/or situation. Researchers, including Maes [45] and Avila-García [29] see it as the problem of how to choose between competing behavioural alternatives so as to enable an agent to survive in dynamic, unpredictable and hazardous environments. In his work, Avila-García suggests that the main problem in this area is to build architectures that will result in an agent demonstrating “adaptive, robust and effective behaviour”. He argues this will be achieved by both solving the problem of action selection — specifically, by using the correct *arbitration mechanism* — and giving agents the ability to learn from experience. This follows quite closely the arguments of earlier researchers such as Meyer [36; 40; 46] and Steels [42; 43; 44].

Action selection mechanisms can be seen as a collection of tools that are used in an attempt to solve the problem of action selection within AI. That is, to allow agents to choose what to do next at a given time. The argument often underlying and used to justify the work of those that concentrate their attention purely on this problem is that if we can develop the right action selection mechanism, it will be possible to produce agents that behave both adaptively and autonomously in their environment [45]. (Though this also implies that the action selection mechanism will be sufficient in itself to guarantee adaptive and autonomous behaviour — something which is still debatable)

Outside of the animat approach, as well as within its research literature, a range of different types of action selection mechanisms have been identified and designed. Yet often these can be characterised as belonging to one of a few main types:

- Classical or symbolic action selection mechanisms, covering those which use more GOFAI[1] symbol-based reasoning

---

[1] Good Old Fashioned AI

- Distributed or connectionist action selection mechanisms, including those using neural networks as animat controllers
- Dynamic control systems, including action selection architectures
- Hybrids of more than one of these types

Animat researchers often try to compare the properties that different action selection mechanisms display so as to see what type(s) of controller or action selection mechanism perform(s) best under different conditions. For example, Tyrrell [14] compared several (ethologically-inspired) action selection proposals/architectures, finding that some performed better than others in his simulated environment.

Other action (or behaviour) selection mechanisms that have been proposed include those outlined by Bryson [52], Guillot and Meyer [53] and Pirjanian [54]. They look at some of the many different action selection architectures that have been successfully developed and tested using both the animat approach and others. As Avila-García [29] notes, such designers as are mentioned can often be seen to differ along three main dimensions in their proposed architectures: “winner-takes-all” versus “voting-based”, “hierarchical” versus “non-hierarchical”, and “reactive” versus “motivation-based”.

Because of the more biological-like focus of the animat approach, animat researchers have often similarly argued for and taken a correspondingly more biologically-inspired approach to the development of action selection mechanisms [26; 55; 56; 57; 58; 59; 60]. Indeed, many researchers effectively treat the action selection mechanism as something like the equivalent of a “robot brain”. Consequently, many also attempt to take what has been discovered about the organic brain by scientists from biology, neuroscience and other similarly biologically-centred disciplines, trying to apply it in order to better develop action selection mechanisms/architectures [58; 59].

This follows the argument that the best and most complex example of action selection mechanisms can be seen in the real world, provided by human and/or animal brains, which can therefore be used, as working examples, so as to develop efficient, working artificial action-selection mechanism(s) for action in the real world. However, perhaps more interestingly, it also makes a case for the argument that the idea of an action selection mechanism being the equivalent of an animat’s brain may be useful in exploring further the possible roles of brain, body and environment, as well as that of their interactions, in generating/designing for adaptive and autonomous behaviour.

### 2.2.2 Bodies and their Importance to Animat Brains

In contrast to those researchers that have focused on the development of controllers for generating adaptive and autonomous agents, another relatively more recent branch of literature has argued against the advisability of focusing purely on such “brains” at the risk of forgetting all about the role of *embodiment* [61]. This includes the body of an agent itself [62]. As Chiel and Beer [21] mention "mechanisms of adaptive behaviour generally focus on neurons and circuits. But adaptive behaviour also depends on interactions among... nervous system, body and environment... body structure... [creating] constraints and opportunities for neural control; and continuous feedback between nervous system, body and environment... essential for normal behaviour." Thus, "one cannot simply 'peel away' the body to understand the nervous system's role in adaptive behaviour".

This broader perspective on adaptive behaviour can also be identified as that underlying ecological psychology, a discipline which has influenced behaviour-based robotics as well as the field of computational neuroethology, which models both neural control and the periphery of animals, and is praised for its position as a promising methodology for the understanding of adaptive behaviour. Other researchers advocating greater consideration of the "body" in the study of intelligence, as well as the subsequent creation of adaptive and autonomous agents include Pfeifer and Scheier ("Understanding Intelligence" [63]) and Pfeifer and Bongard ("How the body shapes the way we think" [62]). This has implications for those studying action selection mechanisms, in that it suggests that adaptive and autonomous agents may best be seen when a researcher manipulates *both* the brain and body component of an agent in order to suit its environment — rather than just the brain. Thus, perhaps better action selection mechanisms might be developed if they were to encompass a body too. (Though we agree it is also possible to see the body simply as a means of solving some of the problems of information-processing, effectively moving such problems "upstream" by designing bodies that will filter out unnecessary information and/or perform some kind of computation.)

Following this, the body of an animat can be thought of as the communication channel between brain and environment: an animat being composed of both brain and body. The animat approach is complementary in this respect as it too forces the designer to consider the body design of an animat beforehand. Yet we must also consider what is covered by the term "body". Taking an animat perspective, perhaps it is best defined as the entire material of physical structure of a given animat. Though this does not mean both an "internal" and "external" body cannot be further identified. For example, researchers might (and do) define internal 'physiological variables' and/or "drives" as is the case with certain motivation-based architectures [60]. An interesting area of research which aims to study more closely how brain and body interact has been more recently termed by Parisi [24; 64] *internal robotics*.

### 2.2.3 Environments and their Dynamics

Another player in the animat approach, though it seems its role is often left underemphasised in the study of action selection mechanisms, is that of the environment of the agent itself. A logical argument follows that, if a criticism that arose from researchers focusing purely on attempting to synthesize some sort of robot "brain" for generating autonomous and adaptive behaviour was that it forgot about the role of the "body"; then focusing purely on the agent (brain plus body) may suffer similar criticism in the way it thereby under-emphasises and devalues the role of the environment in developing interesting animats that make use of better action selection mechanisms. One idea might be for researchers to consider the environment itself as an agent. However, more usefully perhaps, we might also consider more the ideas of embodied, embedded cognition (EEC). Such ideas namely revolve around the premise that, to borrow a phrase from the biologist Gerald Edelman , "the brain is embodied and the body is embedded in its environment. That trio must operate in an integrated way." [65].

Real-world environments can be characterized by many factors, including temperature, precipitation, slope or gradient, sunlight, chemicals and sounds. Though organisms are also considered to have an internal environment, including factors such as nutrients and body temperature, this can be sub-divided using the concept of brain and body (internal). For agents studied in computer science in general, however, their worlds (environments) are often necessarily designed to be much more limited. Often the designer is creating a very specific, constrained and artificial ecological niche to see how agents will manage when placed within it. Attempts have been made to classify such artificial environments. For example, Todd and Wilson [35; 66] started to build a taxonomy of environments as well as a taxonomy of agents that would provide a basis for comparing different action selection mechanisms. However, for our purposes, we have found it more useful to see environments as primarily varying in their level of complexity and/or dynamics.

The "environments" that have been developed (or otherwise singled out) for the study of both action selection and (indirectly or not) the brain-body-environment relationship have varied: from the simulated and created gridworlds of Tyrrell [14], Cañamero [67] and Sutton [68]; Wilson's woods (cited in [14; 69]) and McCallum's maze (cited in [70]); the physically embodied, but more constrained worlds of those such as Avila-García [29; 71]; to the physically embodied, continuous-action robotic environments (or niches) agents such as Genghis [72] have been built to exist in. Physically embodied agents and environments are particularly interesting as agents such as robots are taking, as their environment, a sub-set of the real world as we know it directly — in other words overlapping their environment more directly, if not directly with ours.

Environments may therefore be simulated or embodied, static or non-static (as is the case with dynamic environments). Interestingly, researchers such as Miglino, Lund and Nolfi [73] have also provided evidence that it is possible to build an accurate model of robot-environment dynamics by sampling the real world through the sensors and actuators of a robot (indicating that the study of simulated and real environments need not necessarily be mutually exclusive). The ultimate goal for robotics is for a robot to be able to deal with environments as dynamic as the ones we live in. Thus, a dynamic environment should be of particular interest to those who study action selection and the brain-body-environment relationship. For, in order to develop robots that will one day exist and function adaptively in our environment, we will first have to find ways to allow such agents to cope with and adjust to its complex dynamics. Tyrrell's [14] simulated environment, though he used only one environment and task, is a good example of a dynamic environment that introduces many sources of dynamics into its environment: from consumable resources to moving predators. There are also some researchers who have attempted to incorporate more indeterminate stimuli such as night/day cycles, such as Cañamero [67].

Among others, Maes [45] and Wilson [66] argue that it is not possible to decide whether one action selection model is better than another unless the characteristics of the environment are also considered, along with the task and the agent. Maes gives the example that, in an environment where the cost of making the wrong decision is high, an agent should do more anticipation; while if the cost of making the wrong decision is negligible, it does not matter if an agent often performs incorrect actions. Moreover, in an environment where changes may occur quickly, an agent needs to act very quickly. Another related example is that an agent with noisy sensors should have some inertia in its action selection so that one wrong sensor reading does not make the agent switch to doing something completely new and different. Maes [45] notes that an agent with many sensors can rely on the environment to guide its selection of actions, whereas an agent with fewer sensors will "need to rely more on its internal state or memory to decide what to do next".

Wilson, Somayaji and Yanco [74] go so far as to argue that adaptive behaviour is in fact possible without any knowledge of the environment, at least, over an animat's lifetime. Yet in their experiments they used *evolution* to evolve controllers for their animats, which *were* "blind" in that they did not have any direct contact with the environment — their agents were not given any sensors. However, we would argue that they *were not* truly "blind" in that the process of evolution itself was acting as a more indirect, cross-generational sensor of the environment. Thus, evolution was effectively used by them as a way to avoid using sensors directly. In this way, the environment created the adaptive behaviour because the adaptive behaviour would not have evolved without the environment. That is to say, agent and environment were still interacting, just not as directly as if direct sensors were used.

While their work also fits in with the earlier results of Todd and Wilson [66] (they characterised important features of the environment structure in terms of the adaptive behaviour the elicited) Wilson mentions the fact that, in many systems "multiple creatures with different behaviours never interact in a common world, and as a result much richness of both social behaviour and the effects of whole populations on a shared environment is left out" concluding "creatures need not be able to sense nor remember anything about their world (or themselves) in order to behave adaptively in it, provided the world is generous and benign enough". Floreano's [75] shared environment research examines the differences in behaviour that occur when creatures can or cannot sense each other.

### 2.2.4 The Brain-Body-Environment Relationship: Interactions and Manipulation

Seth [11] suggests "animals do not choose between behaviours per se; rather, behaviour reflects interactions among brains, bodies and environments." One of the useful advantages of the animat approach is that it allows researchers to take a broader, often more abstract look at the relationship and interactions between brain, body and environment under different conditions (that is, provided it is allowed that action selection mechanisms represent some form of brain-equivalent — and similarly allowing for the animat's physical, or simulated, body plus its world to be its environment). Indeed, one of the many hopes of the approach is that it is via interaction dynamics (between these components) that emergent complexity will arise; another being that looking at complete systems will change problems themselves, ideally in a favourable way. This is particularly relevant considering the biologically-inspired aspect of the approach as, in real life, there is an unarguably intimate connection between our brains, bodies and environments — and many scientific disciplines exist to argue this point.

However, still much of the research within AI has focused on one or more components in the brain-body-environment relationship, without necessarily looking at the "bigger picture". Though there is one recent branch of research that has started to focus on looking at all three of these components together [27], it does so using evolutionary algorithms, which often restricts the study to the use of simulations. Perhaps this is why there is not much evidence of systematic studies of the brain-body-environment relationship using real robots. Furthermore, more often than not one aspect of the relationship seems to be neglected — perhaps none more so than that of the environment. Looking at the research, it is in fact very tempting to suggest that the brain-body-environment relationship itself makes a good conceptual framework to use for categorising the work so far. For example, experiments can change or manipulate:

- Brain: for instance, using different controllers/action-selection mechanisms, varying in form from architectures to neural networks
- Body: such as those using different hardware components (manipulating external body) or different physiological variables in a motivation-based system (manipulating internal body)
- Environment: from completely static (such as Avila-García's Two Resource Problem [29]) to increasingly more dynamic (for example, Cañamero's [67] and

Tyrrell's [14] gridworld simulations[1])

Some researchers have chosen to study or manipulate just one of these components. To give an example, Tyrrell [14] studied the performance of different architectures in the same body and environment (though it is possible to argue that the environment is in essence changed, because a different architecture will introduce different dynamics into the environment in return). However, others have varied two or, more rarely, all three together — for instance, Avila-García [29] used the "same" body, with different architectures and environments, while Buason, Bergfeldt and Ziemke [27] chose to vary all three components using evolutionary algorithms to produce different controllers (brains) morphologies (bodies) and environments.

The most interesting research to us, however, is that which not only considers how these components themselves might be changed, but simultaneously attempts to study how interactions between these components can be manipulated or controlled so as to produce adaptive behaviour. In this way, resulting experiments might focus on the study of:

- Brain-Body interactions: for instance, by the addition of feedback mechanisms including hormone-like and/or emotion-like mechanisms, learning mechanisms or other mechanisms such as those as studied in Internal Robotics (bridging the gap between brain and body) [64]
- Brain-Environment interactions: for example, using different bodies or sensor placements [76; 77]
- Body-Environment interactions: including developing more dynamic environments, different bodies, two-way interactions and the ability for agent to manipulate aspects of the environment [78][2]

Again, a researcher may choose to study or manipulate one or more types of these interactions. However, an idea that seems at times to be noticeably avoided or disregarded, particularly when evolution is introduced from such studies, is that of niche construction. This term refers to the consideration that the agent-environment interaction does in fact go two ways. First adopted and defined in 2003 by Odlin-Smee et al. [79] it relates to similar terms and concepts that can be traced further back, driving research using such concepts as stigmergy. As such, niche construction refers to the "neglected" processes used by an agent to manipulate or create its *own* environment and environmental dynamics. Indeed, this idea seems particularly relevant for the fact that it emphasises an aspect of the brain-body-environment relationship we

---

[1] Though we are not forgetting that an agent can itself be considered a dynamic of its environment

[2] NB. Brain-Body-Environment and Body-Environment interactions might alternatively and collectively be thought of as Agent-Environment interactions

have already identified here as being often neglected within the study of AI. Moreover, it further illustrates one of the problems with more traditional approaches and/or those using evolutionary algorithms, being an idea that may have significant implications for evolutionary computer scientists.

### 2.2.5 Mechanisms for Action Selection and their Effect on B-B-E Relationship

For the animat approach, with the focus often being on the controller of an agent, it can be useful to see the action selection mechanism as the "core" mechanism around which other mechanisms may be created to act on and influence. One interesting branch of the research on the action selection problem, as well as the brain-body-environment relationship, can thus be categorised as focusing on improving existing action selection mechanisms by the incorporation or design of mechanisms which themselves (as mentioned) can be seen as manipulators of the brain-body interaction (and thus the brain-body-environment one). The use of such mechanisms might involve the incorporation of any kind of perceptual or other mechanism that will bring the agent information from the outside environment: whether from simple sensor addition such as touch sensors for detecting an obstacle or features of its world, to the use of more complex algorithms to obtain or filter sensory information for use in action selection.

It can be argued that such perceptual mechanisms manipulate the body-environment interaction. This is because they lead to different perceptions of the outside environment being received by an agent. Examples of other types of mechanisms include those imitating some aspect of learning, memory and/or other higher-level functions. Without getting side-tracked into a debate of psychological theories of human emotion, interesting developments have further been made through the introduction of emotion-like mechanisms or components to existing and/or purposely-designed action selection mechanisms. This includes the subset of action selection mechanisms we call in this thesis "emotion-based". Leaving aside the notion of emotion in its entirety, such research is biologically-inspired in that it recognises that one of the likely functions of emotion in humans is to produce adaptive behaviour in a dynamic environment (though it is also interesting to note that adaptive behaviour itself is usually dynamic behaviour).

Various mechanisms and functions of (artificial) emotions have been suggested in relation to autonomous agents and various agents, both simulated and robotic, have subsequently been designed and implemented for study. Bringing the literature together in this area, a good analysis using a selection of existing agent architectures that include an emotion mechanism has been provided by Rumbell et al [80] who have recently attempted to address the difficulty of making comparisons between such agents by proposing a set of architectural qualities as a basis for comparison. The architectural qualities they suggest are the agent architecture, action selection mechanism, emotion mechanism and the emotion state representation, along with the model it is based on.

To give a few examples of the work carried out to-date by those researchers who have chosen to adopt a bottom-up approach to the study of (various mechanisms and functions of) emotion in artificial agents, in the work of researchers such as Gadanho and Hallam [81] we have seen the development of a non-symbolic emotion model in the form of a recurrent artificial neural network. Key to their work is the idea of reinforcement learning and emotions. That is, they explore what they consider to be the most usual role attributed to emotions in the functionality of an artificial learning agent — providing a source of context evaluation or reinforcement. In contrast, Delgado Mata and Aylett [82] examine a slightly different role for artificial emotion (with an emphasis on the specific emotion of fear). Focusing on communication of emotion between agents, they have developed an ethologically-inspired action selection mechanism which integrates emotion with flocking and grazing behaviour in virtual animals. Their suggestion is that emotion reduces the complexity of behaviour, thereby acting as mediator between individual and collective behaviour. More recently, Burattini and Rossi [83] have provided another example of this approach, suggesting that the introduction of emotion as "timed controlled structures" may lead to an adaptation in emergent behaviour without having an explicit action selection mechanism. They have related the emergent properties of their (robotic) architecture to a substrate for emotional processes, again intended as a bottom-up influence on perceptual capabilities and action selection. Their conclusion is that this allows for more flexible/adaptive behaviour, such that their agent is able to react and adapt in real time.

Various researchers, including Cañamero [67] and Avila-García [29], suggest hormone-like mechanisms are likely to be of increasingly more importance and useful for agents operating in more dynamic environments, allowing for temporal patterns of behaviours to be adjusted, often in response to stimuli detected in the environment. Mendao's [84] work on an architecture based on the role of biological hormones in the neural system provides further support for the idea that such mechanisms give us the potential to achieve complex behaviours through the interaction of very simplistic structures. However, he suggests that as the number of behaviours grows so too would the amount of design required to achieve meaningful emergent behaviour.

This research is also supported by the work of researchers in other disciplines, including neuroscientists, such as Lazarus [85], Damasio [86] and Fellous [19]. This is reflected more recently, with Tanguy, Willis and Bryson [17] suggesting the use of emotions as "durative dynamic state(s) for action selection". In this way, perhaps a better solution is to see perception as being the information obtained from the outside world/environment; perceptual mechanisms relating to how this information is obtained; and cognition as more of an emergent property (which seems more in keeping with the nature of the animat approach itself anyway). Thus emotion or hormone-like mechanisms may be considered mechanisms influencing perception, which therefore create or help to create a system. Indeed, this thinking could also be applied to the action selection mechanism itself. That is to say, an action selection mechanism could be the emergent property of an arbitration mechanism, architecture, body and environment acting/interacting together.

## 2.3 The Prey-Predator Relationship

The relationship between predator and prey is one that should be of particular interest to those studying action selection [14; 79; 87]. Indeed, it is of interest across and within many disciplines. While there are many aspects of this scenario to interest researchers, what often stands out is the fact it is a relationship between two agents. Moreover, it is a relationship characterised by a dependency of one agent (the predator) on another (the prey) for its continued survival. This results in interactions between agents that will determine the success of each agent, with a push-pull effect. Where one wins, the other will likely suffer some corresponding cost or loss.

Looking at the literature, research has explored this scenario from various perspectives: from the level of the individual over a lifetime to populations across generations. Yet the way this relationship has most often been studied is through the development of action selection mechanisms for the prey that will result in it fleeing whenever it sees a predator, thereby making immediate flight the more or less automatically optimal or decided choice of action, regardless of the task currently being performed. Strangely, researchers have also commonly continued to focus on one type of agent only (predator or prey) with the action selection problem of the other agent being of secondary to no interest. We regard this as possibly leading to a more superficial look at, or treatment of, the action selection problem for artificial predators and prey. This could encourage researchers to adopt a perspective which may lead to less rich, or realistic, solutions than might be the case or useful in real life and real time.

For example, this emphasis does not take into account or allow for the possibility that in fact there may be times in which the more adaptive behaviour would be for the prey to “take the risk” of being attacked by its predator. Or, indeed, the case that there are some, if not many, environments in which life must constantly be risked in order to achieve long-term survival. Perhaps in favour of satisfying some other survival need or task. Looking towards ethological studies for evidence and inspiration, researchers illustrate this could also be true of biological organisms.

For instance, Cooper [88] found a species of lizard will tolerate predators to come closer before they decide to “flee” under certain conditions, including when they were eating food. Though it could be argued this might also reflect the possibility that the lizard’s attention is more directed on feeding than awareness of or perception of the predator. More interestingly, it could be that some kind of economic model allows for “risk-taking” or a kind of “cost-benefit” analysis in terms of risk assessment that is adaptive for agents. Then too, this could lead to a role for emotion-like states as quick, real-time assessors of risk in relation to certain stimuli.

The predator-prey relationship is an interesting and relatively popular action selection problem, as well as a general scenario for study within AI. For example, it has been explored, using various techniques, by researchers such as Floreano and Nolfi [89]; Buason and Ziemke [4]; Avila-García [29]; Cañamero [67]; Tyrrell [14]; Kelly, Holland, Scull and McFarland [87]. While certainly interesting in its own right, perhaps what makes it particularly useful for studying the problems of action selection is that the introduction of some kind of "predator" will immediately introduce dynamics to a previously static environment, thereby increasing the "pressure" on the prey (or, more specifically, its action selection mechanism) to make the right decision; while also allowing the researcher to explore brain-body-environment interactions further.

As an action selection problem, the predator-prey scenario for the prey mostly consists of having to choose how to respond to a predator so as to survive, whilst not neglecting other essential tasks necessary for its survival. While many researchers have focused on looking at problems of action selection through studying the prey (Avila-García [29], Tyrrell [14]) others have also attempted to consider the predator and its action selection, often by taking a broader perspective. Usually, this involves introducing some kind of evolutionary component so as to simulate some kind of "competitive co-evolution" of/between both predator and prey (For examples, see [4; 89]). What is particularly interesting about this latter approach, however, is that it also enables its researchers to study further the interplay between brains, bodies and environments within the predator-prey scenario.

Some studies [29; 67] have also found, through manipulation of brain-body interactions, action selection and therefore performance of the prey can be improved. For example, via hormone-like mechanisms which release a "hormone" when the presence of a predator is detected. This, in turn, affects interactions between brain and body such as perception (in the case of [29] the hormone-like mechanism acting as a second-order controller to the existing action selection mechanism). This is an interesting development in that it further supports the idea that, by making interactions between brain and body more dynamic, whilst also connecting/rooting this to relevant environmental features (in this case, the presence of the predator) we can produce more adaptive behaviour in more dynamic environments.

Intriguingly, different types of predator-prey relationship may also be seen to arise from the use of different action selection mechanisms, morphologies and environments. Or, in other words, brains, bodies and environments (as demonstrated by Buason and Ziemke [4], for example). This is perhaps reflective of the different types of predator-prey relationship discovered in nature: from the more symbiotic relationship, to the parasitic, to the more traditional life-or-death situation once predator encounters prey. In this respect, continuing research in this area may thus be beneficial not only for what it can tell us about the problems of action selection (for example, in producing adaptive behaviour in dynamic environments) but also for what it can tell us about the dynamics of the predator-prey relationship itself, particularly in the context of the brain-body-environment relationship. That is to say, we want to know how the predator-prey relationship is influenced by and can influence the brain-body-environment relationship. But we also want to see how this can, is and/or might be manipulated.

With regards to niche construction in particular, the predator-prey scenario may also be useful in exploring further this "neglected" aspect of the brain-body-environment relationship. For example, by manipulating a prey agent in ways that will change its behaviour, the environment of a predator agent is also likely to change, which may mean it will also respond differently and/or need to adapt in different ways. This effectively means that prey has some freedom/control over creating its own niche (and vice versa). Or, in other words, the prey can partially influence and shape its own predator-prey relationship. Indeed, when considering real life scenarios, with entire ecosystems having evolved together, it seems likely that such studies will help lead us further to generate ideas as to how such predator-prey relationships as do exist may have actually come to exist. But perhaps more importantly, how each agent adapts to the dynamic behaviour of the other: thereby also allowing us to consider the possible contribution(s) of perceptual mechanisms to niche construction.

What will be interesting to discover is whether more "adaptive" action selection mechanisms will be those that attempt to and keep brain-body-environment relationships stable; or whether they will be those that attempt to actively modify the brain-body-environment relationship itself. That is to say, in the agents' own self-interest. Perhaps this may also be related to the ideas of autonomy/automaticity (such as Steels [42; 43; 44]). For instance, an action selection that attempts to keep the brain-body-relationship stable might be considered an automatic system, while an action selection that successfully (adaptively) manipulates the brain-body-environment might be considered an autonomous one.

In summary, the predator-prey relationship may thus be of interest for action selection researchers for many reasons. However, for us, among the most interesting are:

- Adding a predator (or prey) to a given agent's environment is a way of making that environment dynamic. It leads to changes over time that the agent must respond to adaptively and often increases environmental complexity. Thus, in terms of action selection, it can act as a good test for how well an individual agent (or the action selection mechanism implemented within it) can cope with increases in the dynamics of their environment. Importantly, the nature of these dynamics are typically such that each agent has to make quick decisions in order to make adaptive ones. This leads to a trade-off, where if the agent hesitates or ponders too long, all could be lost anyway (game over, especially for the prey).

- It allows us to study action selection at a higher or more general level, within the context of two agents in a very unique relationship. Typically, one in which, where one agent wins, the other will invariably lose. This may affect the demands for (and guide the design of) the agents and action selection mechanisms themselves, especially as the relationship is characterised by a dependency of one on the other — in as much as predator is dependent on prey. Admittedly, prey might also be said to be dependent on predator. For instance, at the population level, to avoid over-population. Yet such dependency is likely to be much more indirect. This thereby makes the balance of opportunity cost and stakes for each agent in any interaction unequal. Where predator loses a meal, prey loses its life.

- It provides us with (if nothing else a wealth of biological) inspiration for building action selection mechanisms both a) capable of dealing with situations of high and immediate risk (used by prey) and b) capable of adapting to another agent's behaviour (environmental dynamics) for the agent's own advantage (used by prey and predator). It is also a problem that may call for compromises, increasingly specialised or more adaptive behaviours and, more specifically for us, interesting trade-offs. Namely, between the basic choices for the prey of whether it should flee or not, and for the predator of whether it should attack/hunt or not. Somehow, these agents must be able to effectively weigh up and make these decisions in the limited time available.

- It allows us to focus on the interactions that result between (the action selection mechanisms of) two agents with different sensory abilities, brains, bodies, motivations, possibly emotions (especially at the time of interaction) and behavioural repertoires. Starting our own "arms race" between such agents, we can develop and fine-tune features of these agents to enable one to gain an advantage over the other. This could not only produce and drive the production of increasingly more adaptive agents, but also lead to a better understanding of the (different types of) predator-prey relationship(s), as well as the circumstances when certain components of action selection mechanisms might be most adaptive.

- It allows us to look in more detail at the requirements for adaptive behaviour in this context. For example, it allows us to ask whether a predator needs more "brain power" than its prey in order to be able to catch it, or simply different types of behaviours and abilities. Similarly, it allows us to explore those ways in which we might increase or examine the adaptive value of predator and prey action selection mechanisms. This could include the use of methods across disciplines. For instance, we might analyse developed prey agents' behaviour in a similar way to Cooper's lizards: in terms of the assessments of risk or cost-benefit analyses that he suggests can be used to explain their behaviour.

## 2.4 So where does this Thesis fit in?

For an agent, the problem of action selection can be defined as consisting of deciding what behaviour to execute next, so as to carry out conflicting tasks and guarantee survival in a given environment and/or situation [45] which is often dynamic, unpredictable, and hazardous. According to Maes [45] the main task in this area of study is to build architectures or action selection mechanisms that will result in an agent demonstrating "adaptive, robust and effective behaviour". Indeed, many artificial action selection mechanisms have been proposed in the literature, and their researchers often try to compare the properties that different mechanisms display in order to see what type(s) of controller (as, in effect, examples of different animat "brains") perform(s) best under different conditions [40].

However, in contrast to research focusing on the development of controllers for generating adaptive and autonomous agents, another, relatively more recent branch of literature argues against the advisability of focusing purely on such "brains" at the risk of forgetting all about the role of embodiment. This includes the body of an agent itself. For, as Chiel and Beer [21] mention, "mechanisms of adaptive behaviour generally focus on neurons and circuits. But adaptive behaviour also depends on interactions among... nervous system, body and environment... " thus "one cannot simply 'peel away' the body to understand the nervous system's role in adaptive behaviour".

This argument has implications for action selection mechanisms in that it suggests that adaptive and autonomous agents may best be understood when investigating the relation between "brain" and "body". For example, manipulating interactions both between these components, as well as their combined interactions with aspects of an environment, rather than just the "brain" (as represented by the action selection mechanism). In this thesis we too claim that better, more adaptive action selection mechanisms would be developed if they were to encompass bodies. Furthermore, these ideas may have important consequences for the study of the role of emotions in action selection and decision making in general, encouraging researchers to take an "embodied emotion" approach [90; 91] to model the effect of emotions on cognition and behaviour through the body. Related to this, Parisi termed "internal robotics" the area of research which aims to study the interactions of brain and body more closely [24].

Driven by our interest in making our own contribution to the research outlined in this chapter — and because we believe a more systematic study of the predator-prey scenario will lead to the development of better, more adaptive, emotion-based architectures — in our research we have been and are currently developing both a model (and associated robotic implementation) of brain-body-environment interactions in a predator-prey scenario. In doing so, we have chosen to study a specific type of architecture so as to develop insights and explore more general concepts related to the adaptive value of emotion in dynamic environments. The general framework within which we hope to formulate our concrete research question(s) can be summarised thus: Under what conditions will a mechanism simulating neuromodulatory effects (as a proposed substrate of emotion) prove adaptive for agents, in terms of action selection, emotion and brain-body-environment interactions in a predator-prey scenario?

This also equates to: What factors will affect the adaptive value of a mechanism simulating neuromodulation, as a proposed substrate of emotion and biasor of action selection, in a predator-prey scenario? We are interested not only in what this will tell us about the possible adaptive value of emotion, but also its likely link to and dependence on properties of a given body and environment (implementation or task[s]). More specifically, however, we ask: How will changes in the physical (such as sensory-perceptual and motor-behavioural) abilities of predator and prey agents — sensory capabilities, in this case the distance into environment information about stimuli can be obtained — interact to affect the balance and dynamics of their relationship, in terms of the advantage of one over the other in given encounters? This includes who "wins", but more importantly it considers the behavioural interactions and the adaptive value of a mechanism simulating neuromodulation, as a proposed substrate of emotion and biasor of action selection.

Having considered the literature with regards to action selection and emotion in the context of brain-body-environment interactions, as well as the brain-body-environment relationship, we think this question is interesting because it not only explicitly explores the importance of certain specific aspects of body in producing adaptive behaviour, but also considers their importance for the successful integration of emotion and emergence of specific, adaptive behaviours within a predator-prey situation. For it not only looks at what kind of role emotion might play with regards to brain-body-environment interactions, but also how the presence of another agent (prey or predator) might concurrently affect and direct this relationship and its interactions.

To put this more simply, in the rest of this thesis we will ask: What will happen to the dynamics of a predator-prey relationship in terms of physical/behavioural advantage and consequent adaptive value of a mechanism simulating neuromodulation as a biasor of action selection when sensory capabilities (in this case, the distance into environment information can be obtained) are varied?

# Chapter 3

# A Model of Emotion-Based Action Selection

*"Remember that all models are wrong; the practical question is how wrong do they have to be to not be useful."*
**— George E. P. Box**

Adopting a bottom-up, animat approach to investigate intelligence and adaptive behaviour, in this thesis we ask: How will changes in the perceptual abilities of predator and prey agents interact to affect the dynamics of their relationship, in terms of the adaptive value and emergence of adaptive behaviour from an emotion-based architecture for action selection? With regards to our wider research interests, but especially those encapsulated within our research question, the last chapter provided an initial overview of some of the literature relevant to and motivating our research — developing a framework for, as well as positioning our ideas and subsequent work within the wider literature. As we have illustrated, with regards to their roots in the literature, the nature of the goals, encompassing the research question and aims, of this thesis are particularly interdisciplinary. This poses several challenges, one of the main ones being to choose among the many different methodologies available.

The underlying philosophy encapsulated within our model is that by representing and explicitly modelling agent brains, bodies and environments as components, we have the means for producing, using a bottom-up approach, aspects of intelligence including emotion-like phenomena. Being strongly convinced that a bottom-up approach would provide us with a new look at an old problem we first had to consider how to model and explore our main ideas further. In this chapter, we proceed to both narrow and extend our focus on the literature to turn to and describe the model we have consequently adopted in greater detail. In selecting and designing elements of our model of brain-body-environment interactions in a predator-prey relationship, we aimed to construct a useful model for testing our ideas about the emergent properties of a mechanism simulating interoceptive modulation. Adopting an animat approach, in our model we make use of the Hazardous Three Resource Problem (H3RP). This is the formalisation of a specific scenario which appealed to us because we consider it enables the combined study of action selection, brain-body-environment interactions and the predator-prey relationship. Here, we start to describe our model by providing an initial description of the basic elements of the H3RP, looking at the research which generated it.

The specific emotion-based action selection mechanism chosen for use in our investigation throughout the rest of this thesis is also outlined, which we further link to the existing literature by comparing it with some examples of others using several criteria. The chapter concludes with the examination of the different measures researchers have used to evaluate agent performance. The calculations we use to measure "adaptive value" can be differentiated between in terms of the more quantitative measures of viability and more qualitative temporal measures relating to activity cycles. These are the dependent variables measured in our experiments (an outline of which will be provided in the next chapter).

## 3.1 The Hazardous Three Resource Problem

### 3.1.1 The Two Resource Problem (TRP)

The H3RP is perhaps best described as an extension of the Two-Resource Problem (TRP). Used by researchers to study the problem of action selection, the TRP is an embodiment of the simplest model of action selection there is. That is, the situation where an agent, given a choice between two resources, must choose at any given moment which one to search for or consume.[1] However, we focus on one formalisation of this problem in particular, looking to the work of Avila-García [29]. This is because, in keeping with our own interest in the potential role of the body, both internal and

---

[1] With regards to the terms used to describe both our ideas and research within this thesis, we are aware some can have certain biological, or other, connotations we may not intend. However, despite the potential controversy in using them to describe elements which relate to artificial agents, we argue it is less confusing to simply adopt these terms, giving working definitions if necessary, than to invent a new terminology. For, from our perspective, introducing new terms will not only be a potential source of confusion (why invent a new word when the reader grasps what is meant more quickly using an existing one, or will likely only substitute it in their mind for something more biologically-related anyway?) but will likely also disassociate the element it refers to from the possible biological counterparts they have been inspired by. Moreover, it affects the interest and appeal of our work when other researchers have to learn a new vocabulary to understand it. This would not be in keeping with the interdisciplinary nature of this research (in terms of encouraging a two-way flow of ideas). Thus, where we freely use terms such as consume, physiology and nest it not always in their fullest, biological sense, but in a more restricted, often more abstracted, artificial one.

external, for artificial intelligence, in his formalisation of the TRP an internal physiology is introduced. Two physiological variables, internal to the agent, must be kept within a certain range, or else the agent will die (which would not be particularly adaptive!). These two variables are set to naturally decrease with time, following a kind of metabolic activity. However, they can be increased if the agent can reach and consume the corresponding resource in its environment. An action selection mechanism must therefore be devised to keep the agent alive long enough to reach one resource, without dying from lack of consumption of the other.

### 3.1.2 The Hazardous Three Resource Problem (H3RP) — A Dynamic Extension of the TRP

In terms of environment, the original TRP was a static one in as much as the agent's surroundings were unchanging. However, the TRP was extended to create the H3RP. Compared to the TRP, the H3RP incorporates additional elements to make it both a dynamic environment as well as a platform for studying the predator-prey relationship. For here another agent (the predator) is introduced. With this problem, the original agent of the TRP (now a prey agent) must still decide at any given moment which, of two resources, to search for or consume. However, it now also has to deal with the additional environmental dynamic caused by the danger introduced by the predator, which can attack and therefore "kill" its prey by decreasing levels of a third physiological variable. Not only does the now-prey agent gain this additional physiological variable to keep within a viable range, but it also gains another resource type in the form of a nest. This can be thought of as a safe place where it is protected from its predator in as much as it cannot be attacked, but where the levels of its other two physiological variables will still continue to decrease.

The prey's task now becomes more difficult, as it does not just have to balance its time between the two resources of the TRP. Indeed, any action selection mechanism designed for the prey must now allow this agent to balance its time between *three* resources. This means the prey must effectively "decide" how or somehow incorporate a successful strategy to deal with the presence of its predator. Successful action selection mechanisms might use various strategies to do this. For instance, perhaps biasing the prey to spend more time "hiding" in the nest, so that there is less chance of an attack — coming out only to "forage". Or, to give another example, spending less time in any one location — this might keep the prey from stopping too long in one place, which might allow the predator to catch up. Alternatively, this might involve the prey risking attack. Though, if attacked, the prey then also still has to choose whether to spend time searching for its nest and healing any damage or satisfying its other goals first, risking further attacks.

### 3.1.3 Earlier Research

Overall, the H3RP was one of three problems formalised, of which individual robotic implementations of each one were developed and studied. Using LEGO RCX robots, the first to be created was an implementation of the TRP. Following the results of experiments in simulation, this was used to compare two different action selection architectures. The experimental data collected was analysed using the more novel concept of *activity cycles*. That is, analysing and seeing action selection as more of a cycle of activities — which in the TRP must keep the internal environment of the agent stable — than a series of individual actions or decisions made.

The second implementation, building on the first, was an implementation of the Competitive Two Resource Problem (CTRP). Like the H3RP, this problem can be thought of as an extension of the TRP. For, in this scenario, two agents have to effectively perform the TRP in the same environment. Biologically-speaking, this can also be related to the case of conspecifics competing for the same resources.

Avila-García found that the architectures of the TRP were not adaptive enough to deal with the additional dynamics introduced to the CTRP or the H3RP. However, with both the CTRP and H3RP he found the performance of motivation-based architectures could be improved with the addition of a hormone-like mechanism. In the CTRP this was achieved by a process akin to exteroceptor modulation (modulating internal and external sensor readings) and in the H3RP by a process akin to interoceptor modulation (modulating perception of an internal variable — in this case, the variable used to represent the damage sustained by the prey agent).

## 3.2 An Emotion-Based Architecture

Taking inspiration from the work of neuroscientists such as Damasio [86] and Fellous [19; 20] — the former having used studies of brain-damaged patients to suggest emotions have a vital role to play in rational human decision-making — it has been suggested that mechanisms incorporating or otherwise simulating functions equivalent or similar to those of emotion(s) may be useful in the successful design and improvement of existing action selection mechanisms. In computer science, particularly for AI and AL researchers, such evidence has been further used to support the idea that some kind of artificial emotion(s), playing an equivalent function to emotions in human and/or animal decision-making, might be required in order to produce more adaptive and autonomous behaviour from existing robot action selection architectures. For example, by increasing behavioural flexibility or *plasticity*.

Exploring this idea, researchers interested in this area have already attempted, with varying degrees of success, to create more flexible action selection mechanisms by designing what we term “emotion-based architectures”. However, the approaches taken to incorporate the idea into a model of action selection and emotion have been equally varied. From the more top-down approach taken by researchers such as Breazeal [92], to the more bottom-up approaches of researchers such as Cañamero [67]. While all these approaches have produced interesting models of emotion, exploring its possible role in relation to action selection, we think the latter approach may be especially useful to researchers who wish to investigate plausible mechanisms underlying emotional modulation of action selection and their adaptive value in autonomous robots. That is, in as much as it avoids taking the more "black box" approach often advocated by top-down models.

As previously noted, in our research we focus on studying and developing one action selection mechanism in particular[1]. In keeping with our research question, and positioning our chosen mechanism in the emotion-based subset of action selection mechanisms, ours implements the same hormone-like mechanism first developed, and previously used, by Avila-García in the context of the H3RP. An added advantage of using this mechanism is that doing so allows us to compare our results directly with those obtained in the original H3RP implementation. However, we chose this architecture for many other reasons as well. Primarily, because neuromodulation has previously been noted by neuroscientists as a possible "neural substrate" of emotion — and the emotion-like states introduced to an agent using this hormone-like mechanism can be said to simulate the effects of neuromodulation, albeit at a more abstract level than that of the neuron.

The hormone-like mechanism that is adopted here is particularly appealing as the idea (or at least one of the major advantages) is it is one that can be "added" on to any existing control architecture quite easily (in this case, a motivation-based one). It can thereby be used to modulate a given action selection mechanism or controller, changing a more reactive architecture into one which can respond dynamically over time (in this case a simple motivation-based architecture has been selected as the base). It acts as a second order controller to the underlying architecture and can be connected to stimuli/components both internal and external to the agent.

In terms of artificial agents, what is particularly attractive about this hormone-like mechanism is it is used as secondary controller, and modulator of perception, to an existing architecture for action selection. That is to say, it builds on top of a previously "emotion-free" action selection mechanism to make it emotion-based. In our case, the underlying action selection mechanism is a motivation-based architecture. However, the emotion-like component is not specific to this type of architecture. This means it is one which could also be attached to or incorporated in practically any existing action selection mechanism. Another point in its favour is that it does not make as many a priori assumptions about the functions of emotion as many other emotion-based action

---

[1] Though we do study the effect on its performance of changing various parameters, which could also be thought of as creating and comparing different architectures

selection mechanisms, allowing us to take a more bottom-up approach.

In our model we focus on implementing motivation-based architectures as our basic animat "brains". Motivation-based architectures can also be thought of as a subset of action selection mechanisms. However, their defining feature is that they are inspired by ideas about *motivation*, represented in some way as to be used in an action selection mechanism. With regards to the concept of motivation, the motivation-based architecture we have used incorporates notions of both intrinsic and extrinsic motivation. Intrinsic motivation comes from within the individual agent, whereas extrinsic motivation comes from outside. In our H3RP this is reflected in our agents' calculations of motivational intensity, which uses the level of internal drives generated by the current level of physiological variables (intrinsic motivation) and the presence of stimuli, including resources, detected by agent sensors (extrinsic motivation).

Considering the hormone-like mechanism we use to create our emotion-based architecture, we are especially interested in exploring how it might be implemented outside of computer simulations, embodied in the real world, in scenarios which place the architectures and evaluate their performance within complete agents (taking an animat approach) and in dynamic environments.

In the case of the H3RP, this hormone-like mechanism is connected to the perceived presence of the predator. If detected, the prey's hormone-like mechanism will signal a "release" of hormone that will act as a biasor of action selection — in this case encouraging the prey to rest or search for its nest. In this way, it can be likened to the primary emotion or some kind of mechanism of "fear". Acting in this way, this mechanism has been tentatively linked to emotion-like modulation of behaviour by researchers taking a bottom-up approach to the study of affect and action selection (following ideas, such as Fellous', about neuromodulation as the "neural substrate" of emotion for humans and other organisms). Previous experimental work has also found it to be useful for adaptive action selection in a predator-prey context: primarily, as a modulator of interoception (internal perception of the levels of physiological variables).

With this in mind, we have developed our own basic implementation of the H3RP (outlined in greater detail in the next chapter). Initial experiments looked at the performance, as measured by measures of viability also outlined by Avila-García, of a prey robot with and without a hormone-like mechanism, placed in different variations of the H3RP. The idea was to try to vary the relationship between "hormone-release" in terms of perceptual capabilities and outside stimuli, including features of the environment such as the resources of the prey's environment and ability of its predator.

By using the H3RP as a framework we effectively study a society of sorts, composed of predator and prey. Specifically, focusing on developing agent brain-body combinations that create agents able to cope with and behave adaptively in ever more dynamic environments. In the case of robots, this refers to dynamics found in the real world. That is, our own — an aspect that makes it all the more appealing to us especially.

Figures 3.1 and 3.4 present Avila-García's architectures. One of the features of these architectures that particularly appeals to our interests is they do not explicitly label any one component as "emotion" — something we also advocate. Instead, we both prefer a more bottom-up approach: trying to model one of the suggested neural "substrates of emotion" — namely, neuromodulation. We do this in order to examine the emergent properties of a system, which may consequently resemble the "emotion-like" behaviours of real-life adaptive agents. Thus, we have both attempted to simulate the effects of neuromodulation for the benefit (adaptively) of action selection mechanisms. In addition, at a level of abstraction which has resulted in the development of hormone-like mechanisms ("hormone-release" occurring in the presence of relevant external stimuli) which affect action selection over time. In particular, Avila-García examined different ways in which such a mechanism can act as a biasor of action selection, modulator of perception (both interoception and exteroception) and "second-order controller" for existing architectures.

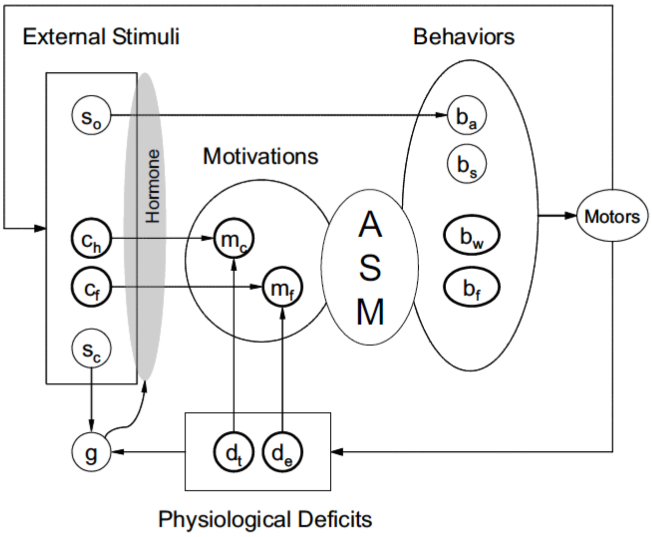

Figure 3.1: Illustration and overview of Avila-García's modulated action selection architecture, used in the CTRP: manipulating perception, and action selection, via hormone-like modulation of *exteroception* using a motivation-based architecture. Emotion-like states are modelled by the addition of a gland (g) releasing a "hormone", the level of which is an additive function of internally-perceived physiological deficits and the presence/perception of a competitor. This affects perception of external stimuli, increasing calculations of motivational intensity. Hormonal concentration decays over time [29].

## 3.3 A Comparison

To give an idea of where we place our emotion-based architecture and resulting work in relation to that of existing emotion-based architectures, it may be useful to conduct a quick comparison using some simple criteria. Here we do so in order to effectively, albeit briefly, contrast our work with architectures created as a result of the work of five different researchers: Breazeal, Arkin, Avila-García, Delgado Mata and Mendao. We chose each of these researchers and their architectures for different reasons: Breazeal [92] provides us with a "classic" architecture for comparison, Arkin [93] with a relatively recent addition for us to compare (TAME being the "state of the art" in the history of his work) and Avila-García's work [29], as we have already mentioned, is in many ways closest to our own, which makes it important for us to distinguish the ways in which our approach and architectures differ. So as to also better illustrate the similarities and differences between our work and that of researchers generally adopting a bottom-up approach to emotion, Delgado Mata [82] and Mendao [84] were then chosen as two further advocates/representatives of this approach.

So as to get more of an overview of these differences, we look at these researchers' work in reasonably broad terms. We do so here firstly in the context of how each of these researchers treat/incorporate ideas about emotion in their architectures. Secondly, we compare their primary motivations, including the problem/domain of interest they are interested in. Finally, we look at what they consider adaptive action selection to be (in other words, their measures of adaptive value).

### 3.3.1 Function and Integration of Emotion

Illustrations of the types of architecture produced by Breazeal, Arkin and Avila-García, along with our own, are provided in Figures 3.2-3.5. To highlight the differences, we can first look at how each one sees "emotion". That is, their ideas as to the function and integration of emotion for action selection mechanisms. As can be seen from Figure 3.2, Breazeal's architecture explicitly introduces emotions as a subset of motivations. Ideas about the function of emotion as being communicative are incorporated through the modelling of emotional expressions (the "actions" selected by her implemented robot Kismet) and internal "emotions" are used to activate a robot's physical "emotional expression" at any given time.

In contrast, Arkin (Figure 3.3) has more recently been contributing towards the development of the TAME architecture. This introduces and incorporates emotions in what might be considered a more "sophisticated" model, where emotion is treated as one of a number of affective phenomena to be explicitly modelled (traits, attitudes, moods and emotions). Similarly to Kismet, the robots (AIBO and Nao) in which TAME has been implemented have used emotion in a communicative context. This differs from some of his earlier architectures, looking "up the food chain", which were

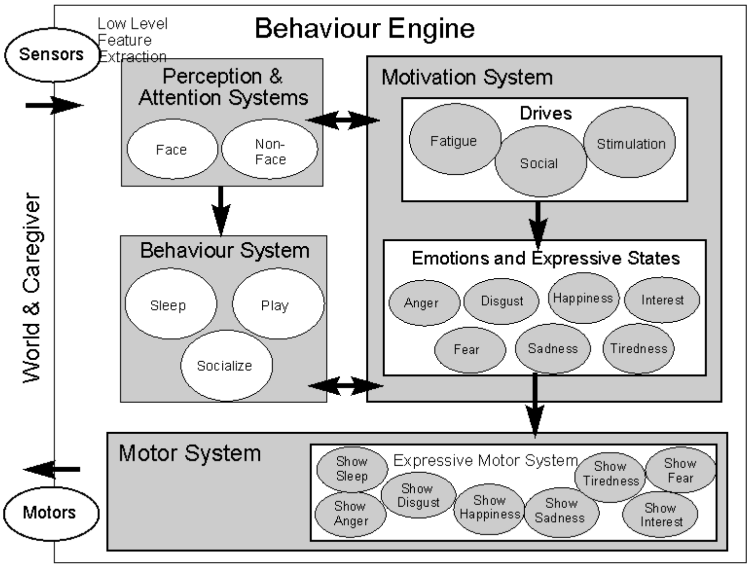

Figure 3.2: Illustration and overview of Breazeal's architecture for Kismet: Incorporating ideas about different types of emotions and connecting them to different motor responses (emotional expressions) [92]

generally based on the ideas of his earliest architecture (AuRA) but also were used to explore the functions of emotion for an individual, autonomous agent.

Looking now towards those researchers specifically adopting a bottom-up approach to emotion (as we are here) in his work Delgado Mata [82] considers one of the functions of emotion as being to act as a mediator between individual and collective (or social) behaviour. To investigate this further, his own architecture models exteroceptors used by real animals to detect the presence of chemicals in the external environment as a virtual nose. Focusing on the communication of emotion between agents, the idea here is that emotion will reduce the complexity of behaviour. In tune with our own work, Delgado Mata demonstrates both a focus on the emotion of fear in particular and a consideration of its role between agents. However, the motivation for our work is different in as much we are not specifically focusing on emotion as a social regulator per se.

Finally, with more relevance for our own work (Figure 3.5), we can refer to Figure 3.4. It is this work that we most closely align ourselves to in the respect of function and integration of emotion. This is because, in his architecture, Avila-García does not actually explicitly label any one component as "emotion" — something we also advocate. Instead, we both aim to adopt a more "bottom-up" approach, by instead trying to model, taking inspiration from neuroscientists such as Fellous, one of the suggested neural "substrates of emotion" — that of neuromodulation [19] — in order

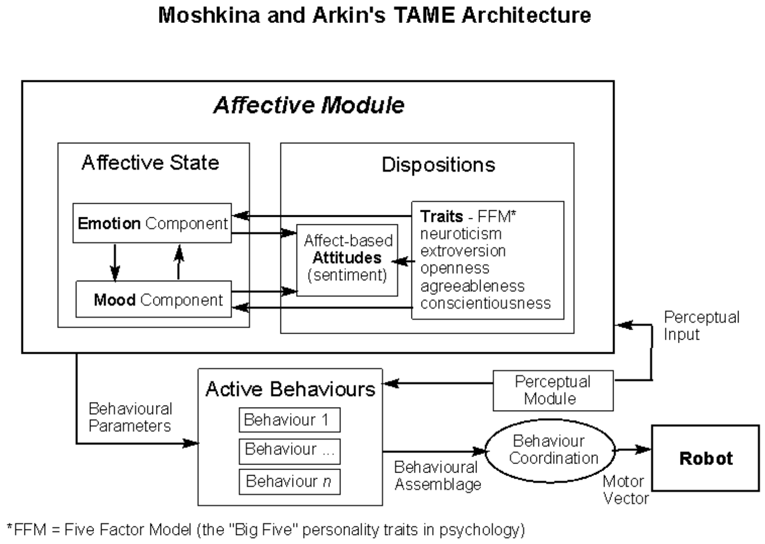

Figure 3.3: Illustration and overview of Moshkina and Arkin's TAME Architecture: Incorporating ideas about and explicitly modelling personality and emotion using concepts connecting Traits, Attitudes, Moods and Emotions (each of these varying in their temporal effects and influence on each other) [93]

to examine the emergent properties of a system, which may consequently resemble the "emotion-like" behaviours of real-life adaptive agents. Likewise, Mendao [84] is similarly cautious about attributing a particular emotion to the emergent behaviour demonstrated by his architecture. Having implemented an artificial system consisting of neural networks that are sensitive to the available concentration of specific hormones (released under certain conditions) his work bears a closer resemblance to our own in our combined efforts to see if hormone-like mechanisms give us the potential to achieve complex behaviours through the interaction of simple structures. Whilst we do not use neural networks in our own work, interestingly, Mendao demonstrates the potential of the types of hormone-like mechanisms we are interested in for architectures other than our own.

In both Avila-García's and our own architectures, we therefore attempt to abstract the effects of neuromodulation for the benefit (adaptively) of an action selection mechanism — which has resulted in the focus on development of a hormone-like mechanism ("hormone-release" occurring released in the presence of relevant external stimuli) affecting action selection over time. In particular, in his architectures Avila-García examined different ways in which such a mechanism can act as a biasor of action selection, modulator of perception (both interoception and exteroception) and "second order controller" for existing architectures (in this case a motivation-based one).

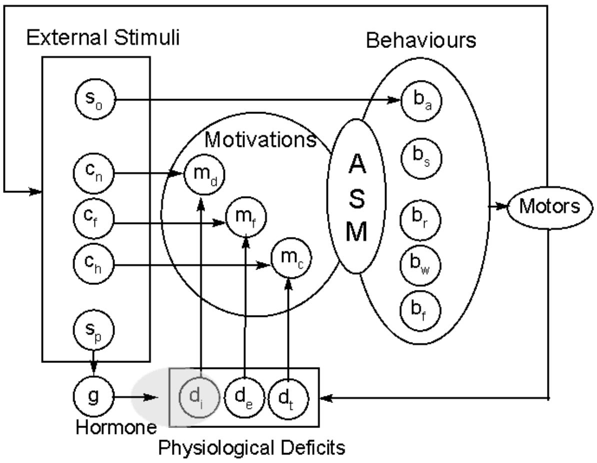

Figure 3.4: Illustration and overview of Avila-García's modulated action selection architecture, used in the H3RP: manipulating perception, and action selection, via hormone-like modulation of *interoception* using a motivation-based architecture. Emotion-like states are modelled by the addition of a gland (g) releasing a "hormone" in the presence of a specific stimulus (in the case of his predator-prey scenario, the H3RP, the predator) which affects perception of internal physiological deficits, increasing calculations of motivational intensity. Hormonal concentration decays over time [29].

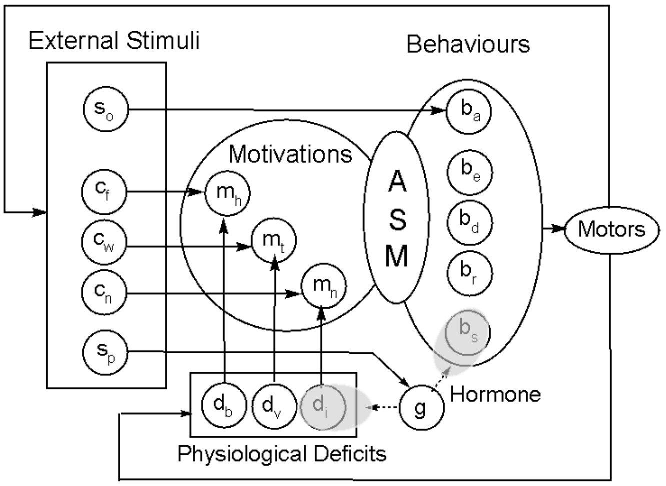

Figure 3.5: Overview of our developed architecture (“brain”) for a prey agent: internal “body” is represented through physiological variables, deficits of which act as drives which, combined with the presence/absence of external stimuli, are used to calculate motivational and behavioural intensity. For example, calculations of motivational intensity for a motivation representing hunger will take into account both physiological deficits such as blood sugar and the presence/absence of food in the environment. In our experiments we vary external “body” using different physical sensors. Emotion-like states are modelled by the addition of a gland (g); releasing a “hormone” in the presence of a specific stimulus (in this case the predator) which affects both perception of internal physiological deficits, increasing calculations of motivational intensity, and the behaviour selected in terms of physical response (speed or tempo of behaviour is increased if hormone is present)

However, one way in which our currently developed architecture differs from others is that we try to integrate this kind of mechanism more pervasively or intricately with the rest of our architecture. As later chapters will show, we have linked our hormone-like mechanism not only to calculations of motivational intensity, but also the intensity of behavioural response. To give an example, in latter experiments, this has translated into an implementation of a prey agent that, when its "hormone level" increases, so too does its physical speed. Thus, we use this "substrate" not only to modulate perception, but to influence behaviour more dynamically and physically, in terms of factors such as time and speed.

We think this has the advantage of effectively making "shortcuts" or more direct links between a perceived external stimulus and physical response/readiness of action, which may especially help in the problem of allocation of limited "energy" resources. Moreover, we go further to consider the interactions between two agents (and their architectures) rather than looking at one individually.

### 3.3.2 Problem or Domain of Interest

Next, we can turn to and compare the particular areas or "problems" that these architectures, or to be more specific their researcher's implementations, have been used or designed to study and solve. We attempt to do so here with regards to each researcher's particular contributions towards the study of action selection. These are reflected in the implementations each researcher has developed, and the particular type of context (environment/scenario/task) they have looked at the role of emotion or emotion-like states in. By doing this, we can extract some of those features of an action selection task that each focuses on.

Whilst each architecture can itself be considered as a contribution to the action selection literature, and all have been implemented in robots which is especially appealing for reasons we have previously mentioned, they have each been implemented for quite different purposes and in quite different environments: Kismet to model social interactions between infant and caregiver (human-robot interactions); Arkin's TAME to model affect more sophisticatedly for human-robot interaction; Delgado Mata's to integrate emotion with flocking and grazing behaviour in communicating emotion between virtual animals; Mendao's to see whether taking inspiration from biological hormones might increase the flexibility of traditional artificial neural networks; Avila-García's to test the properties of architectures across different types of environment/scenarios (only one of which includes a predator-prey type scenario); and ours to study action selection within a very particular context and relationship (predator-prey) in order to examine brain-body-environment interactions.

First, in more general terms, we can see that the primary implementations of both Breazeal's and Arkin's architectures have been in the area and interests of human-robot interaction: the robot head Kismet being a result of Breazeal's; and TAME having been implemented in both Sony's AIBO dog and the humanoid Nao. While human-robot interaction is of course an extremely relevant and interesting area for the study of the role of emotion (particularly with regards to communicative functions and interactions) what sets such architectures apart may actually therefore be that they are designed to say as much, if not more, about our own emotions and interpretation of another agent's behaviour than its generally-applicable function(s) for a robot/agent as an autonomous decision-maker. That is, as an agent interested primarily in its own survival.

In contrast, we regard this as bringing a dimension to the work that we currently prefer to leave out of our own, in favour of focusing our study more on our agents — and, one of the advantages of a synthetic approach, the interactions of two agents we already know the exact internal workings of — and less on their impact on our own (human) behaviours and perceptions of them as agents. Having said this, however, of course we do not rule out the fact that we may inadvertently introduce our own bias as researchers if we are not careful in how we study them.

To a lesser extent, the same can also be said of the work of Delgado Mata — one further objective of this work being to examine how far his animals could increase the feeling of immersion experienced by a human user in such environments. In terms of the primary similarities and differences to his other research interests, whilst we might point out that we can both be said to be generally investigating the function of emotion between multiple agents, one of the most obvious differences between our own work and that of Delgado Mata in terms of the problems addressed/domain of interest is that we also explicitly model the action selection of a predator agent. Furthermore, although we would not consider the two to be mutually exclusive, in our work we have taken the decision to test our ideas directly outside of a computer simulation. That is, we have taken the decision to implement our architecture in robots in the "real world". Though focusing only on the performance of one agent, in comparison to this, Mendao's work is perhaps more closely aligned to our own in the respect of both the domain of interest and problem being addressed, having implemented his architecture in a mobile robot (a Pioneer Dx2) equipped with a camera.

Avila-García similarly goes a different way to Breazeal and Arkin: implementing his architectures across different scenarios, also using LEGO robots (Taurus and Sador being examples of these). However, he focuses instead on developing ways to quantitatively and qualitatively measure these implementations as individual adaptive systems, so as to identify their specific properties in different contexts. That is, regarding other agents solely with regards to how they may add to the environmental dynamics, and possibly environmental complexity, rather than as an agent in a partnership or some kind of artificial ecology, which can affect and be affected by other agents.

So, by not focusing on one particular problem, Avila-García was able to look at the properties of architectures, in particular arbitration mechanisms, across different scenarios. He developed several types of scenario for the study of action selection, including a robotic two-resource problem, competitive two-resource problem and hazardous three-resource problem (H3RP). Yet, even in his predator-prey type scenario (the H3RP) action selection did not involve situations of such high risk as might be expected of such a relationship. This was due to the development of a more "parasitic" type of predator-prey relationship (allowing his agents some leeway in choosing to change activity).

This does not, however, mean that we do not want to, think, or aim to contribute towards developing ideas that may also be of use to these other domains of interest. More, we think by focusing on our particular scenario now — that of predator-prey — we will be able to bring something particularly special or unique to the problems of these other architectures later. Currently, for instance, all of these other architectures, when you consider their researcher's implementations, do not seem capable of producing adaptive behaviour in situations where both the two-way relationship between two agents is accounted for (in other words, considering both agents) and the right decision/action selection is vital for agent survival. That is to say, high risk.

What is primarily different about our own motivation then, regards to the kinds of decision and environmental demands we want our architecture to deal with, including those situations where there may not be enough time or flexibility to allow for mistakes or trial-and-error learning; instead requiring split-second judgements. More to the point, we want to consider this in the predator-prey scenario for a much more in-depth look at this kind of relationship, where a predator is not just an environmental dynamic. For example, if someone were to wave a baseball bat in a threatening manner at a robot, what we would like to see this robot's architecture capable of doing is to use "fear" to better make those split-second decisions that will direct action selection towards the agent's own survival. This could involve some means of "fleeing" the scene, but might even involve something like a suppression of emotional expression so as to calm the human threatening it down.

Another difference can be seen in the type of intelligence or adaptive behaviour studied. For example, Breazeal and Arkin can be said to study action selection and emotion more towards our ideas about human-level intelligence and emotions (though Arkin has in fact previously studied those relating to a lower, more insect-level intelligence). However, again in common with Delgado Mata, Mendao and Avila-García, we attempt to go back to basics more; considering these concepts more in terms of animal-like mechanisms of adaptive behaviour and intelligence — another reason for studying the predator-prey relationship. While Arkin has previously studied architectures aiming towards insect-like intelligence, incorporating and developing ideas about motivation and emotion, in "moving up the food chain" [94] it does appear he left a somewhat expansive gap between the level of insect and that of animals. This is where our work comes in. That is, somewhere between the reactive architecture given to an insect and the more deliberative architectures he chooses for those interacting with humans.

### 3.3.3 Measures (of Adaptive Value)

Finally, we can also compare these architectures in terms of the level of analysis and criteria each researcher expects will be used to measure the adaptive value of their architectures in a given implementation. Without going into unnecessary detail, in this respect, perhaps due to their interest in human-robot interaction, both Breazeal and Arkin can be said to have focused on the use of both internally and externally-derived measures. That is to say they measure, for different purposes, both external effects of the robot's action selection on human response and the internal parameters of the system or architecture over time. When involving observations, this is often a lengthy process with regards to analysis, but has the benefit of allowing us to directly study interactions between humans and robots. Although not as focused on observational analysis, Mendao nevertheless also used data from both internal and external variants for statistical analysis, with Delgado Mata focusing on examining the trajectories of his agents across time steps as well as defining a measures of complexity.

Conversely, Avila-García's architectures have been studied placing focus on the use of more internally-derived and summarative measures: to develop measures of analysis that consider the viability of his agents over an individual life span (as the correct level of analysis to study adaptive value). Avila-García also considered action selection in terms of activity cycles rather than separate decisions.

Similarly, we would like to consider how analysis of behaviour over time might bring us more insights into our architecture's behaviour in different predator-prey scenarios. In our work though, perhaps more in common with Breazeal, Mendao and Arkin, we try to combine the use of both externally and internally-derived measures. However, we also attempt to go further, for a more comparative look. One of our primary concerns is to ask at what level of study will we find out most or understand our systems best, especially with regards to what one might consider adaptive value to be (and in terms of brain-body-environment interactions and their emergent properties). In this way we seek to bridge the gap between these architectures, in respect of the level their researchers have proposed we analyse them at, whilst also trying to develop our own.

One source of inspiration for us in this endeavour again comes from another discipline: ethology. Though dynamic systems theory has developed tools to study the interactions of dynamic systems, we use the analogy of animal-like behaviour to suggest that the ethologists have already developed many tools to be used in the analysis of our animat agents. In particular, many of these methods allow us to combine both considerations of internal and external data (as derived or obtained/collected from experiments).

### 3.3.4 Contributions

Having considered our work in the context of these existing emotion-based architectures, the specific contributions we hope to make with our model, especially towards the literature on action selection and emotion (or affect), can be summarised as follows:

#### 3.3.4.1 For Affective Action Selection:

- Development of our emotion-based architectures and implementation of the H3RP: extending and investigating the scenario further by varying properties relating to the architectures, emotion-like states and embodiment of each agent. In initial experiments, for example, we divided perception into proximal and distal types (combinations of which make further sub-problems or versions of the H3RP). This enables and hopefully justifies direct comparison, especially in terms of the interactions of different physical properties of predator and prey, with previous findings using the same framework (such as the work of Avila-García). At the same time, this also introduces a new dimension for study (combination with different types of embodiment: in this case range of perception or "sensory ability"). Such a comparison will, for example, enable us to identify aspects of the original scenario that may have been crucial for the success of the proposed emotion-like mechanism and overall action selection.

- A more systematic study of the predator-prey type relationship than has been conducted yet in the action selection literature with regards to affect. For example, looking to see the minimal conditions under which our chosen mechanism (or emotion in general) might be adaptive, especially with regards to the capabilities of our agents' "brains" and "bodies", as well as features of the environment, varying both abilities of predator and prey (for, while others have looked at the role of emotion in the predator-prey scenario, they do not necessarily know or have not necessarily taken into consideration how their mechanisms or emotion-based architectures might work, or be developed to work, in increasingly more dynamic environments, or with different types of embodiment such as sensory-range).

- An analysis of costs and benefits of both emotions and decisions in the predator-prey relationship: looking at neuromodulatory effects as the basis for emotion, when used in different ways for agents (for example, aggression for predator and fear for prey). Moreover, looking at action selection mechanisms more in terms of trade-offs, thus using these mechanisms as assessors of risk or opportunity cost — quick or rough-and-ready filters for behaviour and/or representations of the importance and limited nature of time — so as to allow for timely, temporally-adaptive response. That is to say, looking at action selection in terms of a trade-off between the time taken to decide and time taken for environmental circumstances to change adversely.

#### 3.3.4.2 For Analysis of Adaptive Systems:

- A comparison and evaluation of measures of adaptive value (both quantitative and qualitative) that might be used: from internal measures of viability, to internal Markov Models of an individual agent, to Markov Models constructed from external observational data (adopting the idea of activity cycles, thereby looking to analyse temporal behaviour of agents rather than simple life span etc).
- An analysis of the action selection problem in terms of the brain-body-environment relationship. Taking a broader look at action selection, so as to ask whether we should actually be looking at the architecture alone in isolation, or whether we find out more by considering elements together (for example, considering both architecture and body, predator and prey — together, rather than individually). Moreover, looking at how (more realistic) two-way interactions may affect performance of architectures and where emotion might fit in the relationship.

#### 3.3.4.3 For System Design:

- Demonstration of how we might manipulate or adjust parameters so as to better "fine-tune" our mechanism and increase its value for adaptive action selection in this context (of predator-prey H3RP). In particular, looking at how we might benefit from further distributing control and neuromodulatory influence across both agent architecture and agent body (as generators of brain-body-environment interactions).

While each of these contributions alone may provide useful insights into various aspects of the problem of action selection, we suggest that together these contributions will enable us to make an altogether much more comprehensive, greater, even synergistic, contribution to the literature. For example, attempting not only to link important concepts such as action selection and emotion to the predator-prey relationship and brain-body-environment interactions; but, in turn, highlighting their more general contributions to the more intelligent design or creation of artificial life.

### 3.3.5 Our Research

Relating these elements back to our research question then, we are asking "How will changes in the perceptual abilities of predator and prey agents interact to affect the dynamics of their relationship, especially in terms of the adaptive value of an emotion-based architecture for action selection?". To summarise our methodology, to answer this question we have adopted an animat approach. This means we have started by developing agents with less complex abilities — with the aim of incrementally increasing both the complexity and our understanding of our agents. To study a predator-prey relationship in terms of action selection and emotion we have designed an implementation of the H3RP, using the hormone-like mechanism it was first designed to allow study of. To extend this research further, we have also selected and developed a number of measures to calculate "adaptive value" for our agents. With regards to changes in the 'perceptual abilities' of these agents, which encapsulates our questions about the importance of agent brains, bodies and environments — as well as their interactions — in our experiments we manipulate elements connecting our agent "brains" with agent "bodies" and "environments". This includes use of different physical sensors, as well as changing the impact and connections of the hormone-like mechanism we study.

## 3.4 Summary

Within the current literature, Avila-García's work is in many ways closest to our own. This is especially true in terms of methodology in as much as we both study the H3RP and have developed robotic implementations of it to study this particular type of emotion-based architecture. Consequently, in many respects this thesis can be said to build on the findings of his previous work. However, we consider our own study of this problem to be much more in-depth, focusing more on the H3RP as it relates to the predator-prey relationship and less on using it as a platform to compare the performance of different architectures across different environments. One of the aims of our current experiments has been to allow us to see what factors might affect the success of the hormone-like mechanism, previously studied in this implementation.

If emotions were not inherently useful for survival, one argument is that evolution would have wiped them out already. However, just what continues to relate emotion to survival (if indeed it does) is not something researchers yet agree on. Indeed, many different functions of emotion have been proposed. Wanting to explore the idea of emotion (or one of the functions of emotion being to act) as a second-order controller, this thesis partly aims to ask whether hormone-like feedback or *modulation*, in the form of a hormone-like mechanism or module added to an existing control architecture — thereby creating an emotion-based action selection mechanism — will consistently be able to improve performance of that architecture, as an action selection mechanism, so as to enable it to behave more adaptively in more dynamic environments. At the same time, this research is just as concerned with looking at how the performance of the overall action selection mechanism itself can be found to interact with other features of both agent and environment. We do so namely by varying the perceptual abilities of both predator and prey — and in increasingly dynamic environments.

In this way, while taking much of its inspiration from the previous work of Avila-

García, we focus on their work mostly in terms of how it provides us with an interesting and useful platform for exploring dynamic interactions. In this case, those taking place between the perceptual capabilities of an agent and the performance of its action selection mechanism within a dynamic environment. We adopt the animat approach as one that specifically allows us to look at both the "whole" agent, as well as individual "parts" to see how they interact or contribute to overall adaptive behaviour. By developing agents using this approach, we are therefore advocating a bottom-up approach — by which we aim to start by designing and studying simple, but complete agents. The aim here is to then successively build on (our understanding of) these agents to hopefully build even more adaptive agents. To further define our approach, in terms of the four questions of ethology, rather than adopting a phylogenetic perspective — for example, by evolving agents — we focus on adopting a more ontogeny-related perspective. That is, looking over an individual animat's lifetime as opposed to lifetimes of multiple agents (whether generations or populations).

# Chapter 4

# Robotic Implementation of the H3RP

*“At bottom, robotics is about us. It is the discipline of emulating our lives, of wondering how we work.”*
**— Rod Grupen**

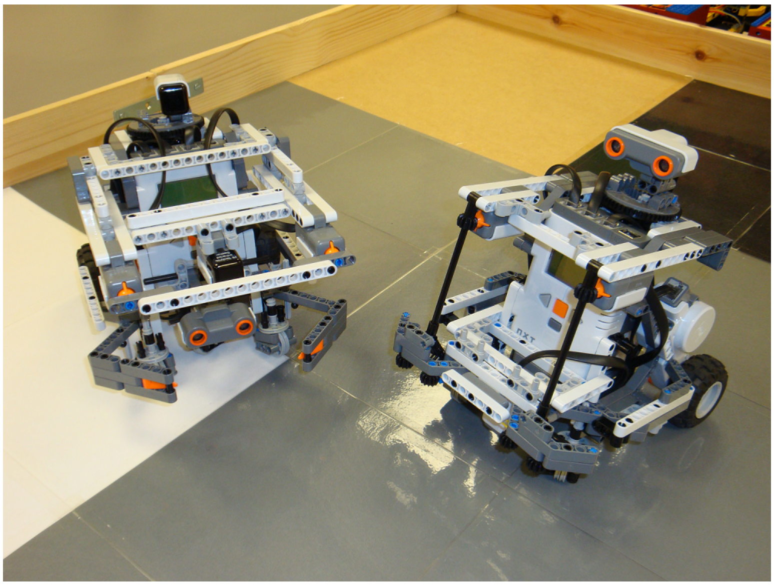

For both the study of artificial intelligence (AI) and artificial life (AL) robots continue to hold a special appeal as a platform for research. Part of their attraction is that, in many ways, building and studying robotic agents allows us to test ideas about our own intelligence directly in our own environment. This may give robotic models of intelligence a key advantage over those developed solely in computer simulations. Not only may they act as better litmus tests for theories about intelligence as they relate to our world, but they may also be able to cope with tasks and model aspects of intelligence that purely computational models of intelligence do not. Using LEGO NXT robots, we have developed a platform for investigating ideas about action selection, particularly with regards to the addition of emotion-like mechanisms in the context of brain-body-environment relationships (including agent-agent interactions).

Moving on from the more abstract or conceptual elements of our model, here we provide concrete details of both the software and hardware of the robots we used to create our physical implementation. In our experiments we have manipulated each of our basic animats (predator and prey) so as to create different variations of the H3RP (and predator-prey combinations) for us to study. By doing so, we wanted to see how our hormone-like mechanism especially would perform in such variations. Here we focus on providing a description of the basic prey and predator robots we have developed.

The last chapter introduced and described the main features of our model of action selection and emotion, introducing both the Hazardous Three Resource Problem (H3RP) and our particular type of emotion-based architecture. Using these to help us create a template for creating and studying different artificial predator-prey relationships, we have also selected several measures of adaptive value to create and measure the performance of agents that act within it, incorporating some ideas of our own so as to develop new measures for comparison. In this chapter we move on from the last to consider and describe the more technical details of our own implementation of the H3RP and measures for evaluating agent performance.

## 4.1 Our variation of the H3RP

Our implementation of the H3RP uses two LEGO robots (one prey, one predator) to create a platform for studying the adaptive value and emergent functions of a hormone-like mechanism. In particular, we have developed our implementation so as to enable us to consider what influence different aspects of embodiment have on this value. To develop our basic H3RP, however, we first created an implementation of the Two Resource Problem (as outlined in the previous chapter). Starting with the development of our prey we began by developing this agent so that it would be competent in solving the TRP, before trying to study it in the more dynamic environment of the H3RP.

The robots used in our experiments were two Mindstorm LEGO NXT Robots, one built as prey and the other as predator, each with 32 bit ARM processors running at 48MHz and direct access to 64KB of RAM and 256KB of flash memory (thus 320KB in total). The NXT model has an advantage over its predecessor, the RCX, in that it has 10 times as much memory, as well as a separate 8 bit AVR processor, running at 8 MHz, to control the servo motors and rotation sensors to guarantee the accuracy of the motor operations. Both robots' configurations were developed incrementally, building on a design called "R2MeToo", described by Brian Bagnell in his book "Maximum LEGO NXT: Building Robots with Java Brains" [95]. Our agents' architectures were programmed using Lejos, an open source replacement firmware for the NXT that aims to imitate the Java VM and API style.

In initial experiments we chose to vary both prey modulation (to see if our results would support the findings of research conducted using the original version of the H3RP, which found modulation helped increase the adaptive value of a motivation-based architecture) and prey perceptual distance (to test our own ideas as to how aspects of embodiment might interact with or otherwise affect any additive value of modulation). So as to better enable us to link our work with earlier work using the H3RP, an attempt was made to base the initial underlying architecture and external environment as much as possible on Avila-García's own robotic implementation. This was incorporated into design of the environment, as well as that of our predator and prey. Both robots were placed in an arena, creating an environment consisting of several different resources for the prey to choose between, thereby reconstructing the necessary environmental stimuli required for a version of the H3RP to be created. In our implementation we named the two physiological variables related to the TRP blood sugar and vascular volume.

Our agent (in this case robot) was placed within this environment, in which it could identify and distinguish between two different types of resource. In our implementation we labelled these resources food and water, with each internal variable being associated with one of the resources (blood sugar and vascular volume respectively). Over time, the levels of these variables naturally decrease, displaying a kind of "metabolic" effect. These can only be increased by the robot finding and "consuming" its associated resource. In this case, food to increase blood sugar levels and water to increase vascular volume. The robot's basic task at any point is to decide which resource to search for or consume, while the ultimate goal for the underlying control architecture is to enable the robot to make the right decisions so as to allow it to keep both internal variables above their set point — and thus enable the robot to "survive".

As outlined in the previous chapter, the H3RP extends the TRP by introducing an additional, third, internal variable. In our implementation we have labelled this "integrity". Though not present in the earliest versions of our implementation (that is, for the TRP) the additional resource we associated with this variable we have called our robot's "nest". However, unlike the other two variables, the integrity of our robot is not set to decrease predictably over time. Instead, this variable is affected by the presence of another robot, now labelled the predator, which is able to cause damage to the robot of the TRP (now the prey) and subsequently add greater complexity to the original scenario by developing a dynamic environment (changing over time depending on the actions of the predator).

In the H3RP, the prey can recover from damage to its level of integrity by escaping

| Animat | PV | Motivation ($m$) | Phys. Drive | PV Min-Max | Ext. Stim ($c$) |
|---|---|---|---|---|---|
| Prey | *Blood Sugar* | *Hunger* | $\downarrow d_{blood\ sugar}$ | 0-100 | *Food* |
| | *Vascular Volume* | *Thirst* | $\downarrow d_{vasc\ volume}$ | 0-100 | *Water* |
| | *Integrity* | *Need to Heal* | $\downarrow d_{integrity}$ | 0-100 | *Nest* |
| Predator | *Blood Sugar* | *Need to Hunt* | $\downarrow d_{blood\ sugar}$ | 0-100 | *Prey* |
| | *Energy* | *Need to Sleep* | $\downarrow d_{energy}$ | 0-100 | *n/a* |

Table 4.1: Summary of each animat's physiology: relating their physiological variables (PVs) to their associated motivations ($m$), drives, minimum and maximum levels and the external stimuli (incentive cues needed for consummatory behaviour) associated with them.

to its nest. However, the opportunity cost of this is, of course, the food or water it could have otherwise consumed. The robot's basic task at any point is still to decide which resource to search for or consume, just as the ultimate goal for the underlying control architecture is still to enable the robot to make the right decisions so as to allow it to keep all internal variables above their set point (thereby enabling the robot to survive). Yet successful control architectures will now have to take into account the dynamics introduced with the presence of the predator. Giving each a basic motivation-based architecture, using Winner-Take-All (WTA) arbitration and following the same design as those studied by Avila-García both predator and prey had to select appropriate actions from a behavioural repertoire. A summary of our robots' underlying physiological variables, motivations and behaviours are given in Tables 4.1 and 4.2.

| Animat | Behaviour | Type | Ext. Stim | Effects on Physiology |
|---|---|---|---|---|
| Prey | Avoid | *reflexive* | *obstacle* | −0.1 to blood sugar and vasc. vol |
| | Eat | *consummatory* | *food* | +1 to blood sugar, −0.1 rate to vasc. vol |
| | Drink | *consummatory* | *water* | +1 to vasc. vol, −0.1 to blood sugar |
| | Heal | *consummatory* | *nest* | +1 to integrity, −0.1 to blood sugar and vasc. vol |
| | Search | *appetitive* | none | −0.1 to blood sugar and vasc. vol |
| Predator | Avoid | *reflexive* | *obstacle* | −0.1 to blood sugar and energy |
| | Attack | *consummatory* | *prey* | +50 to blood sugar, −0.04 to energy |
| | Rest | *consummatory* | none | +1 to energy, −0.04 to blood sugar |
| | Hunt | *appetitive* | none | −0.04 to blood sugar and energy |

Table 4.2: Summary of our prey and predator's behaviours, the stimuli associated with them (incentive cues necessary for consummatory and reflexive behaviours to be executed) and their effects on each animat's physiology.

## 4.2 The Robots

Two main types of decisions had to be considered and made in the process of designing our predator and prey robots: those relating to internal mechanisms and those relating to physical components. In our implementation, this translated to the action selection mechanism chosen and used by us (as described in the previous chapter and Figure 4.1, we chose to implement a motivation-based control architecture/action selection mechanism) and the eventual hardware configuration of our agents.

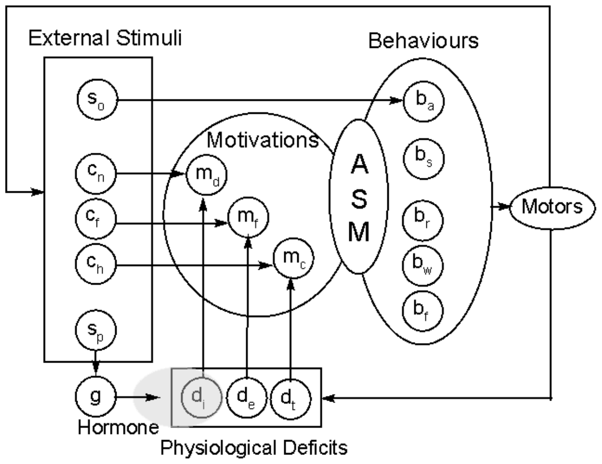

Figure 4.1: Overview of the basic architecture ("brain") of our initial agents: internal "body" is represented through physiological variables, deficits of which act as drives that, combined with the presence/absence of external stimuli, are used to calculate motivational and behavioural intensity. For the prey, a "hormone-like" mechanism is introduced by the addition of a "gland" (g) which releases a "hormone" in the presence of a specific stimulus (in this case the predator) which affects both perception of internal physiological deficits and can create a feedback mechanism between environment and brain via the body that can affect the behaviour selected.

### 4.2.1 Prey

Figure 4.1 again provides an illustration of the basic architecture used for the prey robot, while Tables 4.1 and 4.2 provide more details about the specific motivations and behaviours named in the architecture. That is, the physiological drive used to calculate each motivational intensity, the external stimulus associated with each drive (the presence of which will increase the motivational intensity of the motivation associated with that drive) as well as the optimum values for each physiological variable (set point) and the point at which the system fails (the limit is the point at which one of the physiological variables is less than or equal to 0).

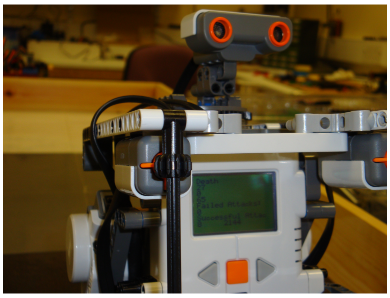

Figure 4.2: Illustration of the physical design of our basic prey agent, used in our initial experiments. We built this basic prey agent, a LEGO NXT robot, so as to take inputs from one ultrasonic sensor (with which to detect the presence of the predator) two touch sensors (with which to detect obstacles, whether that be the presence of the predator or the arena walls) and a light sensor (with which to distinguish between the various "resources" represented in our arena). We programmed our motivation-based architecture and "internal physiology" into the software of our robot using Lejos.

Motivations in this case are tendencies to behave in certain ways as a result of internal and external factors. Internal factors refer to the deficits for each physiological variable at a given time. For instance, when either blood sugar or vascular volume fall to less than their set point (100) in value, the resulting deficit is used as a measure of the internal incentive to increase that physiological variable (as calculated by subtracting the current value for blood sugar/vascular volume from this set point).

However, in addition to internal stimuli, action selection is also influenced in the implemented architectures by the presence of external stimuli or environmental cues that allow the execution of (consummatory) behaviours to satisfy bodily needs. This idea is based on ethological studies that show an animal's behaviour depends on the strength of both external and internal factors. In the current architecture, calculations of intensity therefore attempt to factor in, not only internal incentive cues (which includes the deficits, if any, with regards to each physiological variable at any given time) but also the presence or absence of external cues, including food/water resources. This was achieved using the following formula to calculate intensity:

$$Motivational\ Intensity = Deficit + (Cue \cdot Deficit) \tag{4.1}$$

Motivations in this architecture are calculated at each time-step (a single action loop) by combining external stimuli (incentive cues such as the presence of water as shown by $c_{water}$ and internal stimuli (the physiological deficits such as $d_{vascular\ volume}$. These motivations then direct the competition between behaviours to satisfy the physiological needs (for instance, $B_{drink}$). The winning behaviour, as in this implementation we use winner-take-all arbitration, will modify the physiology of the robot and send commands to the actuators.

Resulting Action Loop: Thus, at every time-step (simulation step or cycle of the architecture):

1. The intensity of each motivation (hunger and thirst) is calculated as the drive/error of its controlled variable.
2. This is used to make a calculation of intensity that factors in the presence or absence of an (external) incentive cue (for example, 0 if no food is present, but 1 if food is present.
3. The behaviour with the highest intensity is selected to be executed (we are implementing an architecture with WTA arbitration)[1]

As much of our research is aimed at the study of brain-body-environment interactions in artificial agents, to be useful case studies for our experiments our agents had to capture (and allow us to describe them in terms of) key elements of each of these core components. As a result, for instance, our predator and prey's "brains" are represented in both our model and implementation by the action selection architecture used (a motivation-based architecture, using Winner-Take-All (WTA) arbitration — see Figure 4.1). Each agent's "body" is defined as encompassing both their software and hardware capabilities, primarily consisting of their internal physiology, "brain", internal and external sensors — which, in the case of the prey, also includes the hormone-like mechanism itself. Finally, the "environment" of each agent can be classified as containing both the presence of the other agent (as a predator-prey scenario) and the distribution of resources in the constructed arena.

Figure 4.2 shows our LEGO prey. The design of this robot, in terms of hardware configuration, was developed incrementally. Following this, our basic prey was equipped with three actuators (motors): two to drive the wheels and body and one to rotate the "head" (though, in the experiments conducted to date, only the two driving the wheels have been made use of so far). In addition, we equipped it with two touch sensors to act as bumpers in obstacle avoidance; a light sensor, used to detect and distinguish between "resources"; and an ultrasonic sensor (which could still be developed further so as to enable the robot to detect or recognise other robots moving towards it). The reading from our basic prey's light sensor is translated into external stimulus or environmental cues for the action selection mechanism, while the motors provide differential steering to the wheels to navigate our agent through its environment, avoiding walls and identifying any resources in its path.

---

[1] In our architecture, if two motivational intensities are equivalent, the system will randomly select one of the joint winners

#### 4.2.1.1 The Hormone-Like Mechanism (Interoceptor Modulation)

As discussed previously, to create our initial "emotion-based" architecture, we designed and implemented a hormone-like mechanism similar to the one used by Avila-García [96]. Implemented in our LEGO prey, this was used to alter the perception that our prey has of its own integrity level when the predator is detected. The mechanism itself models hormonal release and dissipation, taking its inspiration from the artificial endocrine system proposed by Neal and Timmis [97; 98] — Equations 4.2, 4.3 and 4.4.

The first element is a gland ($g$) that releases the hormone as a function of the intensity of the external stimulus predator ($S_{predator}$) and a constant releasing rate $beta_g$:

$$r_g = \beta_g \cdot S_{predator} \tag{4.2}$$

The concentration of hormone ($c_g$) suffers two opposite forces over time. It increases with the release ($r_g$) of hormone by the gland, and dissipates or decays over time at a constant rate $\gamma_g$:

$$c(t+1)_g = min[(c(t)_g \cdot \gamma_g) + r_g, 100] \tag{4.3}$$

The hormonal concentration is limited to a maximum of $c_g = 100$ in order to keep more control on the hormone's dynamics and help in analysis of results. In our experiments we set an initial release rate of $\beta_g = 25$ and a decay rate of $\gamma_g = 0.98$. These values were set by trial and error prior to the experiments. It can be observed how the hormone is released when the predator is detected ($S_{predator} > 0$) and how it decays with time. In our initial implementation, the hormone only affected perception of the level of integrity (the third variable, introduced in the H3RP and affected by the presence of the predator).

The higher the hormone concentration, the lower the reading of the integrity interoceptor $v^1_{integrity}$. That is, the adjusted level of integrity perceived by the prey:

$$v^1_{integrity} = max[v_{integrity} - (\delta_g \cdot c_g), 0] \tag{4.4}$$

$\delta_g$ (here taken to be 0.5) determines how susceptible to hormonal modulation the integrity interoceptor is. We use $\delta_g = 0.5$, which implies that the level of perceived integrity drops by 50 when the hormonal concentration is at its maximum ($c_g = 100$). In other words, although the level of integrity is at its ideal value ($v_{integrity} = 100$), the interoceptor will perceive a level of just 50 if the hormone concentration is at its maximum. Note a constraint exists to avoid the level of integrity to be perceived beyond the lethal boundary or minimum possible value ($v_{integrity} = 0$).

Computation of the integrity deficit $d_{integrity}$ is then carried out as usual, but using the modulated perception of the integrity level:

$$d_{integrity} = (100 - v^1_{integrity}) \tag{4.5}$$

Hormonal secretion follows the detection of the predator and decreases the perceived integrity level. Using the hormone's temporal dynamics, the modulation will be acting in the system after the predator has disappeared from sensory inputs. The concentration of hormone thus modifies one of the inputs to the architecture. As described in [96] and as we have also found from our experiments, this in turn biases action selection. Namely, by increasing the prey's perceived integrity deficit when the predator is nearby. Figure 4.1 shows the prey architecture, used in our initial experiments, when implemented with the hormone-like mechanism. Avila-García called such new architectures *modulated*.

If both prey and predator "brains" can be said to consist of a motivation-based architecture selected as the underlying action selection mechanism, the hormone-like mechanism is somewhat a part of this for the prey: connecting our prey's "brain" to both internal "body" (including its internal physiology — via internal sensors/interoception) and external "environment" (or at least, perception thereof). For example, the hormone-like mechanism from the initial experiments was connected to the external stimulus of the predator via an ultrasonic sensor (used to detect the presence of the predator).

### 4.2.2 Predator

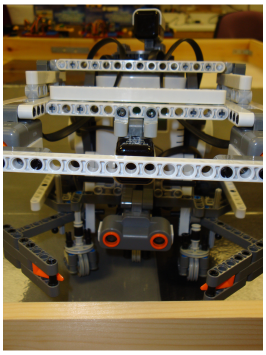

Figure 4.3: Illustration of the physical design of our basic predator agent, used in our initial experiments. We built this basic predator agent, a LEGO NXT robot, so as to take inputs from one ultrasonic sensor (mounted close to the ground so as to detect obstacles, namely the arena walls) two touch sensors (with which to detect and attack the prey) and an infra-red sensor (not used in our initial experiments, but developed for use in later experiments). As with our prey, we programmed our motivation-based architecture and "internal physiology" into the software of our predator robot using Lejos.

Apart from the hormone-like mechanism, Figure 4.1 also provides an illustration of the motivated behaviour-based model underlying the architecture used for the predator robot, while Tables 4.1 and 4.2 provide more details about the specific motivations and behaviours named in the architecture. This includes the physiological drive used to calculate each motivational intensity, the external stimulus associated with each drive (the presence of which will increase the motivational intensity of the motivation associated with that drive) as well as the optimum values for each physiological variable (set point) and the point at which the system fails (the limit is the point at which one of the physiological variables is less than or equal to 0).

In our earliest experiments there was no real choice for the predator's architecture to be used to make, as we gave it only one motivation. However, in later experiments the decision was made to give the predator the task of its own TRP. The rationale behind this was that it would add additional dynamics to the environment and overall system. That is, varying the predator's attack patterns and producing a different variation of the H3RP for the prey to be tested in. However, because the emphasis was more on prey performance, the predator's main task was still to pursue the prey; achieved by not associating an incentive cue for the additional motivation — the motivation representing the need to rest was simply calculated as the drive/error of its controlled variable (energy).

The abilities of our basic predator were (for the initial development of our implementation) kept constant and limited. Thus, our basic predator was given very limited abilities for hunting and an ultrasonic sensor, mounted low to the floor, for obstacle-avoidance upon encountering the arena wall. In fact, the "hunting strategy" of the predator here was to blindly search the arena for the prey agent. Should the predator then come into contact with the prey, a signal was sent via Bluetooth effectively attacking the prey and decreasing its internal physiological variable (integrity) accordingly. Through experiences gained from experimentation, as well as to contrast the results with previous researchers' work, the decision was made to make the damage done by the predator in such attacks quite dramatic. In fact, should a healthy prey allow/experience an attack, this would halve the maximum level allowed for this variable. This means, effectively, the prey would almost certainly die should the predator be allowed to attack more than once in a short period of time (without it having had the chance to spend any time in the nest healing its wounds).

#### 4.2.2.1 Attack of Predator on Prey

In our implementation a successful attack occurs when the predator comes into contact with the prey, which the predator detects through its hardware — namely, through its touch sensors. When successful in attacking, the predator sends a signal to the prey via Bluetooth, which alters the prey's internal integrity variable by the value set as the damage caused by one successful attack. In turn, the variable associated with the predator's need to hunt (blood sugar) will, upon the conduction of a successful attack, also be altered.

## 4.3 Arena

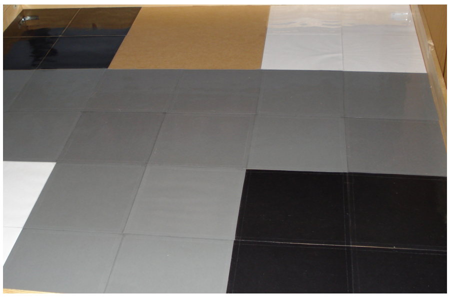

Figure 4.4: Illustration of the basic physical design of our animats' external environment, used in our initial experiments. Different coloured floor tiles represent different resources: black represents food, white represents water, grey represents resource-free (empty/neutral) space and the absence of any tile (the arena floor itself, a tan colour) represents the prey's nest, where our prey is "safe" from attack.

The basic environment (in this case an arena) had to contain both predator and prey, providing the prey with the three types of external resource; food, water and a nest; the latter of which the prey can effectively use to escape from and recover from any injuries caused to it by the predator. As with the original experiments, these resources were kept static. However, from the beginning the plan has been to develop this aspect of the environment in later stages of the research programme. The chosen environment, as experienced by the robot, consists of a 1.2 by 1.2m arena, surrounded by a wall (approximately 7cm in height). The floor is made up of 36 tiles, each one measuring 20x20cm.

In order to replicate the H3RP, the three different resources (food, drink and the prey's nest) had to be represented within the environment. To do this, we used different coloured tiles to represent different resources. Thus black tiles were used to represent food, white to represent water, grey to represent an empty space, free of resources; and the absence of any tile (the arena floor, differentiated by its light intensity, which lies between that of the white and grey tiles) was used to represent the nest. Although the original implementation of the H3RP had used lightness and darkness gradients for resources of heat and energy, it was decided that, at least to start with, distinctions could just as easily be made using these solid blocks of colour.

## 4.4 Measures of Adaptive Value

With regards to our research question, looking at how we will measure “adaptive value” for our agents is actually of as much interest to us — as it is to other researchers as well — as how we create a scenario for manipulating agents perceptual abilities, studying aspects of both an emotion-based architecture and predator-prey environment (the H3RP). Thus, our main research question actually encapsulates another. That is, how can we measure the “adaptive value” of the emotion-like component of our action selection mechanism? In more general terms, this can equate to how we might measure the adaptive value of any mechanism for an agent. Or, indeed, how adaptation itself, in terms of adaptive behaviour, might be measured (that is, given a quantity). To address this in our research, we have used a number of different measures of “adaptive value” to perform our analyses. We have done this, in part, so as to enable a comparison of these measures themselves. Both to get an idea of their relative advantages and disadvantages, as well as to provide us with a more detailed picture of adaptive behaviour with regards to our own artificial agents. This can also be thought of as an attempt to better define what we might mean by adaptive behaviour or what it means for an agent to be “adaptive”.

Within the literature, researchers have used many different kinds of calculations to compare the performance of their agents under different conditions. We have chosen a select group of these, which can be categorised into two sub-groups. The first group we have selected consists of measures that are commonly known and labelled as “viability indicators” or “measures of viability”. By choosing this group, we thereby enable direct comparison between our work and that conducted using the original H3RP. The second group, however, consists of measures that we consider as measures of activity cycles and action patterns. This group was chosen to continue our ideas in light of ideas regarding the re-conceptualisation of action selection in terms of activity cycles. Generally, our aim was to look at the impact across all these measures under different experimental conditions — for the production of adaptive action selection from the prey agent. However, for our own work, we also create graphs mapping aspects of our agents’ (particularly prey) brains, bodies and environment over time. For reference, we call these Brain-Body-Environment maps.

### 4.4.1 Measures of Viability

This group of measures collects together metrics focusing most strongly on the quantification and summary measures of the internal bodily state of the prey across its lifetime(s). That is, ignoring details of interactions between agent and environment — such as which behaviour was selected when and which environmental feature was encountered when — to concentrate on overall viability or homeostatic capacity. Such measures are usually advantageous in as much as they provide a simple, unitary, measure for each instance (or, in this case, lifetime). This allows simple comparisons to be made quickly. For our purposes, four main measures of viability (as measures of prey performance) were used by us to measure and compare the performance of our agents and evaluate prey performance across the experimental conditions described. These are Life Span, Overall Comfort, Physiological Balance and Risk of Death (as defined by Avila-García [29]).

#### 4.4.1.1 Life Span

Perhaps the most obvious measure of adaptive value, this measure can be thought of as a measure of the length of time an agent survives. Time could be measured in different ways. In his work, Avila-García calculated this with the following equation:

$$LS = t_{life}/t_{simul} \tag{4.6}$$

Where $t_{life}$ is the number of simulation steps that the agent lived and $t_{simul}$ is the total simulation time measured in number of simulation steps.

#### 4.4.1.2 Overall Comfort

This is a measurement that can be thought of as a measure of the average level of the internal variables on which the agent's life (viability) depends.

$$OvC = \sum_{1}^{t_{life}} (c_{step})/t_{life} \tag{4.7}$$

Where $t_{life}$ is the number of simulation steps the agent lived, $c_{step} = 1 - (Err/max_{err})$, where $Err$ is the total sum of errors of the agent's physiological variables normalised between $[0,1]$ with $max_{err}$ the worst possible in each step. $Err$ equates to the sum of the intensities of the motivations' drives $(Err = \Sigma(e_j))$, and $max_{err}$ is the number of compatible motivations.

#### 4.4.1.3 Physiological Balance

This is a measurement that can be thought of as how well the agent's internal variables, on which life (viability) depends, are balanced. This average is calculated as follows:

$$PhB = \sum_{1}^{t_{life}} (b_{step})/t_{life} \tag{4.8}$$

Where $t_{life}$ is the number of simulation steps the agent lived, $b_{step} = 1 - (Unb/max_{unb})$, where $Unb$ is the variance of the errors of the agent's physiological variables normalised between $[0,1]$ with $max_{unb}$ (the worst variance possible in each step), which corresponds to the variance of the intensities of the (compatible) motivations' drives $(Unb = \sigma(e_j))$.

#### 4.4.1.4 Risk of Death Management

Finally, "Risk of Death" can be thought of as a measurement of "how close" death is (to the agent) at any given time. Calculated by:

$$RoD_i = steps_i / t_{life} \tag{4.9}$$

Where $i$ is the number of the region ($i$ = 0..10) and $steps_i$ the number of steps the agent's highest physiological need (in this case, the variable with the highest deficit) is within the region $i$. For our experiments, we calculate Risk of Death for ten distinct regions, giving us values for RoD1 to RoD10 for comparisons. RoD1 represents the proportion of time our agent's highest physiological need was in region one i.e. at a deficit of 0-10 percent, with little chance of death; whereas RoD10 represents the proportion of time our agent's highest physiological need was in region ten i.e. at a deficit of 90-100 percent, where death is imminent. This granularity was decided upon both so as to be more comparable to previous research and as a result of our initial experiments. However, it would be interesting to investigate the linear assumptions of this metric, as well as how calculating values for further regions might be more useful later on.

### 4.4.2 Activity Cycles and Action Patterns

This represents a group of measures we introduce to compare the information they give us with that of the first. The idea here was that other measures might be developed to complement — or otherwise provide better information than — those used previously. Inspired in part by the idea of action selection as a series of activity cycles, this led us to suggest perhaps what is missing (to get a more complete picture of the success or not of such agents) are measures that look more at the temporal effects of manipulation on such agents — in terms of behavioural sequence, for example. We therefore created this group of measures to relate more strongly to summarising activity cycles of the prey across its lifetime(s). That is, rather than attempting to provide a more summarative measure of overall success, looking at how the system becomes more or less "stable" (and in what ways) in terms of the behaviour displayed across conditions.

In this way, these measures do not ignore so many of the details between agent and environment as the first group. For instance, taking into account the fact an agent might have started searching for a resource, then switched to satisfying another when the resource representing it was encountered first. This includes which behaviour was selected when and which environmental feature was encountered when and concentrates on building more of a profile of the actions selected than an overall summary of whether that profile was "better" in terms of agent survival. While likely to be more time-consuming in terms of a more length analysis, we argue such measures will be advantageous in as much as they can provide more detail, building a profile that may suggest how the system should be adjusted in order to improve performance in terms of the more summarative measures of viability.

### 4.4.3 Brain-Body-Environment Maps

In order to create a map that incorporates the (to us) crucial aspects of brain, body and environment in our experiments, we create plots for individual runs from the data collected from our implementation of the H3RP. These map our agents' internal variables over time, as well as the environmental stimuli our agents encounter. We have found these useful in our research as they enable us to (effectively) recreate our animat's "life" and brain-body-environment interactions. That is to say, they incorporate and enable us to track key elements of the brain-body-environment interactions represented in our model.

## 4.5 Summary

In this chapter, we have outlined the more technical details of our implementation of the H3RP. That is, the physical instantiation of our model of action selection, emotion, brain-body-environment interactions and the predator-prey relationship (previously outlined in more detail in the previous chapters). To create our own H3RP we have designed two LEGO robots — one predator, one prey — along with an environment for them to interact or "live" in. As discussed, we selected one specific emotion-based architecture because, in contrast to other researchers' implementations of emotion-based architectures, it offers us a possible route for investigating, creating and exploring the possibilities of artificial emotion in a bottom-up way. This is more in keeping with the spirit of AL in terms of allowing us to think about emotion not only as it is, but as it could be. Yet, at the same time, the mechanism we adopt also happens to relate to (as it is essentially a biologically-inspired abstraction of) physical substrates in biological brains. In particular, it relates to ideas about the role of neuromodulation in generating adaptive behaviour.

# Chapter 5

# Brain-Body Combinations and Predictions of Predator Attack

*"The emotions aren't always immediately subject to reason, but they are always immediately subject to action."*
**— William James**

## 5.1 Emergent Functions of Interoceptive Modulation: Results Chapter I.

The literature linked to and concerning our concepts of the advantages and contributions of emotion in the context of adaptive behaviour is not just restricted to humans. For, not only has it been extended by researchers to include evidence from the study of other biological organisms [99] but, more recently, has been yet further extended to encompass evidence provided by artificial agents as well [20; 90]. Whether considering emotions in terms of the past [100], present [101] or future [102] this has led to the development of explanations which suggest, or at least allow, that organisms other than humans might also "have" some kind of emotion(s). For example, from an evolutionary perspective, explanations of our own emotions have been given in terms of a gradual evolution from simpler and/or earlier mechanisms, present in our early ancestors and perhaps still evident in other less "evolved"[1] or otherwise less "complex" species today.[2]

---

[1] By this, we refer to species whose more recent ancestors, at least compared to our own, do not appear to have experienced such rapid change

[2] Here we would emphasise the point that, even if only considered to be "emotion-like" mechanisms, not as "evolved" or intricately "developed" as our own, these give their own evidence for the adaptive

Whether this evidence comes from studying early ancestors [103] or species currently in existence [99], this branch of research has encouraged and led many researchers, including us, to adopt an incremental, bottom-up approach — both for explaining and exploring the origins of intelligence and adaptive behaviour. In other words, exploring the idea that our "modern" human brains have been developed over generations from earlier brains, and presumably less intelligent ones, by incrementally adding or subtracting elements to existing structures. In tune with our research, such explanations often go on to postulate evolutionary pressures being exerted by competition or threats to survival arising from the presence of other organisms. This might be in terms of conspecifics, other species or, most interestingly for our current research, predator and prey organisms [104; 105].

Many different functions of emotion are proposed [106], as well as many different classifications [106; 107; 108] (and inevitable variations on these classifications) of these emotions. However, when we consider the idea that our seemingly complex emotions, or at least a collection of mechanisms serving the same function, have evolved from simpler mechanisms (whether having persisted from one common ancestor or co-evolving independently of each other) it is amazing that they seem to "work" for so many species and across such a wide variety of agent morphologies, niches and capabilities. More to the point, if true, this means we all stand as a very real and practical demonstration of how a bottom-up approach to the design of intelligent and adaptive agents can work, not to mention proof that it does. In its turn, our work echoes this observation and tries to discover whether we can identify (and have identified within our chosen model) similar mechanisms that might likewise work, or otherwise be developed to work, adaptively across many different *artificial* agents (including different types of agent morphologies, niches and capabilities).

## 5.2 Fear and the Fight-or-Flight Response

Perhaps the most common emotion, and associated function of emotion, recognised and studied as being adaptive for action selection in the predator-prey relationship, particularly in terms of the notion of adaptive behaviour, is that of fear in the "fight-or-flight" response upon prey encountering predator. Indeed, in terms of action selection, many have argued one of the adaptive functions of fear is to act as an amplifier of sorts, calling attention to the high risk and dangerous elements of an agent's current environment. In the original Hazardous Three Resource Problem (H3RP) this was in fact noted as an emergent functionality of the emotion-based architecture also implemented in our robots. For this reason, to start with we focus our own attention on this proposed type and function of emotion. To explore the importance of body and environment further, we first use our own model to demonstrate how the emergent

---

powers of emotion in terms of the *persistence* of such mechanisms across generations properties (such as prey “flight”) of our emotion-based architecture will vary with our prey’s brain-body combination. Using our model of action selection and emotion, the rest of this chapter reports our findings. It is in this way we start the investigation into the possible role of our hormonal mechanism, as a modulator of perception — and second order controller of a motivation-based architecture — on brain-body-environment interactions within a prey-predator scenario.

Our experiments show how the adaptive value of our mechanism for action selection varies with a particular aspect of the prey’s embodiment: namely, perceptual distance. By this, we refer to how far into its environment an agent can “sense” or “perceive”. To borrow a term from the field of neuroscience, this can also be thought of as our agent’s *receptive field*. Here we modulated the architecture of our prey robot using two different types of sensory capabilities, proximal and distal, effectively creating combinations of different prey “brains” and “bodies” (where perceptual distance is varied by creating different sensory links between prey “body” and prey “environment”). The results of these experiments were analysed using the different measures outlined more fully in the preceding chapters. These aimed to further explore the way in which an action selection mechanism’s adaptive value and overall performance can be judged.

## 5.3 Brain-Body Combinations for our Prey

We designed variations of the H3RP by varying the sensors available to our the robotic prey for detecting predator presence — in this case either restricting the prey’s perception to sensory information obtained close to and directly from its own body (proximal detection) or allowing it to gather information corresponding to stimuli at some distance from the prey’s own body (distal detection). For our basic implementation, this created two variations of prey, each of which we then tested in our scenario so as to see what effect the addition of an “emotion-like substrate” — our hormone-like mechanism — would have on their performance. In our experiments, we created a total of four different variations of the H3RP by varying both our basic prey’s *perceptual distance* (proximal/distal) and the *hormone-like mechanism* attached to its basic motivation-based architecture (presence/absence). These prey agents were tested against one type of predator (proximal predator). In these initial stages, we thus kept the predator’s abilities at their most limited, creating four experimental conditions (a factorial 2x2 design, summarised in Figure 5.1). To these conditions we also included a control condition, in which the prey was first tested without the predator. This latter condition we used so as to provide us with a baseline with which to better see if/how the addition of our predator would actually affect prey performance.

| 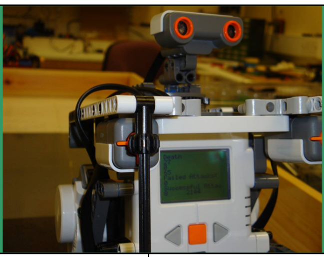 | | *Variations in Prey "Brain" [Emotion-based architecture]* | |
|---|---|---|---|
| | | **No Hormone Release** | **Hormone Release** |
| *Variations in Prey "Body" [Perceptual Distance]* | **Proximal Prey** | Proximal Prey<br>H- | Proximal Prey<br>H+ |
| | **Distal Prey** | Distal Prey<br>H- | Distal Prey<br>H+ |

Figure 5.1: Summary of experimental conditions: In these experiments a factorial 2x2 design was used to investigate how the adaptive value of our emotion-based architecture varies with the perceptual abilities of our prey. Four experimental conditions were created so as to identify the effects of a) the “emotion-like” substrate represented by our hormone-like mechanism, presence (H+) or absence (H-) representing a variation in our prey’s “brain” and b) our prey’s perceptual distance, proximal/distal detection representing a variation in “body”.

## 5.4 Prey

### 5.4.1 Interoceptive Modulation (Absence/Presence)

As the purpose here was to get a basic idea of how aspects of embodiment and modulation might interact to affect performance of the system, we focused first only on the simple absence or presence of the hormone-like mechanism (outlined, along with more technical details of our implementation, in the previous chapter).

### 5.4.2 Perceptual Distance (Proximal/Distal)

As the aspect of embodiment we chose to focus on, we aimed to vary our prey’s perceptual distance by connecting the basic motivation-based architecture to two different kinds of ability. The first we termed “proximal prey” and the second “distal prey”. These can also be thought of as a “blind” and “seeing” prey respectively.

#### 5.4.2.1 Proximal Prey

To recap the key features of this animat, also outlined in the previous chapter, the proximal prey was given a motivation-based architecture consisting of an internal physiology of three physiological variables. Two of these were set to decrease with time at a constant rate (we named these blood sugar and vascular volume respectively) while the third was connected only to “attacks” registered by the prey and made by the predator robot (arbitrarily named integrity). The hormone-like mechanism, if “switched on”, causes “hormone-release” in the presence of the predator, as signalled by an internal sensor, registering a decrease in the level of the third variable. That is, after a physical attack. Our first type of prey was thus given a very limited perceptual system. In this case the only sensory information available to the prey robot from its environment came from touch (via its touch sensors) and internal perception of the physiological variable of integrity. This scenario could be thought of as limiting the prey to a single perceptual system, akin to the somatosensory system present to many organisms. However, this system is particularly limited in the fact that it leaves the prey more or less “blind” to actual events in its environment. For instance, the prey cannot actively (by touch alone) discriminate between the stimuli of a wall or that of the-more dangerous-predator. Indeed, the only signal to the prey that the predator is present comes from internal changes following an attack. Because of this, hormone release, or what could alternatively be thought of as neuromodulation, would occur only after the prey has been attacked (integrity decreases).

#### 5.4.2.2 Distal Prey

The distal prey differed from the proximal in that the hormone-like mechanism, if “switched on”, caused hormone-release in the presence of the predator, as signalled not only by the internal sensor, but also the detection of the predator via use of an ultrasonic sensor, mounted on the prey’s “head”. To give a contrast to the proximal, no-hormone condition the no-hormone architecture for the distal prey condition incorporated a “reflex” behaviour in the form of automatic obstacle avoidance, should the predator be detected by the ultrasonic sensor. In this way, the second type of prey was given a form of “distal” detection to be used to detect the predator at a distance (which, when the hormone-like mechanism was added, would also trigger a corresponding hormone release; otherwise triggering reflexive obstacle-avoidance behaviour). This involved only a very small change; the prey robot now using an ultrasonic sensor to also check with each time-step whether the predator could be “seen” at a distance (distal detection). This scenario could perhaps be thought of as adding a (very) rudimentary auditory-imaging system to that of the prey’s existing perceptual system(s) in as much as the prey can now obtain information from both a form of “somatosensory” system and an additional “auditory-imaging system”.

## 5.5 Predator

### 5.5.1 Perceptual Distance (Proximal)

The perceptual distance of the predator was, for these experiments, an aspect that was kept constant. Here, the predator was given very limited abilities for hunting. In fact, the "hunting strategy" of the predator here was to blindly search the arena for the prey agent. Should the predator then come into contact with the prey, a signal was sent via Bluetooth effectively "attacking" the prey and decreasing its internal physiological variable (integrity) accordingly. Through experiences gained from experimentation, as well as to contrast the results with previous researchers' work, the decision was made to make the "damage" done by the predator in such attacks quite dramatic. In fact, should a "healthy" prey allow/experience an attack, this would halve the maximum level allowed for this variable. This means, effectively, the prey would almost certainly "die" should the predator be allowed to attack more than once in a short period of time (provided it had not spent any time in the nest "healing" its wounds).

As detailed in the last chapter, the decision was made to give the predator the task of its own TRP. The rationale behind this was that it would both add a further dynamic, varying attack patterns and producing a different variation of the H3RP for the prey to be tested in. However, because the initial emphasis is placed that bit more on prey performance, the predator's main task was still to pursue the prey. This was achieved by not associating an incentive cue to the additional motivation (need to rest). Instead, this motivation was simply calculated as the drive/error of its controlled variable (energy). As mentioned, throughout these experiments the predator was, in its turn, also given extremely limited perceptual abilities to hunt the prey. The only way the predator could find the prey being to blindly search its environment and wait for the prey to somehow come into contact with its touch sensors. Unlike the prey however, the predator was able to discriminate prey from the arena wall due to its morphology. That is to say, the touch sensors were mounted higher than the arena wall.

## 5.6 Arena

The layout of the arena (resource distribution and abundance) was kept constant. This consisted of a nest area, in which the prey could "rest" and where it would not be "attacked" by its predator. It also consisted of food and water resources, along with some neutral space in which the prey could not consume a resource.

## 5.7 Results

For each run, both predator and prey were initially placed within the arena in randomly-selected locations (and positions) with randomly-generated starting levels for each internal physiological variable (between 70 and 100, apart from prey integrity, which was always initialised at its maximum level of 100). Each run was considered to be complete, either when the prey "died" (involving at least one of its physiological variables reaching a value of 0) or the prey had reached the end of its pre-determined maximum life span, measured in time-steps. After each run, data was collected from both prey and predator robots with regards to sensory and internal (physiological) data. The prey was set to record, per three time steps, the following data: current time-step, level of each physiological variable, sensory readings, whether the prey had been attacked, the winner motivation, winner behaviour and the times taken to perform each of the previous three time steps. We tested each type of prey robot for a total of ten runs of 3000 steps (which would equate to approximately ten minutes per run if the prey survived), one step representing a loop of the behaviour selection mechanism. Fifty runs were therefore conducted, as ten runs were conducted without the predator so as to provide a baseline for comparison.

### 5.7.1 Measures of Viability

From the data collected, we first calculated and looked at the results in terms of our chosen measures of viability. These included calculations of Life Span, Overall Comfort and Physiological Balance (as defined in the last chapter) as well as calculations of our prey's Risk of Death Management. Figure 5.2 shows how each type of prey performed as measured by our three measures of viability: Life Span, Overall Comfort and Physiological Balance. To enable further comparisons, we have also included the results obtained testing our prey alone, without a predator agent, making this condition essentially the results obtained using our prey to perform a two-resource problem. Figure 5.3 then shows the same results for our predator agent.

#### 5.7.1.1 Life Span, Overall Comfort and Physiological Balance

First looking at Life Span (how many time-steps out of the maximum that the prey lasted for — defined more fully in the previous chapter) Figure 5.2 shows the results in terms of the mean Life Span calculated per condition (error bars show the standard error of the mean). The measure of Life Span, previously used for the H3RP and defined in the last chapter, was originally defined so as to be calculated from the total number of *time steps* that the prey lasted for — that is, the number of individual "decisions" made over this agent's life time. In our experiments what we are particularly interested in is real time data, from a real time environment. Because of this, another decision was made to record the real-time Life Span of our agents as well. We did this primarily in order to compare the two measures, to act as a further control. However, we also wanted to see whether there might be any interesting differences. In fact, both these calculations produced very similar results for these experiments. Nevertheless, both measures of Life Span hold a special appeal, in as much as they are particularly reflective of the aspect of action selection we are focusing on. That is, action selection in terms of *temporal space*.

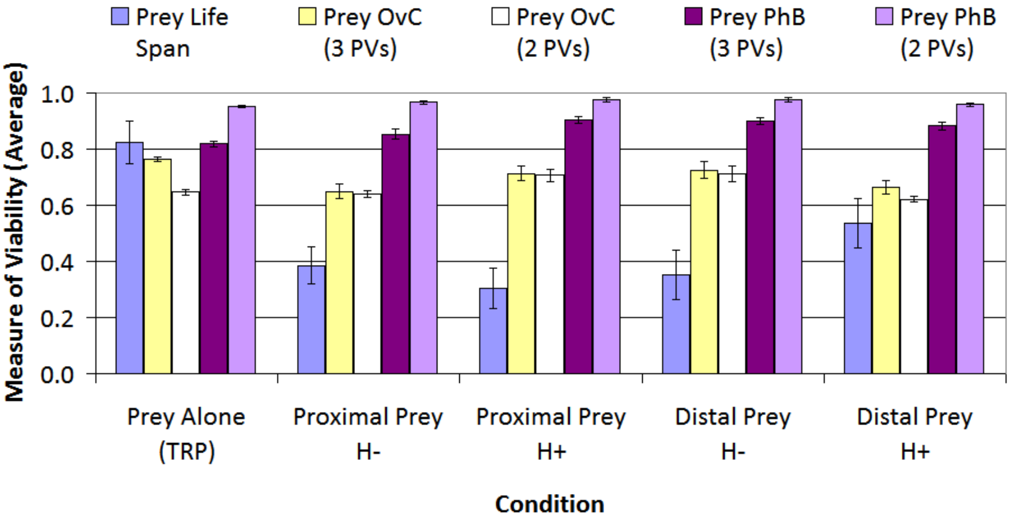

Figure 5.2: Summary of results showing average Life Span (LS), Overall Comfort (OvC) and Physiological Balance (PhB) of our prey across conditions. Error Bars show the Standard Error of the Mean (SEM). We compare calculations of OvC and PhB both including (3PVs) and excluding (2PVs) the variable integrity. In terms of LS, the longest lived prey tended to be those with both "brains" in which our emotion-like substrate was implemented (present) and "bodies" which enabled distal detection of the predator. The other conditions show more similar tendencies (less difference) between them for LS. However, we observe how the activation of our emotion-like substrate, when combined with proximal detection only, actually led to the lowest average LS. In terms of OvC and PhB, we see less obvious variations than for LS both within and between conditions (Error Bars being smaller). N.B statistical tests nevertheless indicate that, though small, some of these differences are statistically significant.

Realistically, once a given "decision" is made (at least, in the sort of architectures we are interested in looking at) it might take longer to perform the associated action, especially in "real life", than it would a different decision or even the same decision at a different time. For example, the reflex action of "obstacle-avoid" might, on average, take a different amount of time than the consummatory or appetitive action of eating or searching for food. Indeed, this was the definitely the case in our own implementation. Regardless of the specific calculation used, however, both do show the same trends across conditions for these experiments. The average life span for our prey when placed in the arena without the predator was found to be significantly different to all the individual experimental conditions tested (average life span decreasing). No significant differences were found between the experimental conditions. However, for the first type of prey (proximal) when the underlying architecture was a modulated one, our prey's average life span was shorter than that of the one without modulation. In contrast, for the second type of prey (distal prey) the opposite trend was actually found. That is to say, when the underlying architecture was modulated, our prey's average life span was longer than that of the one without modulation.

Looking at the means for both calculations of Life Span, our distal, modulated prey thus appeared to perform the "best" for this measure (on average) while our proximal, modulated prey performed "worst". Testing for homogeneity of variances, these were found to be homogeneous both within and between our conditions (p=0.85 across conditions). Interestingly, but perhaps not surprisingly, when we looked further at the distribution of our data, we found that while our data set did pass a test for normality (Shapiro-Wilk result of p=0.08 across all our conditions) Q-Q plots indicated that some of our samples (within-conditions) suffer from a moderate to strong skew, suggesting some non-normality. For example, we calculated a skew of 1.32 for the distal prey, without the hormone-like mechanism.

The reason why this is perhaps not surprising is due to the nature of our measures of viability. This is because it is a calculation that is guaranteed to be between 0 and 1. This may produce a more easily understood measure, but in doing so constraint/binding is imposed on the resulting calculations for comparison. Thus, a freely generated distribution around the mean may not be possible. This would be the case, for example, if the mean falls close to either limit (0 or 1). Moreover, if we have a particularly successful prey, which only dies prematurely once, we may well find a strong positive resultant skew for the Life Span results of this condition. Conversely, if we have a particularly unsuccessful prey that lives a bit longer in one run, we will likely see it result in some kind of negatively skewed distribution. Bearing this in mind, as well as the relatively small sample sizes of these early experiments, the decision was made to transform our data using arcsine transformation before performing our inferential statistics (so as to better meet the assumptions of the parametric tests we planned to perform).

Performing a multi-way (2x2) ANOVA for Life Span (as calculated by time steps) across these conditions showed no significant main effects for prey perceptual distance, (p=0.187) or the hormone-like mechanism (0.636). Likewise, there was no significant interaction found between the two (p=0.121). However, in terms of calculations of Overall Comfort (the average of the average level of all internal physiological variables, defined in the previous chapter) and Physiological Balance (the average of the

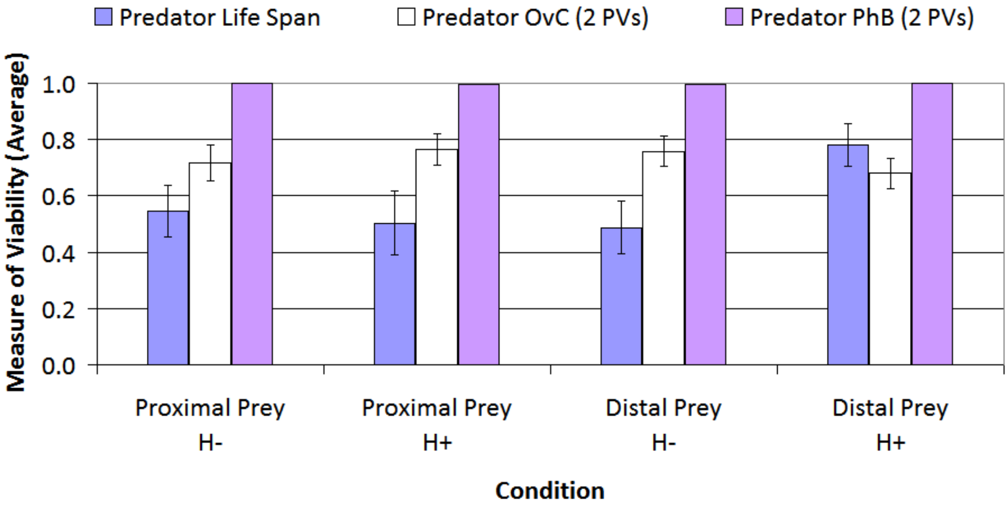

Figure 5.3: Summary of results showing average Life Span (LS), Overall Comfort (OvC) and Physiological Balance (PhB) of our predator across conditions. Error Bars show the Standard Error of the Mean (SEM). In terms of LS and OvC, the trends here generally follow those for the prey (perhaps not surprising considering these trials usually ended with the death of the prey). However, it might be observed that the lowest average LS for the predator was actually not seen in that condition where the activation of our emotion-like substrate was combined with proximal detection only. Instead, this was true of that condition where our emotion-like substrate was not activated, but our prey had distal abilities. This can be explained by the use of time-steps in calculating LS (a single decision/time-step might take different amounts of real time to complete depending upon what action is being selected). In comparison, PhB for our predator hardly varied at all between conditions.

average variance of all internal physiological variables across all time-steps) as Figure 5.2 shows, the results are slightly different. As with Life Span, we calculated both Overall Comfort and Physiological Balance in two different ways to see whether differences between them could be interesting for further research: first, combining all three of our physiological variables and then (as our own variation of the measure) concentrating only on those two physiological variables that decrease consistently with time. That is, those we named blood sugar and vascular volume. We hoped this might achieve support for our idea that the addition of our extra variable of integrity alone (and thus extension from the two resource problem to a hazardous resource scenario) might act to increase adaptive value of our agent. Moreover, we wanted to find out whether this would actually be a more informative measure as far as both Overall Comfort and Physiological Balance are concerned (the additional variable of integrity being, temporally-speaking, very different in terms of increase/decrease).

Figure 5.2 shows the results of the latter calculations. The results of this exercise do provide some interesting data for comparison. For, in terms of both Overall Comfort and Physiological Balance, it looks like the inclusion of integrity in these calculations may have the effect of decreasing the average shown, whilst increasing the standard error of the mean for our experimental conditions. This is not surprising considering the nature of this variable. That is, it will be subject to dramatic decreases, rather than the consistent, predictable decreases of the other two variables.

Performing a multi-way (2x2) ANOVA for Overall Comfort and Physiological Balance (calculated using all three physiological variables) showed significant differences between conditions. For Overall Comfort, after arcsine transformation, there were no main effects for perceptual distance (p=0.606) or hormone presence (p=0.991). However, the results did suggest an interaction between these variables (p=0.018). Again after transformation, for Physiological Balance there were also no main effects found for perceptual distance (p=0.452) or hormone presence (p=0.220). But again, an interaction was suggested between the two (p=0.022).

#### 5.7.1.2 Risk of Death Management

In terms of the prey's Management of Death Risk (as defined in the previous chapter) the individual regions for this measure can be used as an indication of how well viability is preserved. Using this measure to explore our data provides us with an interesting alternative look, and further insight into, our previous results. Figure 5.4 displays our results, in terms of Risk of Death (RoD), providing a breakdown for each experimental condition as well as our control (where the prey was acting alone). Here, each bar represents the average percentage of time our prey was in each "zone". A RoD of 1 means that the prey is almost, if not the most "healthy" it can be, whereas a RoD of 10 means the prey is "near death". However, what is interesting to us here is the similarity between our control condition and that of the prey given "distal" perception, connected

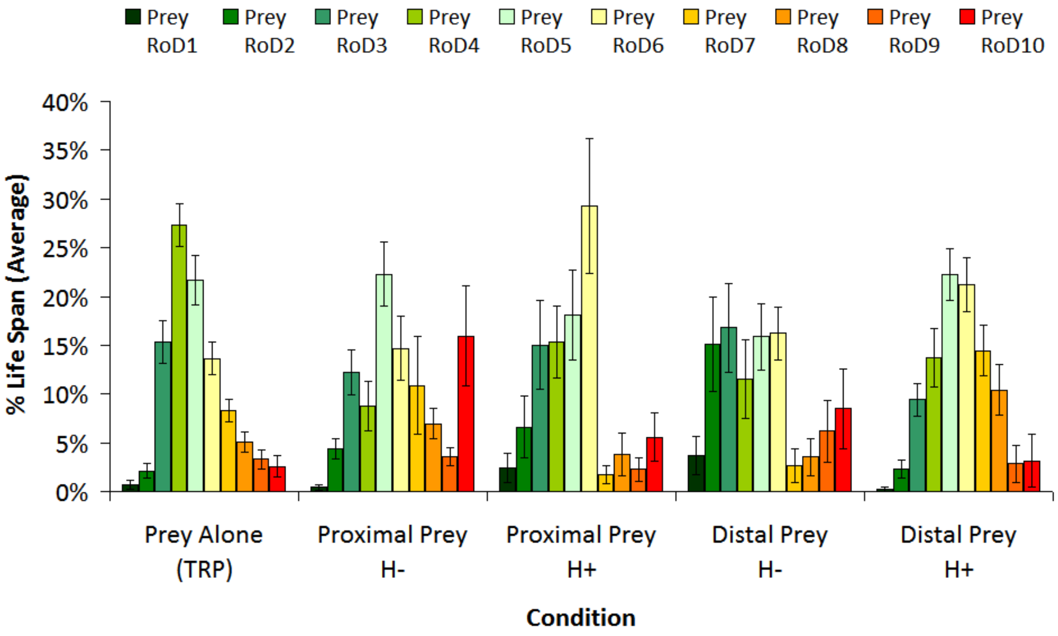

Figure 5.4: Summary of results showing Risk of Death (RoD) Management across conditions for our prey, also including the results for our control condition (Prey Alone). Each bar represents the average percentage of time our prey was in each RoD "zone". Error Bars show the Standard Error of the Mean (SEM). A RoD of 1 (green) means that the prey is almost if not exactly the most "healthy" it can be, whereas a RoD of 10 (red) means the prey is "near death". Interestingly, these results demonstrate greater similarities between our control condition and that of the prey given both "distal" perception and our emotion-based architecture with hormone-release enabled. We suggest that this is because these are the two conditions that show the most stable/adaptive activity cycles. In contrast, the other conditions show more variance, though Proximal Prey H- is interesting because it suggests that this type of prey tended to spend most time in either RoD5 or RoD10. That is, either fairly comfortable or very near death.

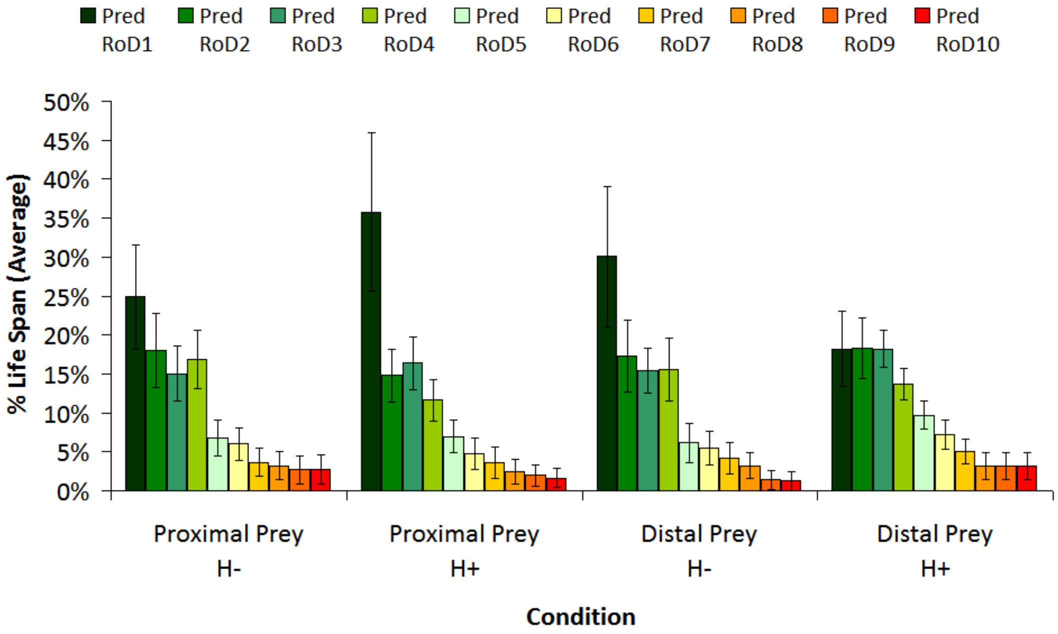

Figure 5.5: Summary of results showing Risk of Death (RoD) Management across conditions for our predator. Each bar represents the average percentage of time our predator was in each RoD "zone". Error Bars show the Standard Error of the Mean (SEM). A RoD of 1 (green) means that the predator is almost if not exactly the most "healthy" it can be, whereas a RoD of 10 (red) means the predator is "near death". These results show that, perhaps not too surprisingly, our predator actually spent the majority of time with a low Risk of Death. It appears that the more "risky" conditions for our predator to be in were those for our non-modulated proximal prey and modulated distal prey.

to our emotion-based architecture in terms of variations in viability between runs. For, both these agents seem to produce more "stable" agents in terms of viability. That is, compared to the other conditions. Figure 5.6 illustrates this in more detail, with each line per condition representing Risk of Death Management across each prey's lifetime, while Figure 5.5 shows the average Risk of Death (RoD) results for our *predator*.

Looking more closely at our data, Figure 5.7 shows some example runs, displaying the levels for each physiological variable of the prey as they changed throughout a single lifetime (increasing with consummatory activity and decreasing with appetitive or other activity) for three of the five conditions. As shown, the architecture of the prey generated reasonably stable cycles of activity when the prey was acting alone. However, upon introduction of the predator for each experimental condition, the architecture's performance caused these cycles of activity to be disrupted (the prey having to recover integrity, initialised for all our runs at its set point). Again, what is just as interesting to us here are the differences in the patterns of behaviour shown. For instance, the distal prey with modulation is seen to "respond" more rapidly to an attack, thereby also recovering integrity levels sooner. These differences, particularly when placed in context with our observations, would seem to suggest that modulation tended not to be adaptive, and was possibly even maladaptive in the case of the simpler prey with proximal perception only, not providing any advantage and carrying only costs. However, in the case of the distal prey it was adaptive as it provided additional anticipatory capabilities (a prediction of predator attack) that helped the prey avoid/escape the predator more easily.

## 5.7.2 Activity Cycles and Action Patterns

### 5.7.2.1 States and State Transitions

One of the aims of this thesis was to also consider previous work advocating the reconceptualisation of the action selection problem in terms of activity cycles. As outlined in the previous chapter, one of the ways we look to explore this idea further is (taking inspiration from ethologists) by looking at our agents' performance in terms of states and state transitions. State transitions in particular are useful calculations in terms of problems involving movement from one "state" to another. In context, a "state" may refer to many different variables, with the models which result from such data (including Markov Models) being conceptual tools that can both elegantly describe and enable further analysis of the nature of changes generated by the movement of such variables. What we are particularly interested in using them for here, in the purpose of our analyses, is to find and explore further the patterns which appear over time (and that might differ across our experimental conditions).

To do this, we must attempt to model our developed system (in other words, to build models of our model). In this case, by focusing on the action selected by our

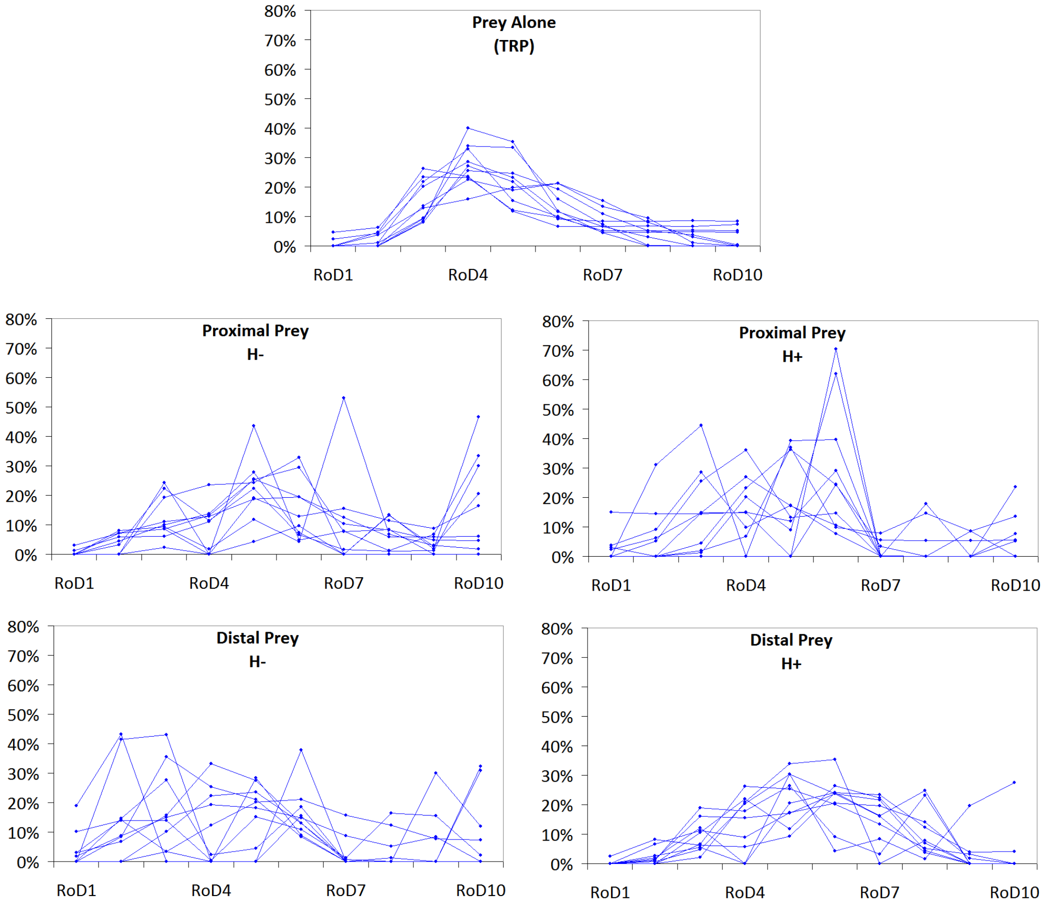

Figure 5.6: Summary of results showing Risk of Death (RoD) Management across conditions for our prey, also including the results for our control condition (Prey Alone). Each line here represents the results from one run and shows the proportion of our prey's lifetime spent in each RoD "zone". A RoD of 1 (green) means that the prey is almost if not exactly the most "healthy" it can be, whereas a RoD of 10 (red) means the prey is "near death". As with Figure 5.4, these results demonstrate greater similarities between our control condition and that of our modulated distal prey. In contrast, the other conditions show more variance, though the results for our non-modulated proximal prey are interesting in that they suggest this type of prey tended to spend most time in either RoD5 or RoD10. That is, either fairly comfortable or very near death.

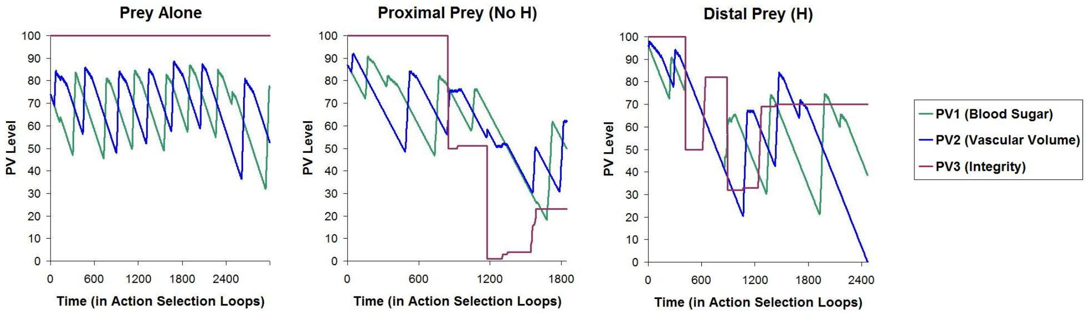

Figure 5.7: Illustration of differences observed across conditions between individual runs. This figure better highlights the types of differences we typically see across conditions by drawing upon some example runs, displaying the levels for each physiological variable of the prey as they changed throughout a single lifetime (increasing with consummatory activity and decreasing with appetitive or other activity) for three of the five conditions. The architecture of our prey generated reasonable stable cycles of activity when the prey was acting alone. Upon introduction of the predator, the architecture's performance (without the presence of our hormone-like mechanism) caused these cycles to be disrupted (prey having to recover integrity). However, upon introduction of both our hormone-like mechanism and distal abilities, the distal prey with modulation was seen to "respond" more rapidly to an attack, thereby recovering integrity levels sooner.

prey robot at any given time, we can construct state space diagrams to explore activity cycles. We could call this "behaviour" but as we will construct these models from data collected "internally" — from the prey "brain-body" — it seems more fitting to call it an action for several reasons. One of these is that behaviour tends to be considered the same as that which is "observable" while the data we collect is more "intentional" in as much as the prey "intends" to perform that action, but in reality what emerges from interaction between it and the environment may not be the same each time. For example, while the prey may record itself as searching for food, it may be that its predator somehow has it "pinned" against the arena wall and that it is actually flailing helplessly in one place.

As "simple" as our agents are, they nevertheless already produce non-deterministic patterns of actions selected. For one action does not always follow after another one hundred percent of the time. To explore this aspect of our agents, from our data we have therefore also constructed state transition diagrams. As a result of the design of our prey's architecture (detailed further in the previous chapter) at any given time it could be "performing" one of six "actions": the consummatory actions of "eating", "drinking" and "resting"; or the appetitive actions of searching for "food", "water" or the "nest". Thus our initial focus was on these states and the transitions between them. Figure 5.8 shows some sample state transition diagrams from different runs. These also show the percentage of the prey's life spent in that state.

Though no clear differences were found between conditions just by looking at these, as there was too much variation within conditions to make this supportable by statistical test from our sample, it was nevertheless interesting to note that the "busiest" prey (in terms of the diagrams showing connections from one state to another) were found in the condition for distal perception, with the hormone-like mechanism. For the proximal prey, without the hormone, it was interesting to see that these diagrams were often very similar to those found in the "prey alone" condition. Putting this into the context of the experiment in terms of our observations, this prey's brain-body combination is such that the prey only considers integrity as it would the other two variables — ignoring attacks from the predator unless it just so happens upon the nest resource. Indeed, this seems a good strategy in terms of our "blind" predator. For, from observation, it seemed that the predator was just as likely to attack, then either decide to "sleep" or else completely miss the prey the second time.

While these diagrams are useful in allowing us to see what happened in each individual run, they also allow us to consider how it might be more meaningful to categorise these "actions" so as to group together any dependent or similar sequences of action. For example, we can see in our control condition eating is usually preceded by "searching for food", but that this tendency dissipates in our experimental conditions. To explore the idea of state transitions further, we have also looked at *averages* across conditions of both how long our prey spent in each state and how often one state followed another. In particular, Figure 5.9 shows the average number of different types of "state transitions" of our prey across conditions. Exploring the idea of action selection as a cycle of activities, rather than a series of individual decisions and having constructed a state transition matrix for each individual run to show how often one action followed another, we specifically focus here on the idea of "opportunism" and "persistence". While we define opportunism as the number of times our prey consumed one resource while searching for another, persistence represents the number of times our prey ended up consuming the resource it was searching for. To this we have added a further definition (a subset of persistence) focusing only on those instances where our prey both searched for the nest and ended up resting in its nest). Here we see that, on average, all of our prey demonstrate more opportunism than persistence. Interestingly though, our modulated distal prey actually tended to display less persistence towards searching for and resting in the nest than our non-modulated prey.

To test our initial thoughts with regards to state transitions, in terms of an activity profile, Figures 5.10 and 5.11 show the average time spent on each action. Together, they show both the differences in terms of the average absolute number of time-steps that each type of prey spent on each activity, and the percentage of their "lives" that this made up. These we find interesting, if nothing else due to the fact that the differences between them show how it could potentially be very misleading to judge performance of our prey by one of these aspects alone. Figure 5.12 shows the same results for our

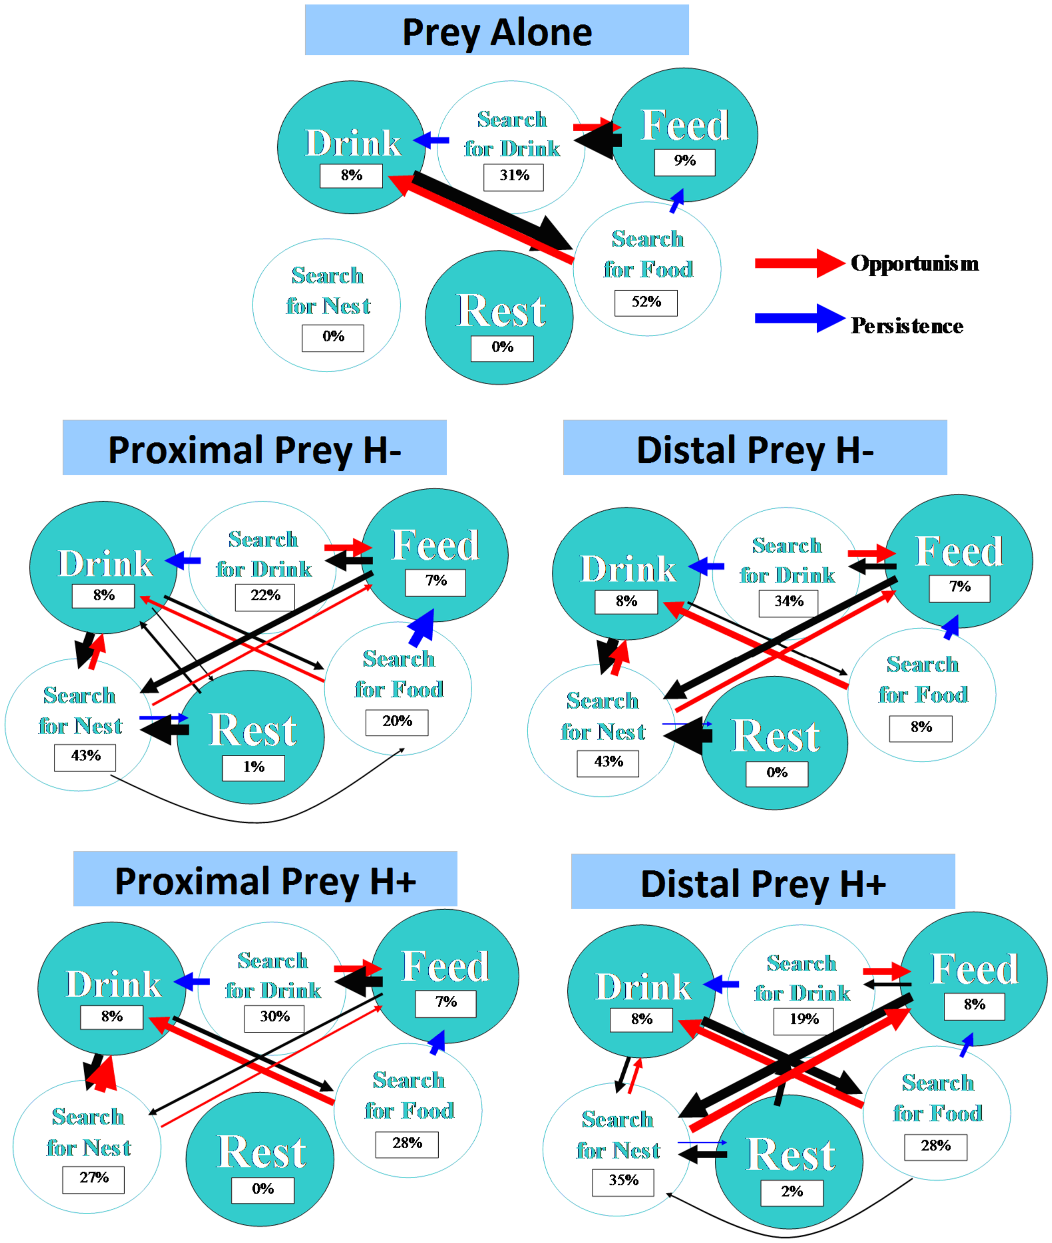

Figure 5.8: Summary of results showing “state transitions” of our animats — one example is given here from each condition to show differences between individual runs. These diagrams were constructed as a result of Markov analysis and show the different types of “action switches” made by our animats. This represents another aspect of our investigation, exploring the idea of action selection being best thought of as a cycle of activities, rather than a series of individual decisions. The thickness of the arrows connecting one state to another represents the transition rates for each run. Here, we distinguish further between those which represent “persistent” (blue arrows) and “opportunistic” (red arrows) patterns of action.

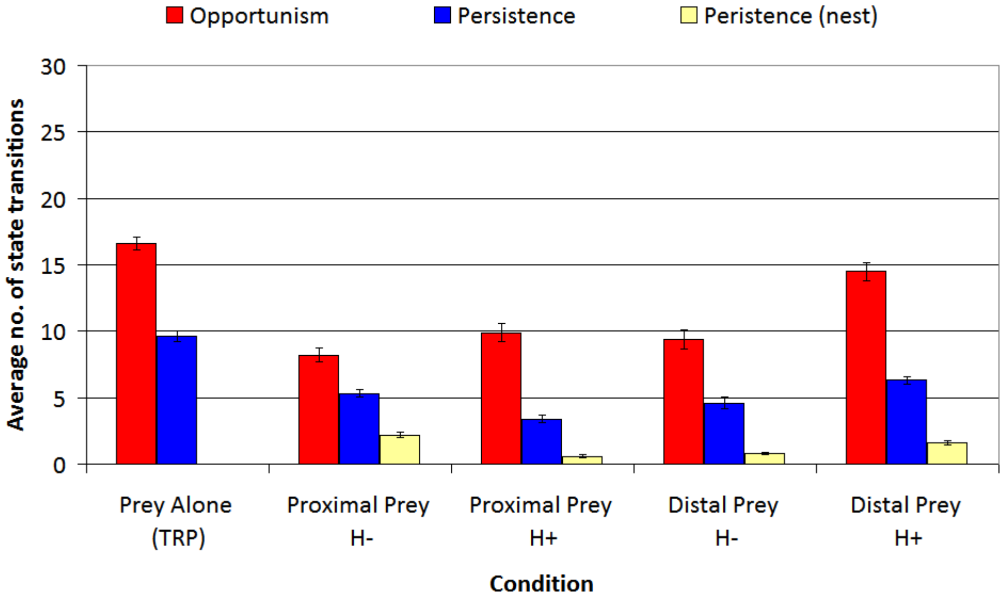

Figure 5.9: Summary of results showing the average number of different types of "state transitions" of our prey across conditions. Error bars show the Standard Error of the Mean (SEM). Exploring the idea of action selection as a cycle of activities, rather than a series of individual decisions, state transitions occur whenever our agent switches from one of the six available actions to another. Having constructed a state transition matrix for each individual run, here we define Opportunism as the number of times our prey consumed one resource while searching for another. Persistence represents the number of times our prey ended up consuming the resource it was searching for and we add a further category to this, Persistence (nest), which is actually a subset of Persistence (focusing on those instances where our prey both searched for the nest and ended up resting in its nest). Here we see that, on average, all of our prey demonstrate more opportunism than persistence. Interestingly though, our modulated distal prey actually tended to display less persistence towards searching for and resting in the nest than our non-modulated prey.

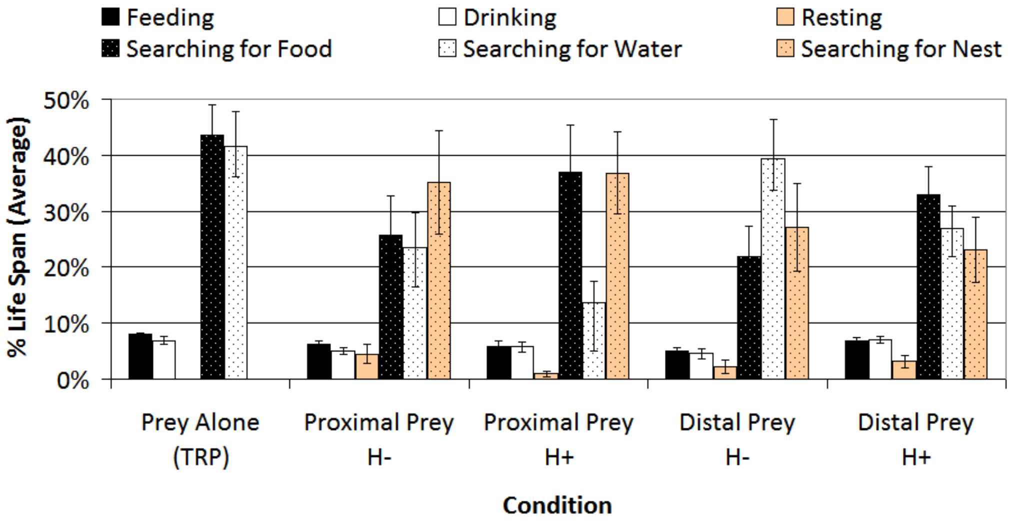

Figure 5.10: Summary of results in terms of “state sinks”. In this case, the average time (as a percentage of overall life span) spent by our different prey animats performing each possible activity/state out of the six available for our architecture to select from. Error Bars show the Standard Error of the Mean (SEM). As the activities that our emotion-based architecture is most likely to influence directly, we can focus our attention here on the proportion of time spent by our prey either resting or searching for the nest. From this we can see that, on average, our two types of proximal prey spent the largest proportion of time searching for their nest. Conversely, our distal modulated prey spent the least amount of time searching for its nest. However, it actually spent more time actively resting than the modulated proximal prey. This also appears to result in an increased proportion of time on other consummatory activities, namely feeding/drinking.

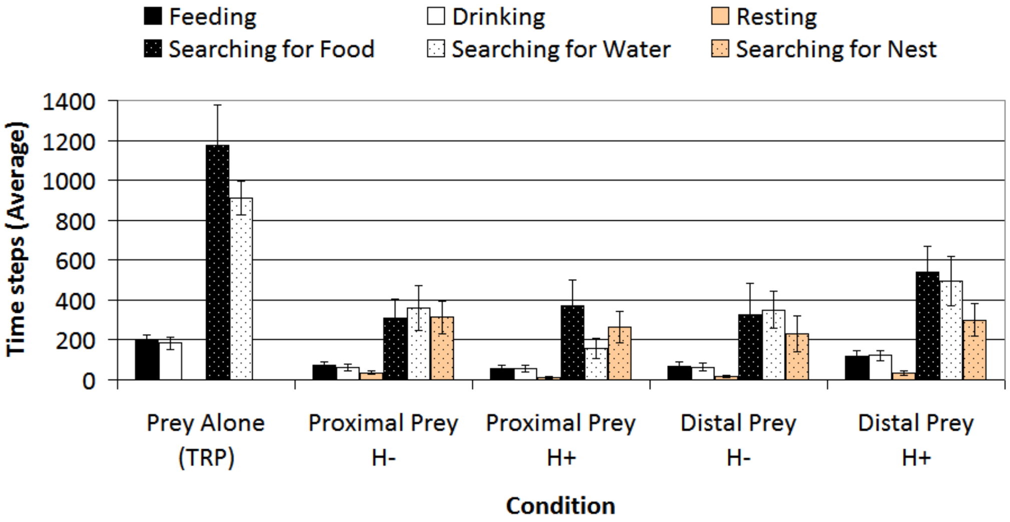

Figure 5.11: Summary of results in terms of "state sinks". In this case, the average time (in terms of action selection loops) spent by our different prey performing each possible activity/state out of the six available for our architecture to select from. Error Bars show the Standard Error of the Mean (SEM). As the activities that our emotion-based architecture is most likely to influence directly, we can focus our attention here on the proportion of time spent by our prey either resting or searching for the nest. From this we can see that, on average, our proximal non-modulated prey and distal modulated prey spent the most time-steps searching for their nest. Conversely, our distal non-modulated prey spent the least amount of time searching for its nest. We can actually see the greatest similarity in terms of absolute time-steps spent resting/searching for nest between our non-modulated proximal prey and modulated distal prey. However, the latter type of prey shows a higher average number of time-steps being spent on consummatory activities, namely feeding/drinking.

predator.

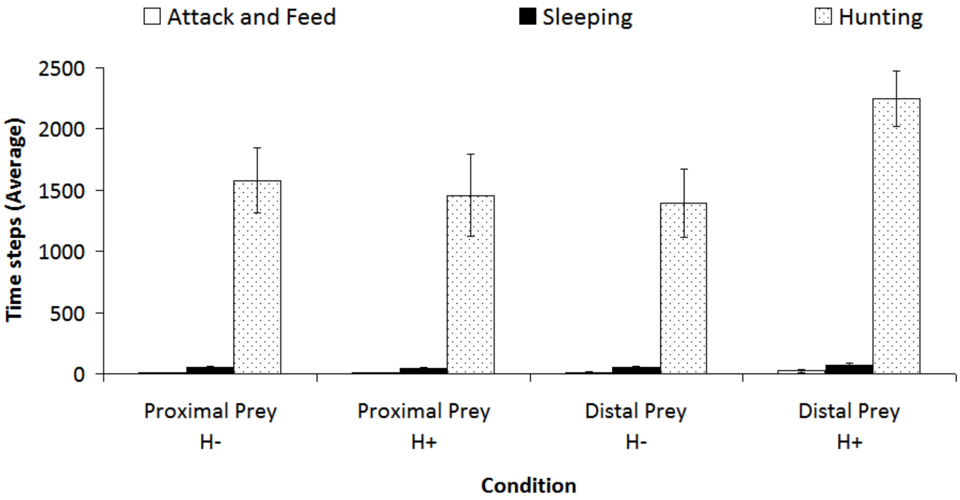

Figure 5.12: Summary of results in terms of "state sinks". In this case, the average time (in terms of action selection loops) spent by our predator animats performing each possible activity/state out of the three available for its architecture to select from. Error Bars show the Standard Error of the Mean (SEM). From this we can see that, on average, our predator spent the most time-steps hunting in the modulated distal prey condition. A greater number of attacks were also performed by the predator in this condition.

### 5.7.3 Brain-Body-Environment Maps

Though we have already looked at our results in the context of some example runs, as a method of visualising the data shown it is interesting to see how our maps show "more of the story". Indeed, from our maps it is also possible to see how the prey's body varied with the external landscape (that is, the environmental dynamics). As Figure 5.13 shows, the prey's movement in the environment was such that it was not necessarily guaranteed to find the resource it was searching for within a predictable time (the prey's movements, particularly obstacle avoidance, having a stochastic element that introduced further noise here). However, this demonstrates another emergent function of our distally-connected hormone-like mechanism; with the "sight" of our predator effectively prompting our prey to keep moving, encouraging both our prey's movement through the environment towards the next resource and away from danger (at times enabling it to both escape attack and keep "topping up" its physiological variables along the way).

To provide more of an overview of this data, Figure 5.14 shows the average number of state transitions or "action switches" of our prey across conditions. Here we can see that, on average, our modulated distal prey performed the greatest number of action switches.

Finally, looking further at the movement of our different animats and time spent on different resources, Figure 5.15 shows the average time spent on different resources/areas in the arena. This shows the differences in terms of the time our prey spent on each resource, looking at the percentage of their "lives" that this made up. On average, our non-modulated proximal prey and modulated distal prey can be seen to have spent more similar proportions of their lives on each type of resource (more closely mimicking the distribution displayed in our control condition, where the prey was acting alone).

## 5.8 Discussion and Conclusion

This chapter builds upon previous studies and constitutes a first step towards a systematic study of how variations in the brain-body interactions of prey and predator can affect the dynamics of prey-predator interactions. In this study, we take in particular what in the literature has been called the Hazardous Three-Resource Problem (H3RP) in which a prey must survive by achieving the right balance between consuming two resources needed to maintain the levels of its two essential internal variables (to satisfy its internal needs or motivations) and protecting itself from a predator.

Varying our basic implementation, our experiments look at whether the addition of our hormone-like component (as the suggested "emotion-like" substrate of our architecture — the details of which are given, along with an outline of our basic implementation, in the previous chapter) proves adaptive for prey performance regardless of its embodiment. That is, in terms of emerging additional functionality, we see how this varies in relation to the sensory connections connecting the prey's brain to its environment. We propose that varying our prey's perceptual abilities by changing its physical "body" will affect the adaptive value and emergent functions of our mechanism simulating interoceptive modulation (which, in turn, also varies our prey's perceptual abilities by acting on the "brain-body" relationship).

Variations in the brain-body interactions were caused here by manipulation of two elements of the prey: its perceptual mechanisms and those mechanisms' connections to features of the environment — in this case the predator. Specifically, we look to see how the performance of the algorithm itself, in the context of the same motivation-based architecture, might be affected by its link to the environment in terms of the senses used for predator detection (triggering hormone-release).

In essence, we have compared the adaptive value of adding a hormone-like mechanism (affecting interoception of an animat prey's internal physiology) to an existing

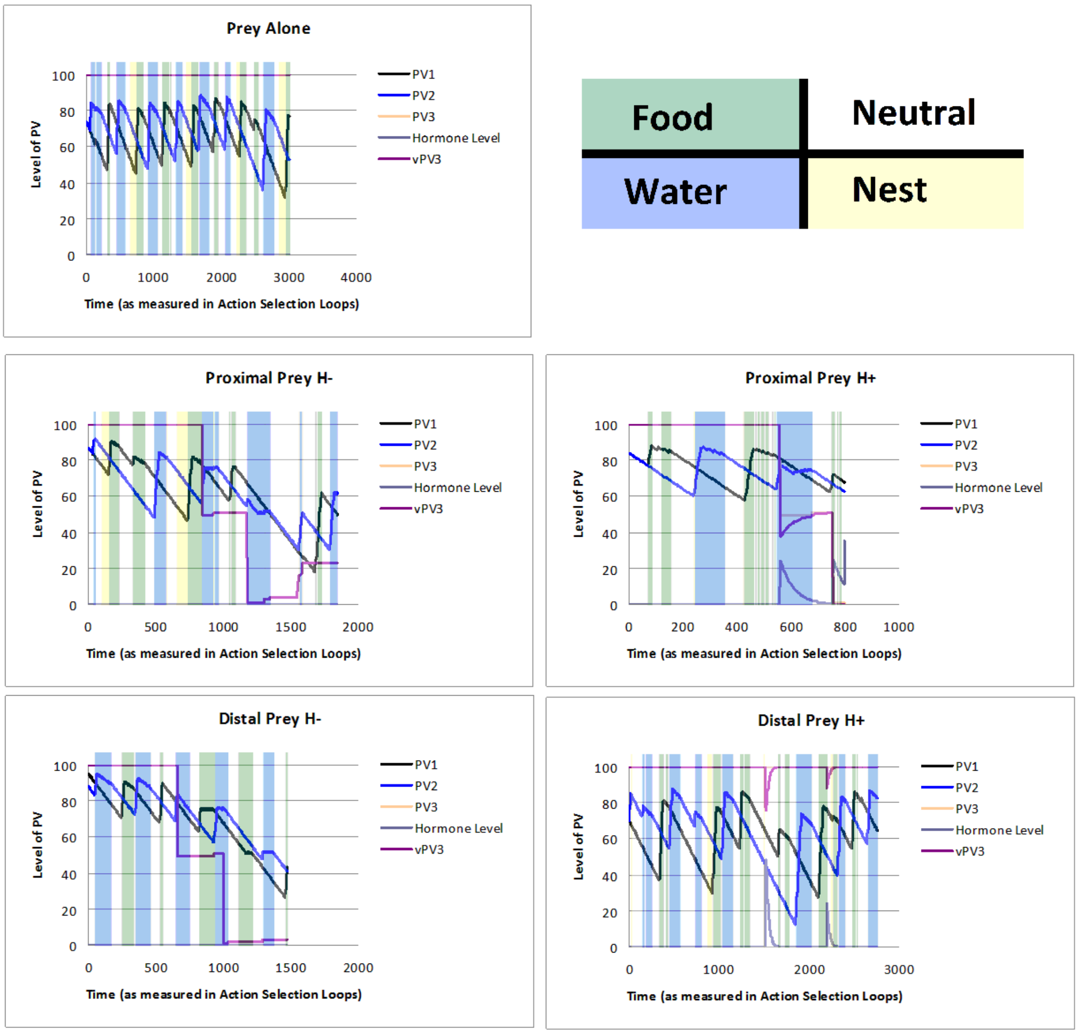

Figure 5.13: Illustration of differences, relating to aspects of prey Brain, Body and Environment over time, as observed across conditions between individual runs and shown by our BBE Maps. The background bars here represent our prey's position within its environment (green bars show where our prey detected "food", blue bars show "water", the absence of any bar indicates the prey detected no resources present in its immediate environment at that time and yellow bars indicate the prey was in its "nest") while each line represents the internal variables of our prey, including each of its physiological variables (black for blood sugar, blue for vascular volume, yellow for integrity) and hormone levels/effects. Additionally, here vPV3 represents the prey's *perceived* level of integrity, which in the modulated conditions differed from the *actual* level of integrity when concentrations of hormone were released.

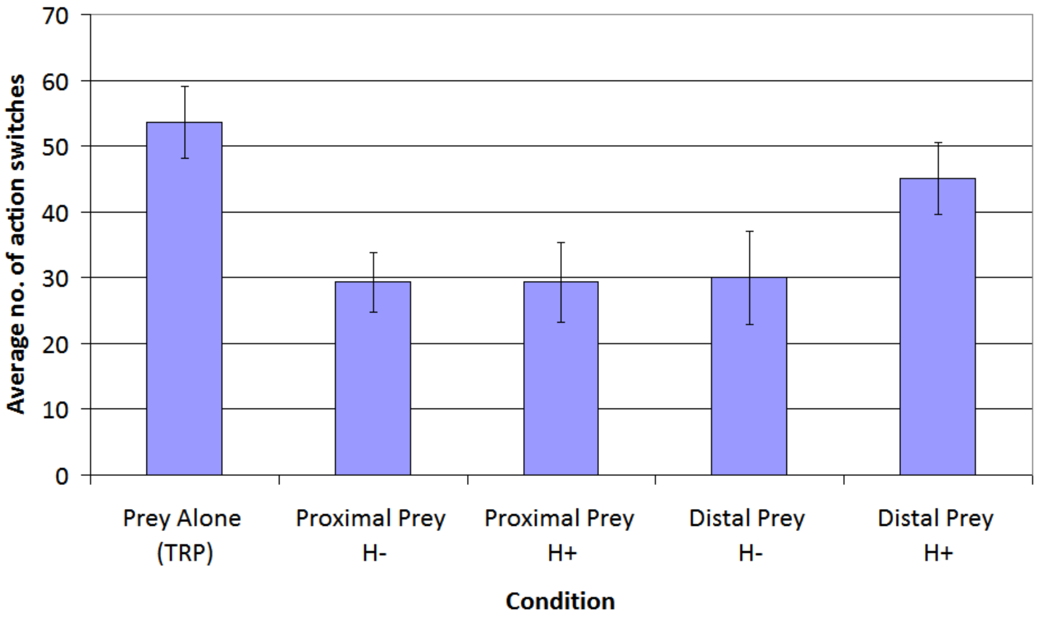

Figure 5.14: Summary of results showing the average number of state transitions or “action switches” of our prey across conditions. Error bars show the Standard Error of the Mean (SEM). Exploring the idea of action selection as a cycle of activities, rather than a series of individual decisions, we use this simple measure to compare our prey. An “action switch” occurs whenever our prey makes the decision to switch from one of the six actions available to it to another. Here we can see that, on average, our modulated distal prey performed the greatest number of action switches.

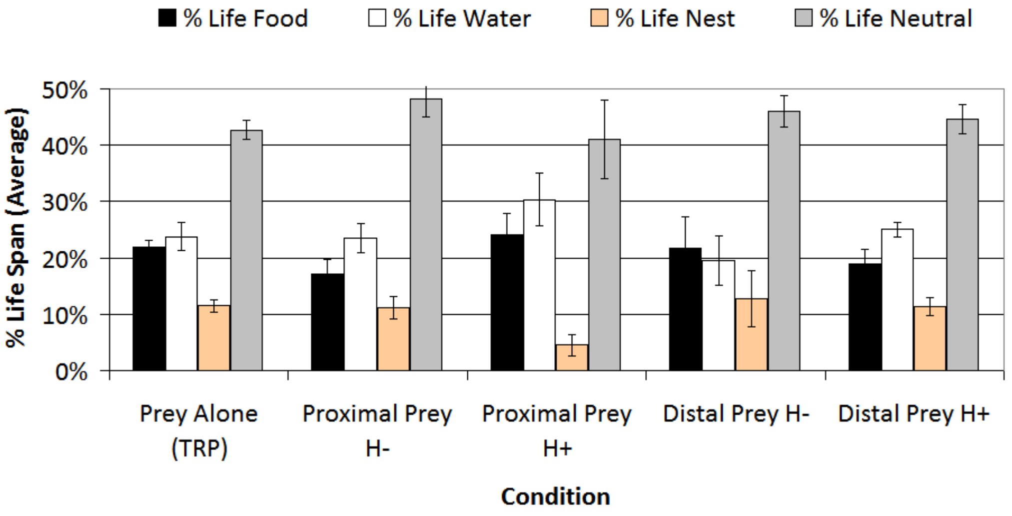

Figure 5.15: Summary of results showing time (as the average proportion of our prey's life span) spent on different resources/areas in its environment/arena by our prey across conditions. Error Bars show the Standard Error of the Mean (SEM). This shows that, on average, our non-modulated proximal prey and modulated distal prey spent more similar proportions of their lives on each type of resource (more closely mimicking the distribution displayed in our control condition, where the prey was acting alone).

motivation-based architecture, when it was linked to proximal detection of a threat to internal physiological variables (the predator) to the adaptive value of adding the same mechanism to the existing architecture when it was linked to both proximal and distal detection. The trends seen in our data (see Figure 5.2), particularly when looking at individual runs (see Figures 5.6, 5.7 and 5.13) would seem to suggest the addition of a hormone-like mechanism was not "adaptive" for the prey when it was linked to only proximal detection (triggering hormone-release) of a predator.

However the addition of our hormone-like mechanism was seemingly adaptive, in terms of a number of our measures and observations, for the prey when it was given further, more distal, abilities for predator detection (triggering hormone-release). From this, we conclude that the way our hormone-like mechanism is connected (or connects brain and body) to the environment, and therefore the stimuli and cues within it, will affect action selection — supported by the differences we have found (thereby further emphasising the importance of body in designing adaptive action selection mechanisms). Our model allows that our agents' environment (composed of the arena and each other) shapes their action selection as much as their action selection shapes the total environment.

In this chapter we use our model to show how the emergence of adaptive responses via interoceptor modulation is dependent on (as it is affected by) perceptual ability. The first set of experiments here were a demonstration of this — which further allowed us to see how we might measure changes in agent performance (thus also providing us with a baseline to compare later results against). Such results as have been explored in this chapter are also interesting in that they allow us to see the "costs" of living longer, using the measures of viability of Overall Comfort and Physiological Balance (Figure 5.2).

This provides initial support for the idea that the emergent functions of interoceptive modulation, in terms of the "flight" response and adaptive behaviour, will be dependent on the brain-body combination. In this case, the most successful combination appeared to be that which allowed the mechanism to act as a temporal predictor of environmental dynamics (in the predator-prey relationship, represented by the prediction or anticipation of predator attack). However, this was not something that we were able to conclude or gain any insights about from either our measures of viability or state transitions. It was only by looking at what happens more closely in each of our agents' "lives" that we could determine this link to, and dependence on the interactions of body with environment (Figure 5.13).

These results thereby provide an initial proof of concept which might be of special interest to researchers in our area including Avila-García [29] and Mendao [84] (providing support for and indicating as it does a role for agent body in emerging emotion-like behaviour that is also adaptive). More generally though, we believe our work should be of interest to the wider research area for the way in which we have extended the existing work on the use of hormone-like mechanisms for action selection.

To summarise, this chapter supports the idea that one of the emergent functions of interoception is to act as a predictor of environmental dynamics (in terms of environmental dynamics such as the predator's behaviour). In the experiments reported here, we have focused on the importance of a specific aspect of embodiment: perceptual distance. Getting this right will be vital for interoception of the prey to be adaptive.

# Chapter 6

# Brain-Body Combinations and Risk Assessment

*“Control your emotion or it will control you.”*
**— Bertrand Russell**

## 6.1 Emergent Functions of Modulating Interoception: Results Chapter II.

Looking to the natural world for examples of intelligent and adaptive behaviour, we see that predators are often described in terms of their highly evolved (as well as adaptive and seemingly intelligent) “hunting” abilities. Whether this is helped along by the evolution of specialist bodies (for example, the evolution of lots of sharp, pointy teeth perfect for making short work of devouring prey) or specialist behaviour produced by specialist brains (brain-body combinations) such as that of ambush predators, these agents, as well as their prey, come in many different shapes and sizes. Considering this further, we get into perhaps more complex and even more specialised (or perhaps just special) realms of predator when we consider species such as lions, which have evolved such that they can successfully coordinate a group attack on their selected prey. Some biological species even assign specific roles to members of their species (ants, for instance) thereby assigning them their “tasks” and, it could be argued, their relative “importance” in life. However, all these species can be linked by their ability to both survive and persist (successfully forage and reproduce) to the present date.

Our research focuses on action selection and our ideas about emotion within the predator-prey relationship. Thus, while the last chapter started our investigation with the initial idea that we would see our emotion-based architecture’s performance (and thus the predator-prey relationship) vary with different perceptual distances and modulation for the prey, we have nevertheless simultaneously, if deliberately for the initial experiments at least, excluded consideration of an equally important and active agent composing this relationship: the predator. Because our aim is to explore further our ideas about both agents, for reasons mentioned in the previous chapters (among these being that it is usually the case that one or the other of these agents is neglected in other predator-prey research) this chapter turns to other experiments that, among other things, look to remedy this. That is, putting the focus back on prey *and* predator. This chapter therefore outlines and looks at the results of experiments designed to explore how both prey and predator perceptual abilities (including aspects of each agent's "body") might or might not interact to affect the adaptive value of our prey's emotion-based architecture.

## 6.2 Emotions in Risk Assessment and Risk Taking

The last chapter started our investigation into the possible role of our hormonal mechanism, as a modulator of perception — and second order controller of a motivation-based architecture — on brain-body-environment interactions within a prey-predator scenario. From these experiments we argue that one of the most important functions of the kind of "emotion" (or potential for "building" adaptive emotion-like mechanisms) that we are studying may be to act as a *temporal predictor* of environmental dynamics (particularly those introduced by the presence of the predator) for prey agents. It also introduces the key concept for us of *predictability* in the predator-prey relationship, especially in the context of the interactive relationship between brains, bodies and environments.

In the rest of this chapter our implementation of the Hazardous Three Resource Problem (H3RP) is extended again to explore further the potential adaptive value of our chosen hormone-like mechanism. In particular, how our prey architecture might react (whether in an adaptive or maladaptive way) in the presence of predators with different physical abilities. Here then, we show how the adaptive value of our prey architecture varies with both perceptual distance and a particular aspect (or dynamics) of its environment; namely, its predator. This is achieved by similarly varying a particular aspect of *predator* embodiment: perceptual distance. In this way, these experiments further aim to investigate the relationships created by different predator-prey combinations. We show that it does not necessarily follow that the greater the perceptual ability, the better interoceptor modulation will work/adaptive responses will emerge. Putting this another way, these experiments show that the power of hormone-like mechanism over action selection must take into account what the perceptual information is indicating about future environmental dynamics. We argue this may point us towards another emergent function of or role for "emotion" in our H3RP. That is, as a *risk assessor*.

## 6.3 Brain-Body-Environment Combinations for Prey and Predator

In our experiments this time we designed variations of the H3RP by creating and combining different prey *and* predator agents, varying *both* agents' *perceptual distance* (proximal/distal). In the process, we therefore created four main different predator-prey combinations of these agents for testing: proximal prey versus proximal predator, proximal prey versus distal predator, distal prey versus proximal predator and distal prey versus distal predator (see Figure 6.1). In order to investigate the impact of our hormone-like mechanism further, however, we also varied other parameters in these experiments that focused on the prey alone, including both hormone *strength* and *decay*. (Though the latter variation is not one that varies within this chapter's experiments, rather it is one that varies between the experiments of this chapter and those of the last chapter — something we will come back to later).

## 6.4 Prey

### 6.4.1 Interoceptive Modulation (Hormone-Release of 0, 25, 50 and 100)

Building on the results of the last set of experiments, one of the main purposes here was to go further than previously, to take a more in-depth look at (and get a broader snapshot of) how aspects of embodiment and modulation might interact to affect the performance of our system. In our experiments, we therefore created further types of prey by varying the parameter representing hormone release rate, which we tested at four levels: 0, 25, 50 and 100. A level of 0 is effectively a non-modulated architecture, while a level of 25 means that only 25 percent of the maximum concentration of hormone is released each time the predator is detected.

#### 6.4.1.1 Hormonal Decay Rate

For this type of prey we also tried modifying our hormone-like mechanism in another way. For, in these experiments, we adapted our existing hormone-like mechanism using a different type of decay rate. Rather than the exponential decay rate previously used, we implemented one that was more sigmoidal in its effect, allowing the hormone to influence the prey's action selection/system for longer. This came from the observations of early experiments that the exponential decay rate used previously led to a very short "emotional memory" for our prey — which meant that our early prey was easily "distracted" by a food/water resource when "fleeing" the predator after an attack. While usually enough to get the prey to move off a resource after attacked and

|  | | 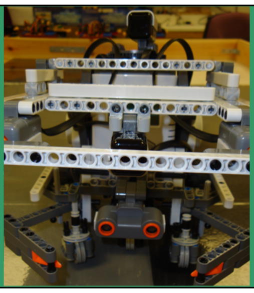 | |
|---|---|---|---|
| *Variations in Prey "Body" [Perceptual Distance and Hormone Strength]** | | *Variations in Predator “Body” [Perceptual Distance]** | |
| **Perceptual Distance** | **Hormone Strength** | **Proximal Predator** | **Distal Predator** |
| **Proximal Prey** | **H-** | Proximal Prey<br>vs<br>Proximal Pred | Proximal Prey<br>vs<br>Distal Pred |
| | **H25** | | |
| | **H50** | | |
| | **H100** | | |
| **Distal Prey** | **H-** | Distal Prey<br>vs<br>Proximal Pred | Distal Prey<br>vs<br>Distal Pred |
| | **H25** | | |
| | **H50** | | |
| | **H100** | | |

Figure 6.1: Summary of experimental conditions: In these experiments a factorial 2x2x4 design (resulting in sixteen experimental conditions) was used to investigate further how the adaptive value of our emotion-based architecture varies when we both vary the strength of the hormone-like mechanism and the perceptual abilities of our prey *and* predator. This design enables us to identify the effects of a) the “emotion-like” substrate represented by our hormone-like mechanism, using different hormone release rates to represent further variation in our prey’s “brain” b) our prey’s perceptual distance; proximal/distal detection representing a variation in prey “body” and c) our predator’s perceptual distance; proximal/distal detection representing both a variation in our predator’s “body” and prey’s “environment”. *Variations in the body of one agent potentially leading to a variation in the other’s environment and vice versa.

search for its nest, we noticed its exponential nature meant the “fear” of the prey acting to increase the motivation to search for its nest and rest was soon overcome by any consequent encounters with a food or water resource. That is to say, the incentive cue of an existing resource won out over the adjusted interoception of the physiological variable integrity (which was used to calculate the motivational intensity for the need to rest).

Thus, based on our intuitions from previous experiments, we wanted to see whether changing a specific aspect of our hormone-like mechanism would prove more adaptive for our prey. In terms of the original mechanism, as outlined previously, the concentration of hormone suffers two opposite forces over time. That is to say, it increases with the release of hormone by the gland, and dissipates or decays over time. While previously, this was at a constant rate, here we modified the decay rate of our mechanism, using a sigmoidal decay rate which allows the perception of its predator to affect/bias our prey’s action selection for longer (which could be thought of as akin to increasing the “persistence” of our prey’s fear in biasing the underlying motivation-based architecture). We wanted to see what would happen when the hormone-like mechanism was that bit more persistent over time (and whether this would be adaptive in terms of prey survival).

### 6.4.2 Perceptual Distance (Proximal/Distal)

Again, as an aspect of embodiment we chose to continue to focus on, we aimed to vary our prey’s perceptual distance by connecting the basic motivation-based architecture (also comparing different hormone-release rates) to two different kinds of ability (identified in the last chapter “proximal prey” and “distal prey” or “blind” and “seeing” prey respectively).

#### 6.4.2.1 Proximal Prey

Our first type of prey’s perceptual distance was severely limited. This was achieved by making the only sensory information available to our prey from its environment come from touch (via its two touch sensors) and internal perception of the variable integrity (a sudden decrease signalling an attack by the predator). As mentioned previously, this scenario can be thought of as limiting the prey to a system akin to the somatosensory system available to many biological organisms. However, this system is also further limited in that the prey cannot even actively (by touch alone) discriminate between the stimuli of a wall or that of the (more dangerous) predator. For this prey, the only way it has for “perceiving” the predator (and thus the only signal affecting our mechanism/triggering hormone release) is via internal changes following an attack.

#### 6.4.2.2 Distal Prey

As before, our second type of prey's perceptual distance, especially with regards to detection of the predator and the proximal prey's ability, was increased considerably (to approximately 45cm away from its body). This was achieved by the addition of an ultrasonic sensor. With regards to our hormone-like mechanism, information from this sensor would thus give the prey two signals available for potentially detecting predator presence/triggering hormone release. As with the proximal prey, we created further types of this prey by also varying the parameter representing the level of hormone released, tested at the four levels specified.

## 6.5 Predator

### 6.5.1 Perceptual Distance (Proximal/Distal)

This time, we also aimed to vary our predator's perceptual distance by connecting the basic motivation-based architecture to two different kinds of ability (which, similarly to the two types of prey agent, can be identified as "proximal predator" and "distal predator" or "blind" and "seeing" predator respectively).

#### 6.5.1.1 Proximal Predator

The first type of predator was the same as previously studied and given an extremely limited perceptual distance. Thus, its hunting abilities were extremely crude, to say the least. The only way this predator can find its prey is to blindly search its environment and wait for the prey to somehow come into contact with its touch sensors. When it does so, this triggers an "attack" on the predator's part, which sends a signal to the prey via Bluetooth to decrease its level of integrity.

#### 6.5.1.2 Distal Predator

However, the abilities of the second (new) type of predator, like the second type of prey, were increased considerably. This was achieved by the addition of an infra-red (IR) seeker to the predator — and infra-red ball mounted on the prey. This allowed us to program a more sophisticated "hunting strategy" by getting the predator to follow any IR light detected via its infra-red seeker/otherwise using the same "blind search" as the proximal predator until the prey is encountered. This predator can therefore be considered much more *dangerous*, but at the same time more *predictable*. It is this predictability that we think will prove an important factor in determining the emergent properties of our prey's emotion-based architecture.

## 6.6 Arena

As before, the layout of the arena for these experiments was kept constant, consisting of a nest area, food, water and neutral space.

## 6.7 Results

To recap, using our two main types of prey and two main types of predator, we thereby created four main predator-prey combinations for study. Further conditions were created for comparison by varying the level of "hormone" released by our algorithm (at four levels — 0, 25, 50 and 100). With regards to the procedure for each prey run, we replicated that described in the last chapter. Similarly, the same data was collected from both predator and prey after each run with regards to both sensory and internal (physiological) data. Combining our variables at each of their levels this created sixteen experimental conditions for our initial comparisons (summarised in Figure 6.1 and creating a balanced, 2x2x4 factorial design). We collected data from a total of fifteen runs per condition (with a total of seventeen conditions — as this also includes a control conditions, where we also tested the prey alone — this generated data from a total of 255 runs). This experimental design allowed us to conduct a much more in-depth study than previously, looking at the effect of different hormone-release rates on the emergent properties of our developed system.

### 6.7.1 Measures of Viability

From the data collected, we again first calculated and looked at the results in terms of our chosen measures of viability (previously outlined along with our basic implementation). Figures 6.2-6.5 shows how each type of prey performed as measured by Life Span, Overall Comfort and Physiological Balance. Again, for purposes of comparison, we have also included the results obtained from testing our prey alone, without a predator agent; making this essentially the results obtained using the prey to perform a two-resource problem. The purpose of this was to again check that our predators were actually having an impact on the prey in terms of representing a hazardous resource.

#### 6.7.1.1 Life Span, Overall Comfort and Physiological Balance

In terms of Life span (how many time-steps out of the maximum that the prey lasted for/remained viable) Figure 6.2 shows the average Life span, using the same method outlined in the last chapter, calculated per condition (error bars show the standard error of the mean). Looking for trends in terms of Life Span (specifically, as calculated from the average minutes the prey "lived" for each condition) for the results of those conditions testing the proximal predator we can see that, on average, our distal predator appeared to consistently shorten the average Life Span of our prey, regardless of condition. This is perhaps none too surprising in itself. For a predator that can "see" its prey will surely be expected to catch it more often! However, it is interesting to note that the modulated prey in the distal predator conditions all show a tendency to last longer than the prey without the hormone-like mechanism. Might this show a greater use for our mechanism in circumstances where both the danger and predictability of dynamics in the environment (that is to say, the predator) are increased?

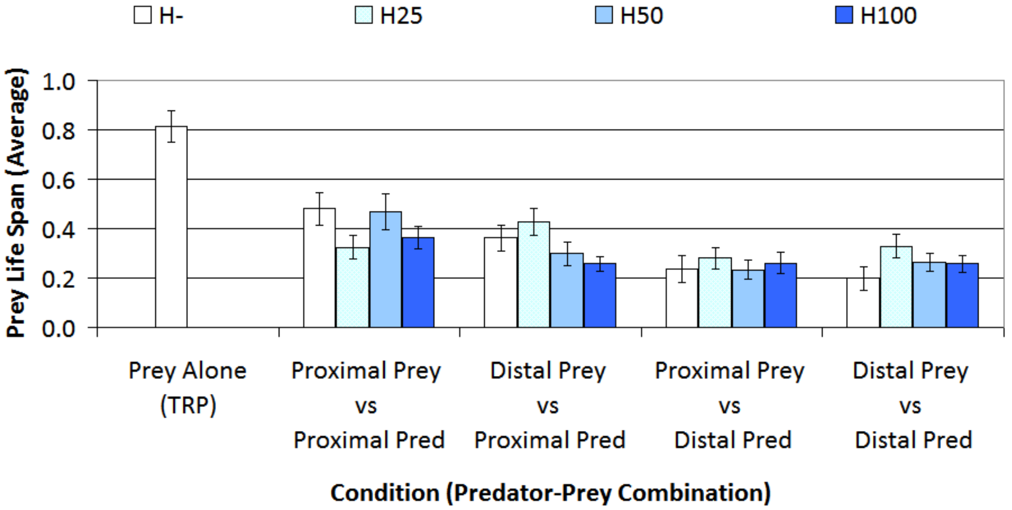

Figure 6.2: Summary of results showing average Life Span (LS) of our prey across experimental conditions. Error Bars show the Standard Error of the Mean (SEM). Looking at the trends, on average, the presence of our distal predator perhaps not surprisingly tended to shorten the average LS of our prey. More interestingly to us though, we see our modulated prey in the distal predator conditions displaying a tendency to last longer than the prey without the hormone-like mechanism. The results of our inferential tests suggest a main effect for predator ability, which we consider to lend support to our idea that our mechanism generates more adaptive behaviour when both the danger and predictability of the environment (namely the predator) is increased.

Looking at the trends shown, at first glance it seems that the introduction of our distal predator caused a general decrease in average Life Span. Though we continue here to focus on Life Span as calculated from the number of time steps our prey survived for, for these experiments — in contrast to those of the previous chapter — when looking at the real-time survival of our prey, we now began to see some differences between our calculations of Life Span, depending upon whether we are looking at real time (msecs) or the number of time steps survived. This is most likely due to some of our prey agents having spent more time performing actions that take longer (namely, obstacle-avoidance). However, it is an interesting difference to note. Thus, to be consistent with our previous results, for this data set we have again focused our inferential tests purely on calculations made using the recorded number of time-steps the prey "survived". However, this time we also performed a further test on Life Span as calculated by the total number of minutes our prey survived. Again after arcsine transformation, we performed a multi-way (2x2x4) ANOVA for Life Span across our sixteen experimental conditions. This was to see if there were any main effects and interactions between prey perceptual distance (proximal/distal) predator perceptual distance (proximal/distal) and hormone strength (0, 25, 50, 100).

In terms of Life Span calculated using time-steps, the results of this test showed no main effect for prey ability (p=0.211). However, a significant main effect was found for predator ability (p=0.000). While significant interactions between our three independent variables (prey perceptual distance, predator perceptual distance and hormone strength) were not found, we nevertheless found the result for interactions of predator perceptual distance and hormone strength to be interesting (p=0.071). Indeed, though the inferential tests performed as a result of this thesis have continued to be judged by the conventional alpha level of 0.05, in exploring our results, a further question arising from our research, is whether, in "noisy" robotic systems such as ours, the conventional alpha would lead us to reject make more or less Type II errors. Though currently beyond the scope of this thesis, it is something which we would be keen to see addressed in future. Doing the same for our second calculation of Life Span in minutes, we found no main effect for prey ability (p=0.319). However, we did again find a significant main effect for predator perceptual distance (p=0.00) as well a main effect for hormone-strength (p=0.007). In turn, the results further suggested a significant interaction between prey ability and hormone strength (p=0.038).

In terms of Overall Comfort and Physiological Balance, Figures 6.3 to 6.5 do not show any obvious decreases between their averages with regards to the introduction of the predator. However, this could be due to shorter average Life Spans in some conditions (leading to a prey "killed off" in its prime displaying similar levels to a prey that has successfully survived and "thrived" in its environment). Performing a multi-way ANOVA for each of these measures, after arcsine transformation no main effects or interactions were found for prey or predator perceptual distance or hormone-strength in terms of Overall Comfort. However, main effects were shown for all three independent variables in terms of Physiological Balance (main effects for prey perceptual distance (p=0.001) predator perceptual distance (p=0.002) and hormone-strength (p=0.000) but not significant interactions between these conditions). Interestingly, when looking at Overall Comfort in terms of blood sugar and vascular volume only, significant effects were found to be the same as for Physiological Balance (main effects for prey perceptual distance (p=0.012) predator perceptual distance (p=0.034) and hormone-strength (p=0.015) but no interactions between groups).

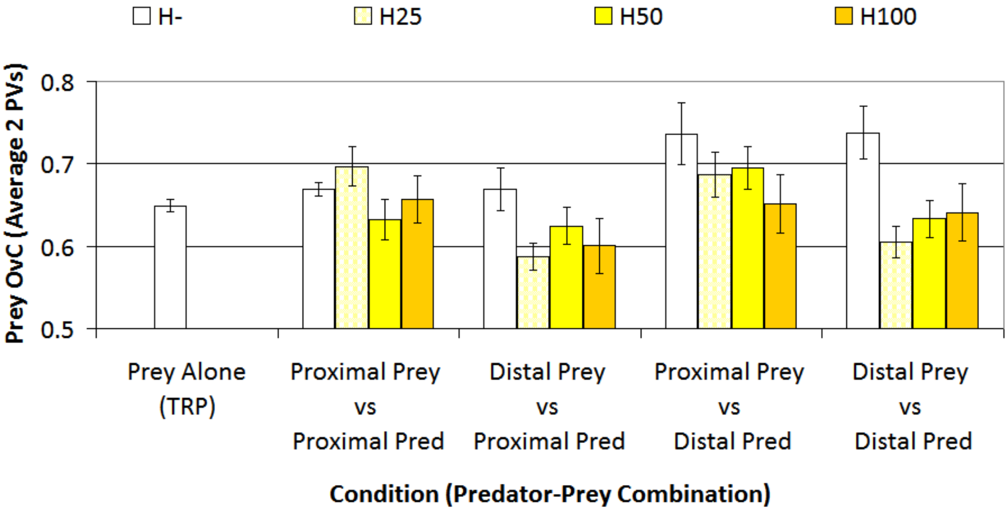

Figure 6.3: Summary of results showing average Overall Comfort (OvC) of our prey across conditions. Error Bars show the Standard Error of the Mean (SEM). In this case, calculations of OvC exclude the third variable of integrity so that we can first consider the effect of adding this variable on our architecture's performance of the original two resource problem. Unlike with Life Span, in comparison with our control condition (Prey Alone) here we do not observe a distinct decrease in OvC with the introduction of our predator. This may be due to shorter average Life Spans — a prey "killed off" in its prime displaying similar levels to a prey that has successfully survived and "thrived" in its environment. For each predator-prey combination though, comparing the modulated to the non-modulated prey, in all but one condition the modulated prey tended to display a decrease in OvC in comparison to its non-modulated counterpart.

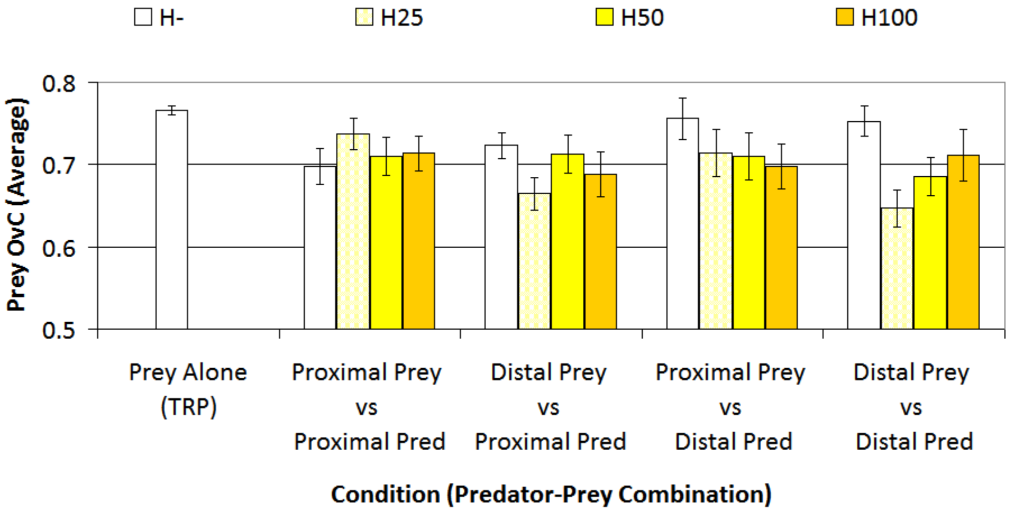

Figure 6.4: Summary of results showing average Overall Comfort (OvC) of our prey across experimental conditions. Error Bars show the Standard Error of the Mean (SEM). In this case, calculations of OvC include the third variable of integrity. Unlike with Figure 6.3, in comparison with our control condition (Prey Alone) here we do observe a decrease in OvC with the introduction of our predator. For each predator-prey combination, comparing the modulated to the non-modulated prey, in all but the proximal prey versus proximal predator condition (which showed the opposite trend) the modulated prey tended to display a decrease in OvC in comparison to its non-modulated counterpart.

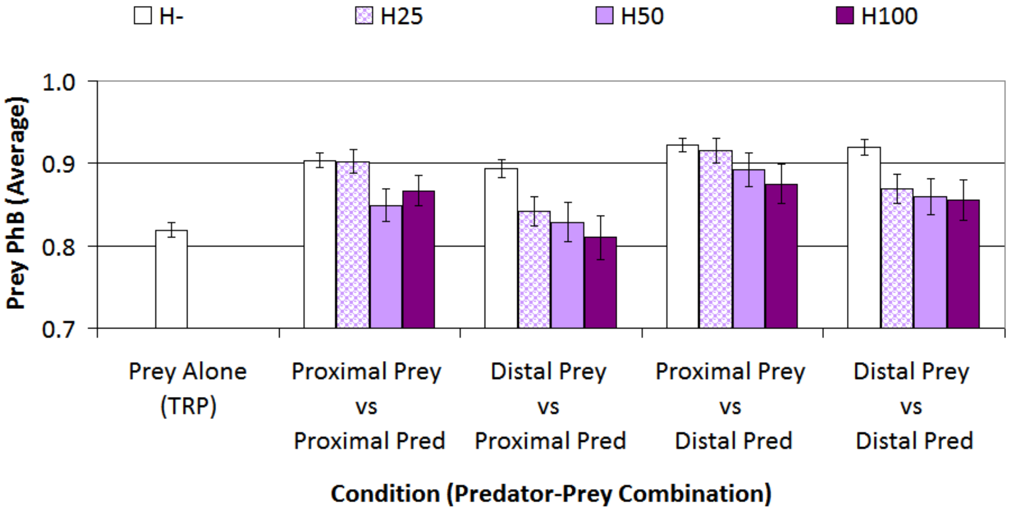

Figure 6.5: Summary of results showing average Physiological Balance (PhB) of our prey across experimental conditions. Error Bars show the Standard Error of the Mean (SEM). In this case, calculations of PhB include the third variable of integrity. In comparison with our control condition (Prey Alone) here we actually observe a tendency to increase average PhB with the introduction of our predator. However, this can also be explained by the fact that in our control condition the variable integrity stays constant, while the other two variables fluctuate. That is to say integrity will never decrease because the predator is absent, thereby skewing (artificially lowering) calculations of PhB. For each predator-prey combination, comparing the modulated to the non-modulated prey, the modulated prey tended to display a decrease in PhB in comparison to its non-modulated counterpart. It also seems that, the greater the level of hormone released following detection of the predator/an attack, the greater the decrease in PhB for all conditions.

#### 6.7.1.2 Risk of Death Management

In terms of Management of Death Risk (as defined in the previous chapter) the individual regions for this measure can be used as an indication of how well viability is preserved. Again, using this measure to explore our data more visually provides us with an interesting perspective on the performance of our different types of prey. Figures 6.6 to 6.10 show our results, in terms of Risk of Death (RoD) showing a plot for each experimental condition as well as our control (where prey was acting alone). Each bar represents the average percentage of time each type of prey was in each "zone". A RoD of 1 means the prey is almost, if not the most "healthy" it can be, whereas a RoD of 10 means the prey is "near death". While the differences here are not as striking those in our initial experiments, the differences in the shapes our results create are still of interest.

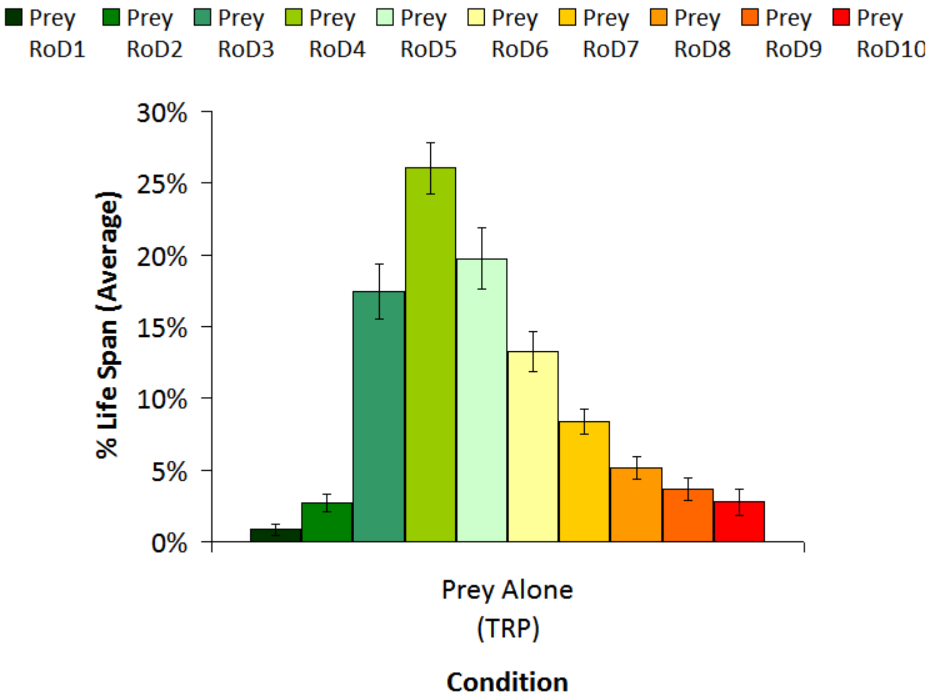

Figure 6.6: Summary of results showing Risk of Death (RoD) Management in our control condition (Prey Alone). Each bar represents the average percentage of time our prey was in each RoD "zone". Error Bars show the Standard Error of the Mean (SEM). A RoD of 1 (green) means that the prey is almost if not exactly the most "healthy" it can be, whereas a RoD of 10 (red) means the prey is "near death". As seen, in this condition, on average our prey tended to spend most of its life in a "healthy" state.

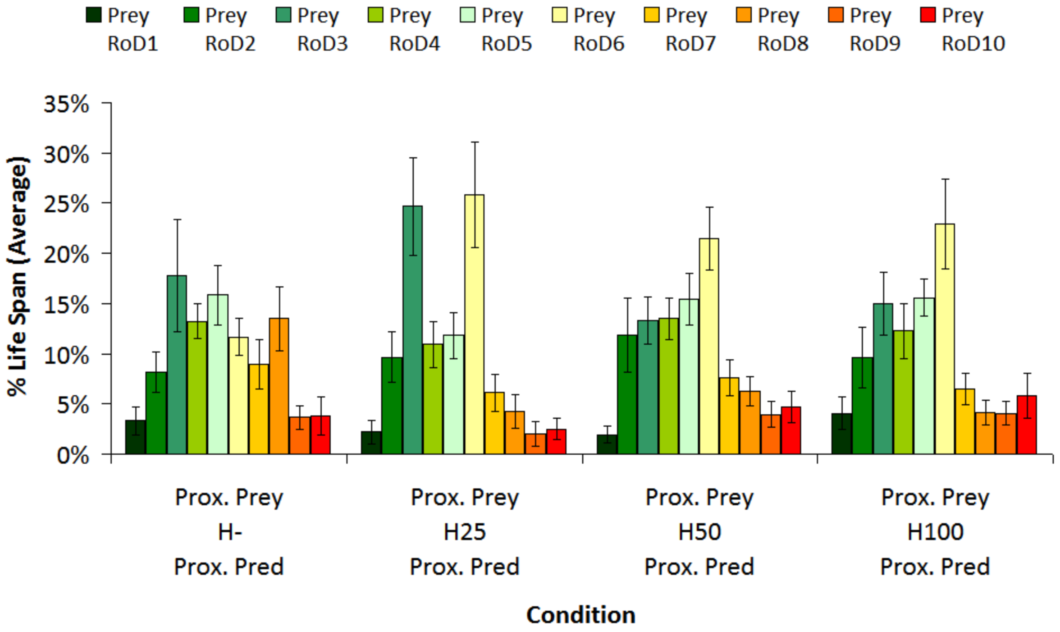

Figure 6.7: Summary of results showing Risk of Death (RoD) Management across conditions where proximal prey versus proximal predator. Each bar represents the average percentage of time our prey was in each RoD "zone". Error Bars show the Standard Error of the Mean (SEM). A RoD of 1 (green) means that the prey is almost if not exactly the most "healthy" it can be, whereas a RoD of 10 (red) means the prey is "near death". Here, we see our proximal prey with a hormone-release of 25 seemed to spend the least amount of time close to death. However, considering Life Span, this was also the condition characterised by a life met with a quick end. On average, the longest lived prey was the non-modulated one. This appears to have also resulted in a more even distribution across zones, which is more similar to that of our control condition.

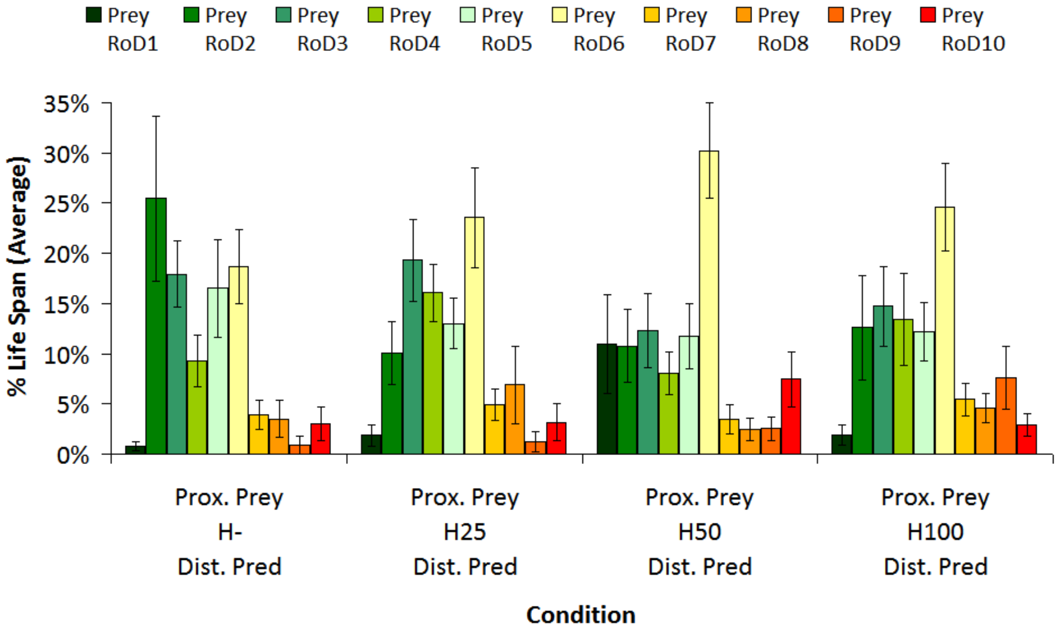

Figure 6.8: Summary of results showing Risk of Death (RoD) Management across conditions where proximal prey versus distal predator. Each bar represents the average percentage of time our prey was in each RoD "zone". Error Bars show the Standard Error of the Mean (SEM). A RoD of 1 (green) means that the prey is almost if not exactly the most "healthy" it can be, whereas a RoD of 10 (red) means the prey is "near death". Here, we see our non-modulated prey seemed to spend the least amount of time close to death. However, considering Life Span, this was again the condition characterised by a life met with a quick end. On average, the longest lived prey here was the one with a hormone-release of 25. This appears to have again resulted in a more even distribution across zones for this condition, which is more similar to that of our control condition.

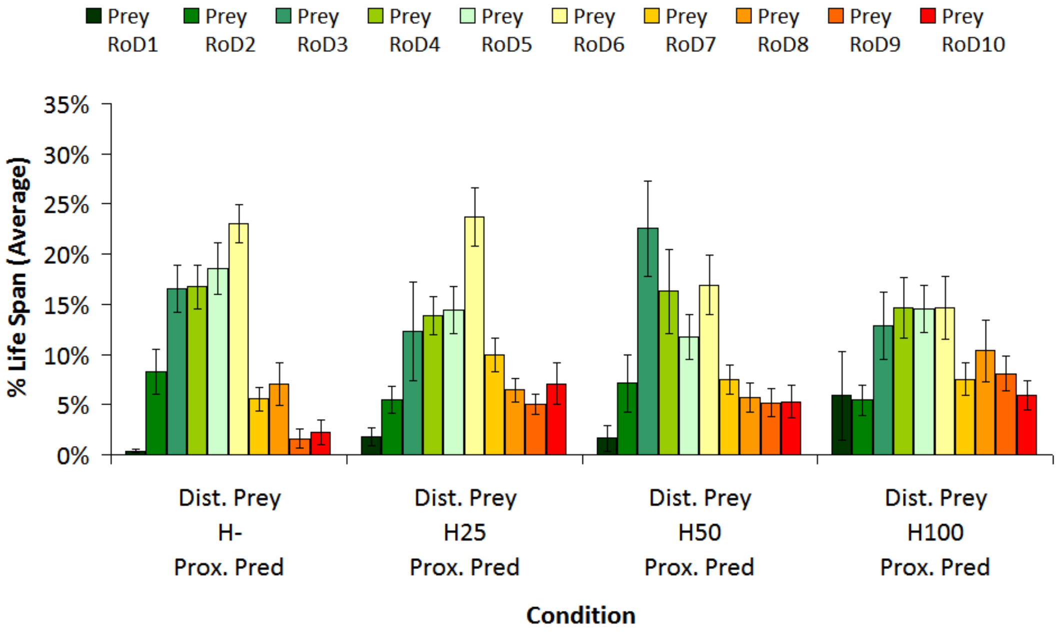

Figure 6.9: Summary of results showing Risk of Death (RoD) Management across conditions where distal prey versus proximal predator. Each bar represents the average percentage of time our prey was in each RoD "zone". Error Bars show the Standard Error of the Mean (SEM). A RoD of 1 (green) means that the prey is almost if not exactly the most "healthy" it can be, whereas a RoD of 10 (red) means the prey is "near death". Here, we see our non-modulated prey seemed to spend the least amount of time close to death. On average, the longest lived prey here was the one with a hormone-release of 25. This appears to have again resulted in a more even distribution across zones for this condition, which is more similar to that of our control condition.

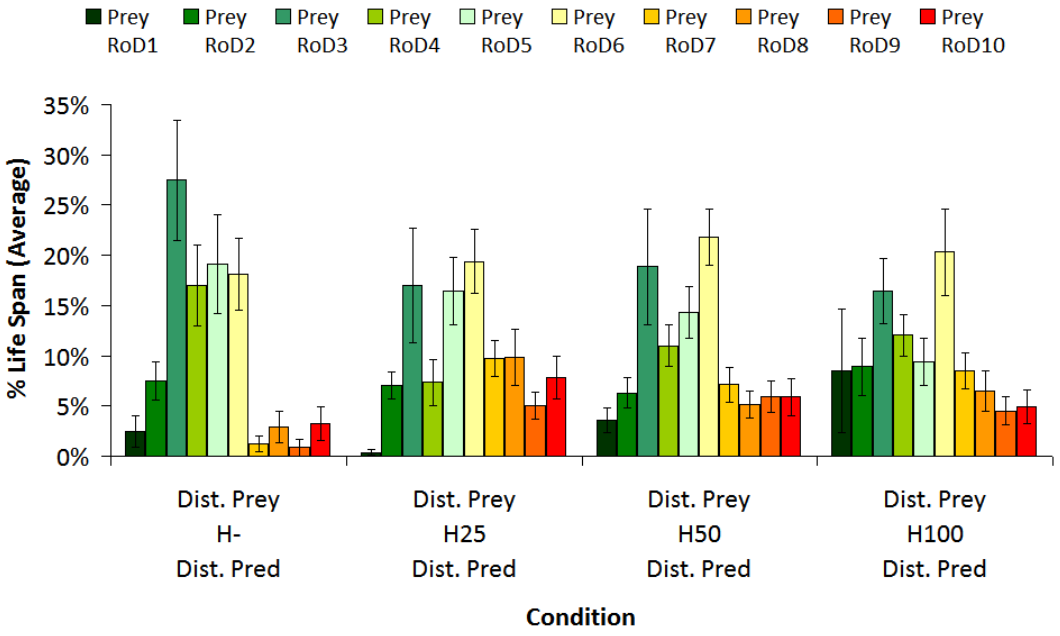

Figure 6.10: Summary of results showing Risk of Death (RoD) Management across conditions where distal prey versus distal predator. Each bar represents the average percentage of time our prey was in each RoD “zone”. Error Bars show the Standard Error of the Mean (SEM). A RoD of 1 (green) means that the prey is almost if not exactly the most “healthy” it can be, whereas a RoD of 10 (red) means the prey is “near death”. Here, we see our non-modulated prey seemed to spend the least amount of time close to death. However, considering Life Span, this was also the condition characterised by a life met with a quick end. On average, the longest lived prey here was the one with a hormone-release of 25.

For these graphs, as opposed to the prey alone condition, we see some runs where the prey seemed to spend all of its life in the same zone. We found that these are most likely to be those runs characterised by a life met with a quick end. That is, death inflicted by the predator within a short period. Indeed, the distal predator conditions seem to be characterised by more of these peaks for the proximal prey than for the distal. Interestingly, it is the distal prey, distal predator conditions where we again see signs of the "calming" or stabilising effect of our hormone-like mechanism (looking at the graph for this condition, from left to right, shows the effect of increasing hormone strength).

## 6.7.2 Activity Cycles and Action Patterns

### 6.7.2.1 States and State Transitions

Looking next at how this performance translates to a pattern of action, Figure 6.11 shows some averages for different types of state transition across conditions. Here we see that, on average, all of our prey continue to demonstrate more opportunism than persistence. Considering the impact had on prey state transitions and state sinks, investigating this further, in terms of an activity profile, Figures 6.12 to 6.15 show the average time spent on each action per condition, as a percentage of overall life span. Together with Life Span, these graphs better show both differences in terms of the absolute number of time-steps that each type of prey spent on each activity, and the percentage of their "lives" that this made up. For our proximal prey against our proximal predator we actually still see quite similar patterns in terms of state transitions as for our control condition. However, with the H100 release level, the amount of time spent in the nest dramatically increases.

For each condition, we can see the effect of increasing the strength or "power" of our hormone-like mechanism. Firstly, across all conditions, we can see (looking at each plot from left to right) that increasing hormone release not unexpectedly generally increased the proportion of its lifetime our prey spent on activities associated with recovering "integrity". Comparing the results for the proximal prey versus proximal predator conditions with those for the distal prey versus proximal predator, we then see how introducing distal perception increases these proportions yet again. With our proximal prey versus proximal predator condition though, when our hormone strength is at its highest, we see our prey spending most of its time either in the nest or searching for it (in the context of our state transition graphs seeming then to make only "opportunistic" stops for food/water if it should so happen to be encountered).

Turning to our distal predator conditions, we see the same general trends. However, taken together with our other observations, in this case it is most likely that the more advanced hunting abilities of our predator is responsible for increasing the proportions of time spent performing activities associated with recovering integrity (that is, when

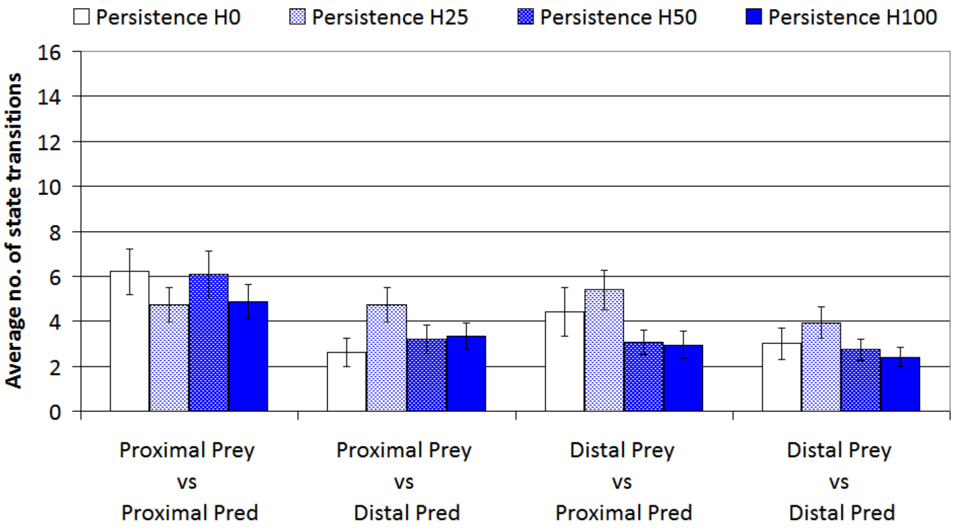

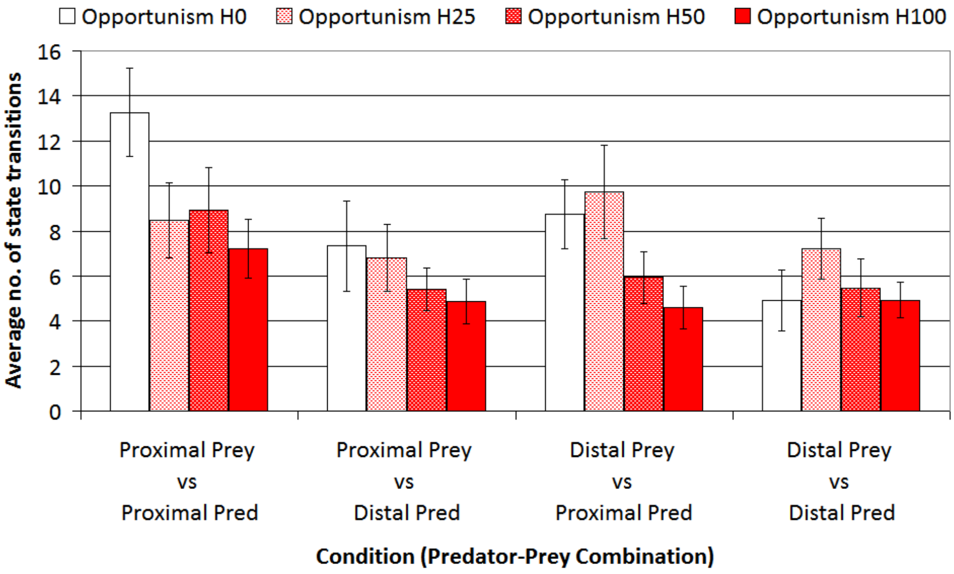

Figure 6.11: Summary of results showing the average number of different types of "state transitions" of our prey across conditions. Error bars show the Standard Error of the Mean (SEM). Exploring the idea of action selection as a cycle of activities, rather than a series of individual decisions, state transitions occur whenever our agent switches from one of the six available actions to another. Having constructed a state transition matrix for each individual run, here we define Opportunism as the number of times our prey consumed one resource while searching for another. Persistence represents the number of times our prey ended up consuming the resource it was searching for. Here we see that, on average, all of our prey continue to demonstrate more opportunism than persistence.

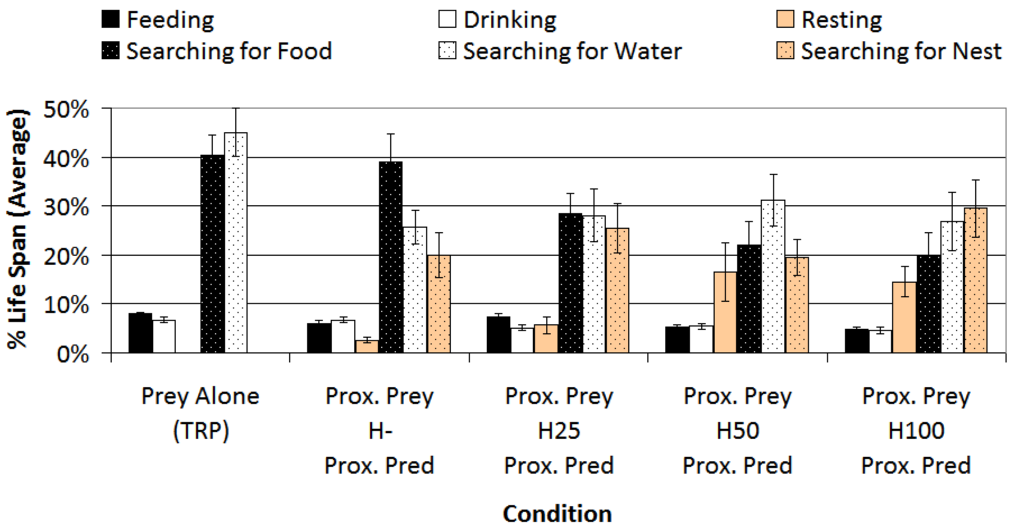

Figure 6.12: Summary of results in terms of “state sinks” across conditions where proximal prey versus proximal predator. In this case, the average time (as a percentage of overall life span) spent by our different prey animats performing each possible activity/state out of the six available for our architecture to select from. Error Bars show the Standard Error of the Mean (SEM). As the activities that our emotion-based architecture is most likely to influence directly, we can focus our attention here on the proportion of time spent by our prey either resting or searching for the nest. From this we can see that the prey which tended to spend most time doing the latter was, somewhat unsurprisingly, that with a hormone release of 100. Conversely, our non-modulated prey and prey with a hormone release of 50 spent the least amount of time searching for the nest. However, the prey with a hormone release of 50 actually spent the most time actively resting (albeit, looking at the error bars, not quite as consistently across life times as we saw for our prey with a hormone release of 100).

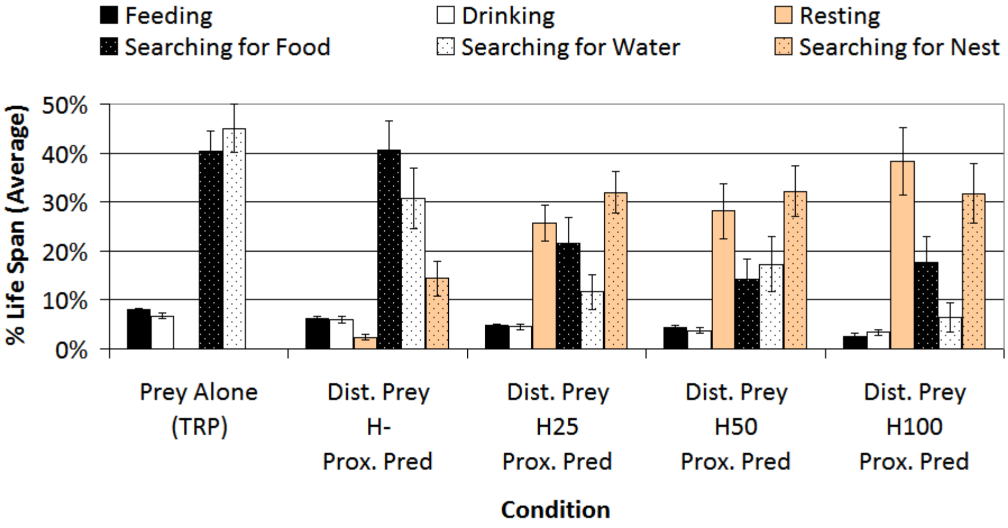

Figure 6.13: Summary of results in terms of "state sinks" across conditions where distal prey versus proximal predator. In this case, the average time (as a percentage of overall life span) spent by our different prey performing each possible activity/state out of the six available for our architecture to select from. Error Bars show the Standard Error of the Mean (SEM). As the activities that our emotion-based architecture is most likely to influence directly, we can focus our attention here on the proportion of time spent by our prey either resting or searching for the nest. From this we can see a definite impact on these activities upon having increased the perceptual distance of our prey. That is, in comparison with our proximal prey in the previous figures, while our non-modulated prey still spends more time trying to feed/drink than rest, our modulated distal prey now tends to spend more of its life span on nest-related activities. The modulated prey here also now spend similar proportions of their life times searching for the nest. However, we now see a corresponding increase in the amount of time spent actively resting as hormone strength was increased. Conversely, our non-modulated prey spent the least amount of time on nest-related activities. Our prey with a hormone-release of 25 was earlier identified as that with the longest average Life Span in the distal prey versus proximal predator conditions.

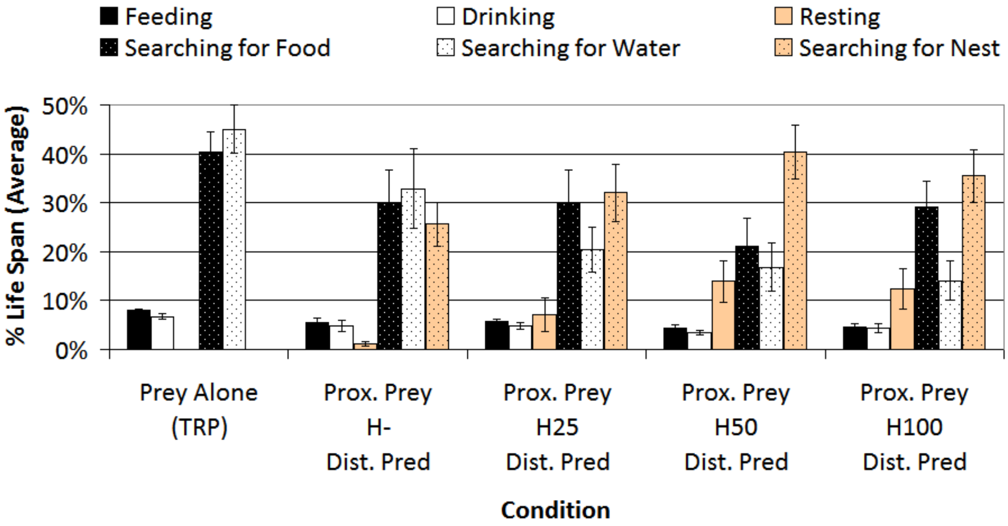

Figure 6.14: Summary of results in terms of "state sinks" across conditions where proximal prey versus distal predator. In this case, the average time (as a percentage of overall life span) spent by our different prey performing each possible activity/state out of the six available for our architecture to select from. Error Bars show the Standard Error of the Mean (SEM). As the activities that our emotion-based architecture is most likely to influence directly, we can focus our attention here on the proportion of time spent by our prey either resting or searching for the nest. From this, we can consider the impact upon having increased the perceptual abilities of our predator. For instance, in comparison with our prey in the previous figures (proximal predator conditions) we now see that our non-modulated prey spends a greater proportion of its life searching for the nest. This prey also still spends more time trying to feed/drink than rest. The prey which tended to spend most of its time on rest-related activities is our prey with a hormone release of 50. This prey actually spent the most time actively resting as well (similarly to the proximal prey versus proximal predator conditions). Our non-modulated prey tended to spend the least amount of its time on the nest-related activities.

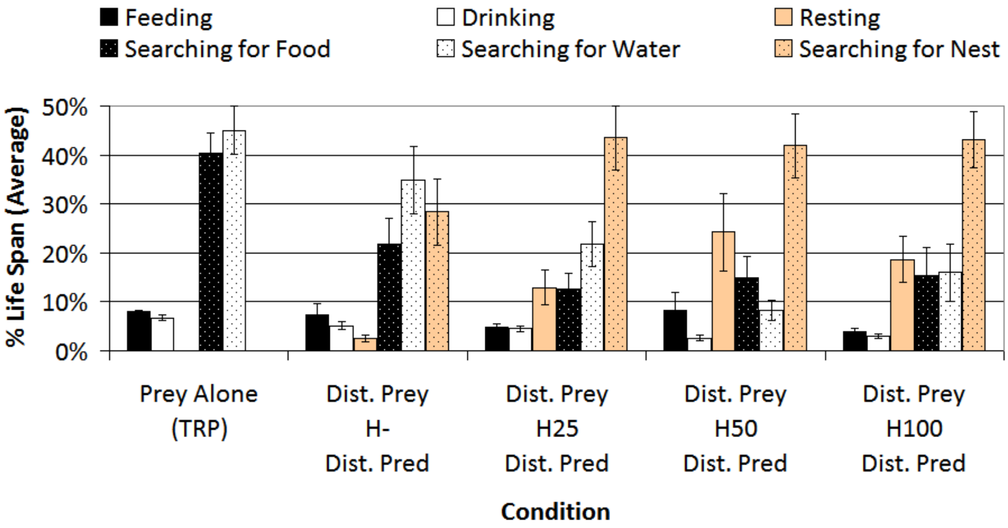

Figure 6.15: Summary of results in terms of "state sinks" across conditions where distal prey versus distal predator. In this case, the average time (as a percentage of overall life span) spent by our different prey animats performing each possible activity/state out of the six available for our architecture to select from. Error Bars show the Standard Error of the Mean (SEM). As the activities that our emotion-based architecture is most likely to influence directly, we can focus our attention here on the proportion of time spent by our prey either resting or searching for the nest. From this, we can consider the impact upon having increased the perceptual abilities of both our predator and prey. In comparison with our other prey, we again see that, in common with our distal prey versus proximal predator conditions, the modulated prey here also now spend similar proportions of their life times searching for the nest. Our non-modulated prey still tended to spend the least amount of its time on nest-related activities. Again, the prey with a hormone release of 50 actually tended to spend a larger proportion of its life actively resting.

compared to our proximal prey versus proximal predator results). In this way, while the increase in these proportions for our distal prey versus proximal predator is due to the body/distal perception of our prey, the corresponding increase for our proximal prey versus distal predator is due to the body/distal perception of our predator (our prey's environment). The distal prey versus distal predator condition then shows that when these elements are brought together, we see the majority of our prey's lifetime is spent either seeking or in its "safe place".

### 6.7.3 Brain-Body-Environment Maps

To visualise the dynamics of the prey's brain-body-environment interactions over its lifetime, Figure 6.16 shows the Brain-Body-Environment maps for a selection of runs from our experimental conditions. These go some way into better/more qualitatively illustrating the differences we have commented on across conditions. To provide more of an overview of this data, Figure 6.17 shows the average number of state transitions or "action switches" of our prey across conditions. Here we can see that, except where we had our proximal predator paired with our proximal prey, our modulated distal prey with a hormone release of 25 tended to perform a greater number of action switches. Conversely, where we had our proximal predator paired with our proximal prey, our non-modulated prey actually tended to perform a greater number of action switches.

Finally, to look further at the movement of our different animats and time spent on different resources, Figures 6.18 to 6.21 show the average time spent on different resources/areas in the arena. Together with the results for Life Span, they show both the differences in terms of the absolute number of time-steps that each type of prey spent on each resource, as well as the percentage of their "lives" this represents. The results here also show how increasing the strength the "emotion" part of our emotion-based architecture leads our prey to spend less time "risking" attack and more time "hiding" in its nest and effectively changing the "strategy" of our prey in terms of coping with the most hazardous element of its environment (the predator).

## 6.8 Discussion and Conclusion

Here, we use our model of action selection and emotion to explore the importance of both body and environment further. This chapter builds upon our previous work, constituting the next step in a systematic study of how variations in the brain-body interactions of prey and predator can affect the dynamics of prey-predator interactions. In this chapter, we continue to use and explore our implementation of the Hazardous Three Resource Problem. In these experiments the hormone-like mechanism used to simulate "fear-like" states, affecting the prey architecture through alteration of the level of our third internal variable (integrity).

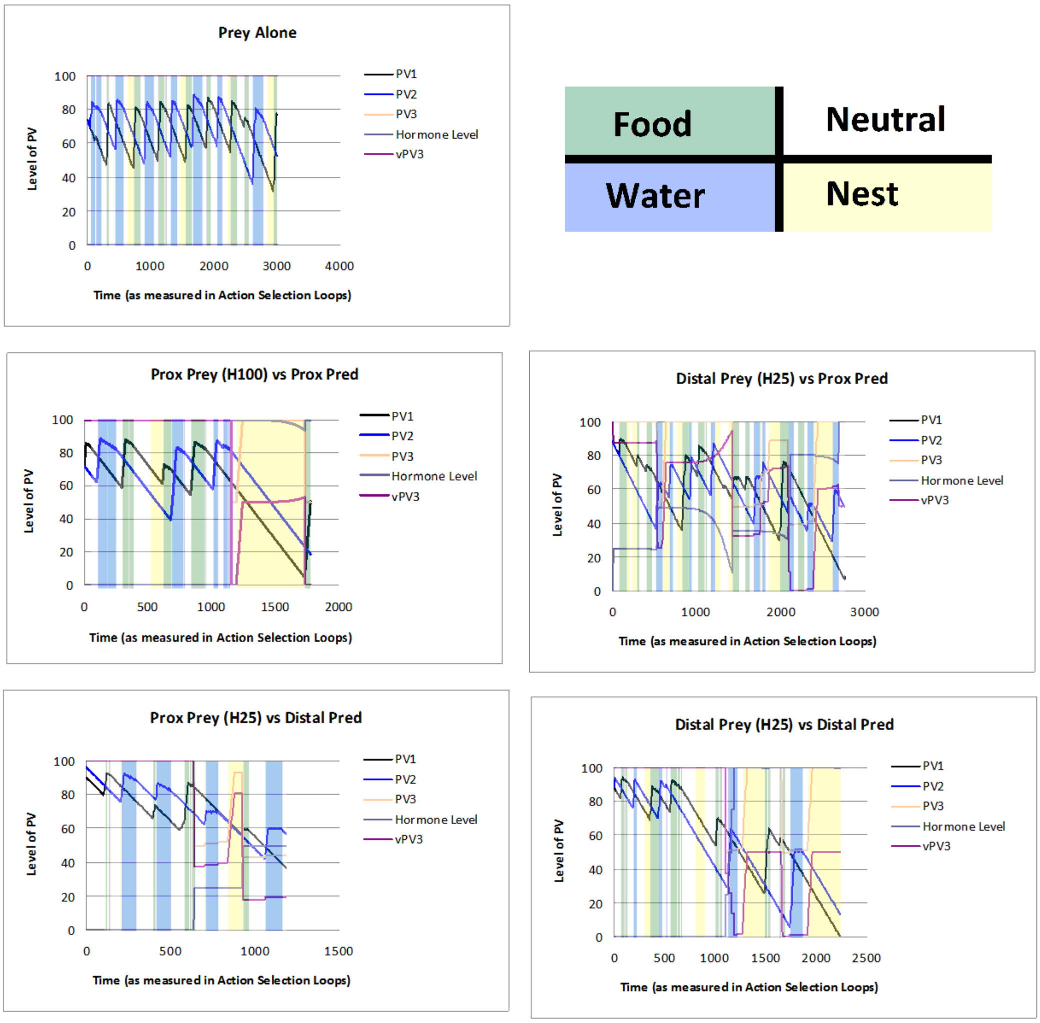

Figure 6.16: Illustration of differences (relating to aspects of prey Brain, Body and Environment over time) as observed across conditions between individual runs and shown by our BBE Maps. The background bars here represent our prey's position within its environment (green bars show where our prey detected "food", blue bars show "water", the absence of any bar indicates the prey detected no resources present in its immediate environment at that time and yellow bars indicate the prey was in its "nest") while each line represents the internal variables of our prey, including each of its physiological variables (black for blood sugar, blue for vascular volume, yellow for integrity) and hormone levels/effects. Additionally, here vPV3 represents the prey's *perceived* level of integrity, which in the modulated conditions differed from the *actual* level of integrity when concentrations of hormone were released.

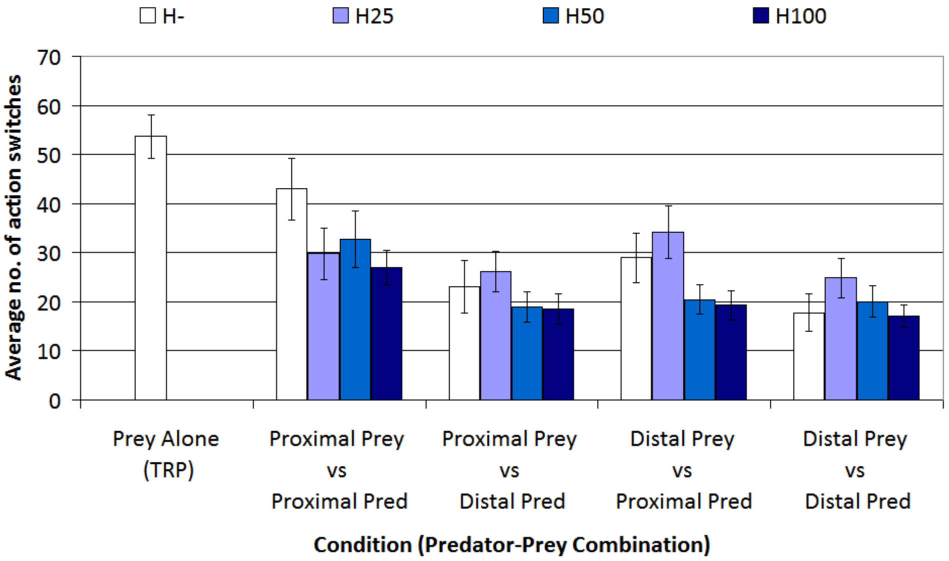

Figure 6.17: Summary of results showing the average number of "action switches" of our prey across conditions. Error bars show the Standard Error of the Mean (SEM). Exploring the idea of action selection as a cycle of activities, rather than a series of individual decisions, we use this simple measure to compare our prey. An "action switch" occurs whenever our prey makes the decision to switch from one of the six actions available to it to another. Here we can see that, except where we had our proximal predator paired with our proximal prey, our modulated distal prey with a hormone release of 25 tended to perform a greater number of action switches. Conversely, where we had our proximal predator paired with our proximal prey, our non-modulated prey actually tended to perform a greater number of action switches.

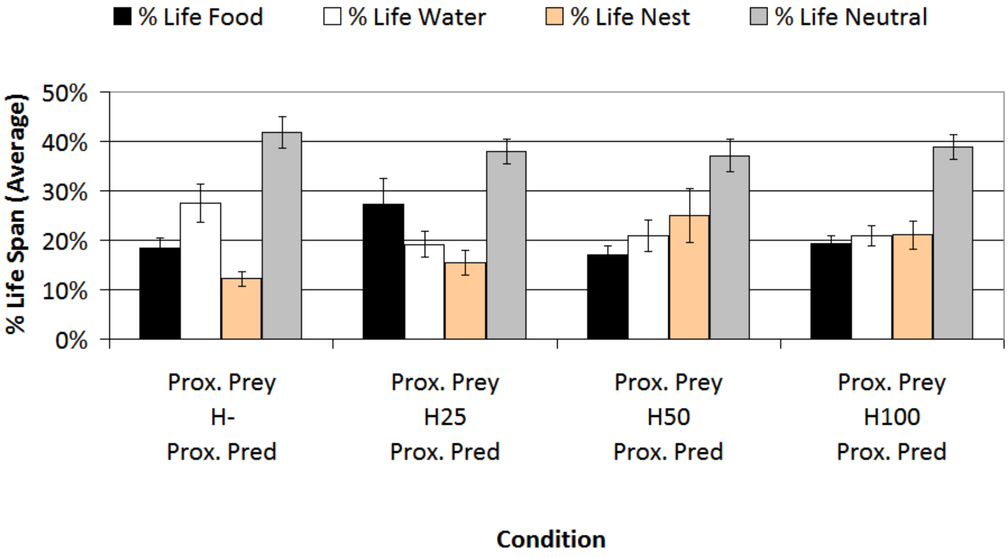

Figure 6.18: Summary of results showing time (as the average proportion of our prey's life span) spent on different resources/areas in its environment/arena by our prey across conditions where proximal prey versus proximal predator. Error Bars show the Standard Error of the Mean (SEM). This shows that, on average, our non-modulated prey tended to spend a greater proportion of its life span on neutral spaces within the arena whereas our modulated prey tended to spend a greater proportion actually on a resource.

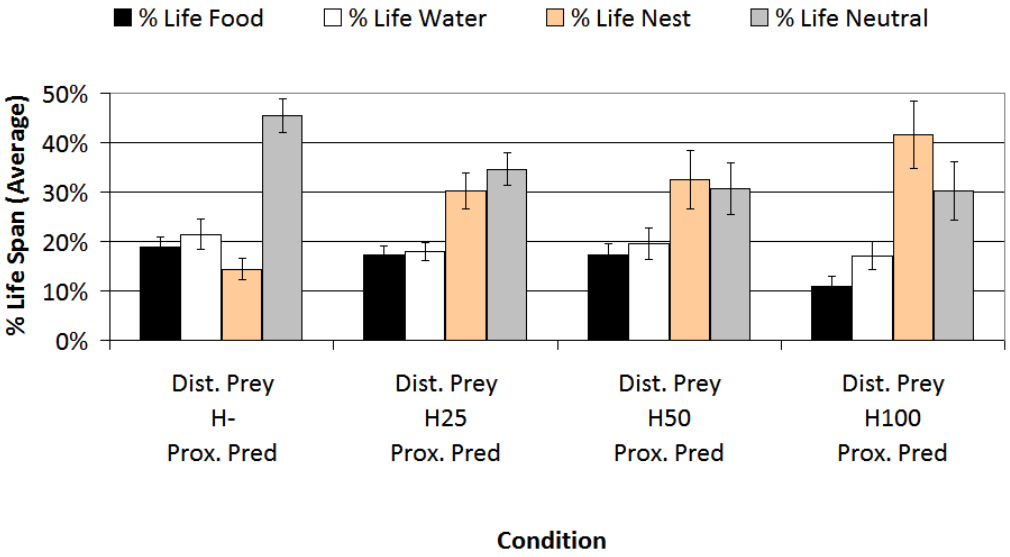

Figure 6.19: Summary of results showing time (as the average proportion of our prey's life span) spent on different resources/areas in its environment/arena by our prey across conditions where distal prey versus proximal predator. Error Bars show the Standard Error of the Mean (SEM). This shows that, on average, our non-modulated prey still tended to spend a greater proportion of its life span on neutral spaces (approximately 40%) whereas our modulated prey tended to spend a greater proportion on a resource. As the prey with the longest average Life Span in this condition, our prey with the hormone release of 25 spent similar amounts of time on either neutral space or in the nest (approximately 30%) and less (but similar, smaller proportions of time (just under 20%) on either food or water spaces.

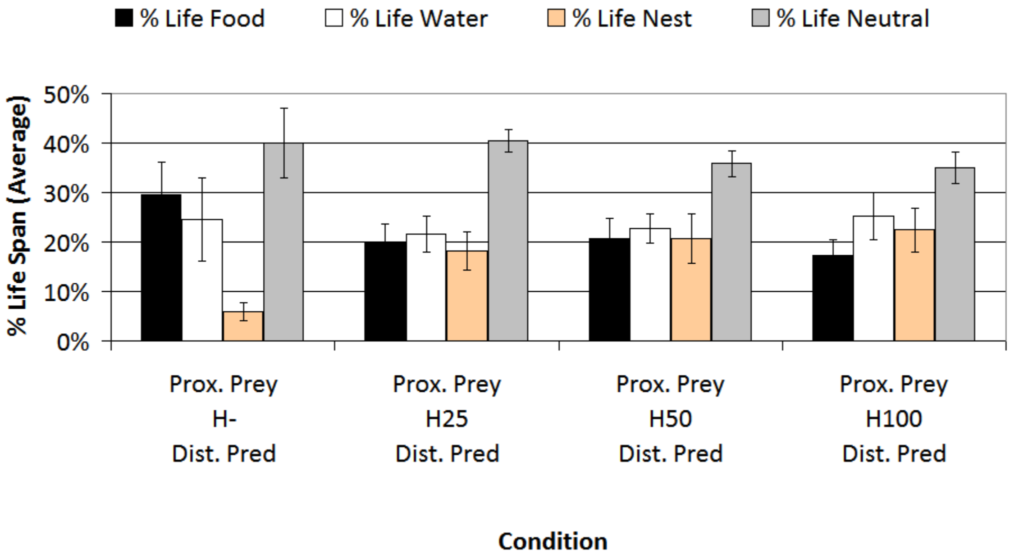

Figure 6.20: Summary of results showing time (as the average proportion of our prey's life span) spent on different resources/areas in its environment/arena by our prey across conditions where proximal prey versus distal predator. Error Bars show the Standard Error of the Mean (SEM). As the prey with the longest average Life Span in this condition, our prey with the hormone release of 25 can be seen to have spent similar amounts of time on neutral space to our non-modulated prey (approximately 40%) but rather than spending, on average, around 5% of its life span in the nest, this prey spent similar proportions of its time on each resource (approximately 20%).

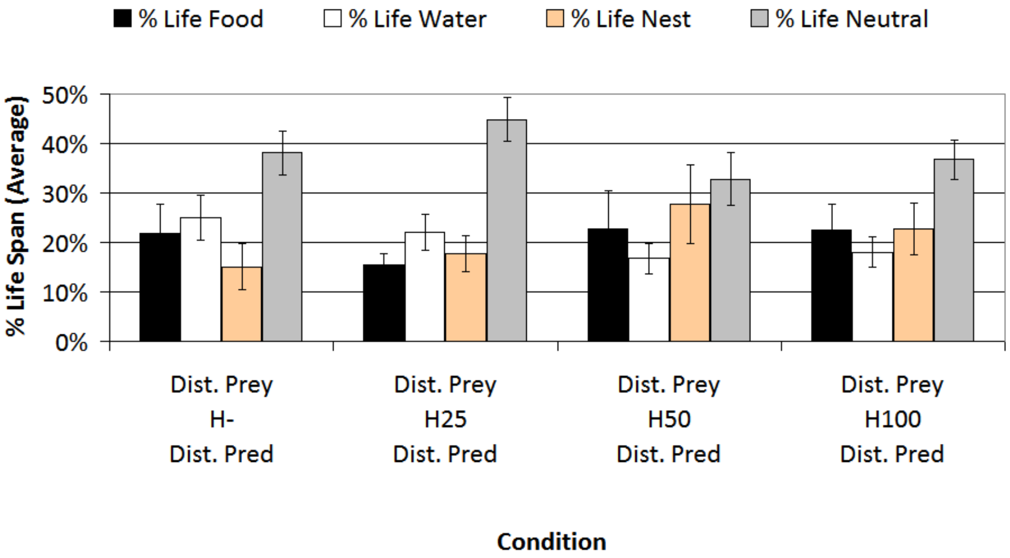

Figure 6.21: Summary of results showing time (as the average proportion of our prey's life span) spent on different resources/areas in its environment/arena by our prey across conditions where distal prey versus distal predator. Error Bars show the Standard Error of the Mean (SEM). As the prey with the longest average Life Span in this condition, our prey with the hormone release of 25 can be seen to have spent greater amounts of time on neutral space and more time in the nest than our non-modulated prey (spending similar amounts of time on each resource).

We have compared the adaptive value of our chosen mechanism (affecting interoception of our prey's internal physiology) across different combinations of predator and prey: also varying parameters and modifying this mechanism itself to see how this would affect our prey life span. We found significant differences between conditions, particularly in terms of our predator's abilities (see Figures 6.2). Taken together with further descriptive statistics, we suggest these results indicate that our mechanism could be more "adaptive" — or in terms of our model have a particular emergent function or specific system property that is adaptive — in cases where our prey is against predators that are both dangerous and predictable. We conclude that in the context of a predator-prey relationship, the success of the use of such a mechanism as ours will in fact most likely depend on getting the balance "just right" between aspects of prey brain, body and environment. In terms of body, this includes the aspect of prey perceptual distance, hormone strength and type. In terms of environment, this includes the aspect of predator body (particularly perceptual distance).

More generally, we can again use our results (Figures 6.2 and 6.16 drawing the majority of our interest here) to suggest that the adaptive value of our prey's emotion-based architecture varies with aspects of body such as perceptual distance (as an aspect of embodiment) demonstrating how the emergent properties (such as prey "flight") of our emotion-based architecture varied with our prey's brain-body combination. Interestingly, this fits in with ethological data which has found that certain prey animal's flight initiation distances (that is, the point at which they "decide" to "flee") vary with predator lethality [109]. Taken together, however, we suggest this could indicate a wider role or function for our hormone-like mechanism in terms of acting as (effectively) a risk assessor, assessing the risk of the environment dynamics (predator) as well as determining how willing our prey is to "risk" attack in order to satisfy its other survival needs (such as foraging).

Thinking back to our earlier review of the literature, in terms of the wider research area, we hope that the results of this chapter in particular may provide further inspiration and encouragement for greater collaboration and exchange between biological and artificial research areas. For, in attempting to more systematically explore the possible role of brain, body and environment in shaping an artificial predator-prey relationship, we can also see a possible application for our robotic implementation for more specifically allowing ethologists to better test and develop their ideas and theories such as Cooper's on flight initiation distance [88].

# Chapter 7

# Sensory Integration and Resource Allocation

*"The degree of one's emotions varies inversely with one's knowledge of the facts."*
— **Bertrand Russell**

## 7.1 Emergent Functions of Modulating Interoception: Results Chapter III.

One of the arguments inherent and central to this thesis is that, by increasing our understanding of brain-body-environment interactions, such knowledge will enable us to build agents that both solve the problem of action selection and, in doing so, will one day display adaptive behaviour on a level with our own. From our own review of the existing literature, and including the results of our experiments, we argue this could well include the incorporation of mechanisms that generate behaviours which can be considered functional equivalents to biological emotions. So far, we have focused on our ideas about both predator and prey agents in terms of the importance of aspects of both these types of agent's embodiment (that is, their perceptual abilities) and their potential effects on the adaptive value of our emotion-based architecture. This includes the consequent effect on the predator-prey relationship itself.

However, here we now proceed to consider how we might develop our actual mechanism to make the most of its potential. That is to say, we look at how we might use it so as to achieve further (multiple) functionalities. In terms of artificial emotion, such functions might include those that somehow act to make connections between brain, body and environment more useful, whether in terms of efficiency (such as speed of response) or cost-effectiveness (as with risk assessment). Relating to our concepts of adaptive emotion in brain-body-environment relationships, in this chapter we explore how we might use our conceptual framework to extend our ideas and develop our agents to manipulate this relationship and its interactions adaptively. In this chapter we continue to build and improve on our system step-by-step similar to the way in which neuroscientists imagine our own brains (and other species) have evolved emotions that are adaptive. As part of our incremental approach, we therefore think about what we might begin to add or take away next.

## 7.2 Sensory Integration and Flight Initiation Distance (FID)

Biologically, sensory integration can be defined as the neurological process that organises an agent's senses from its own body and the environment. Thus it deals with how the brain processes multiple sensory modality inputs into usable functional outputs. Interestingly, researchers note that detecting the positions of objects in the environment, including predators or prey, is of crucial importance to animals. It further seems that several of the sensory systems, especially vision and hearing, are particularly specialized for this task. In this chapter, continuing our biologically-inspired approach, we borrow from biological ideas of sensory integration and focus our attention on trying to demonstrate, as well as understand, the potential of our emotion-based architecture with regards to adaptive action selection. Specifically, in this chapter we demonstrate how our "emotion-like" substrate might successfully be integrated further with and, in turn, actively integrate our prey agent's brain, body and environment.

To do this, we have developed our basic emotion-based architecture — specifically, its hormone-like mechanism — by integrating it in such a way as to produce more *gradient* responses/effects. Looking at ethological research, we considered such a development might have an advantage in being somewhat more biologically plausible and, in turn, ultimately more adaptive for our prey. In the next chapter, the work of this chapter will enable us to not only compare all our findings together, but to better compare our own system with those found in biology. In particular, with regards to how our work might fit into the very same concepts and framework that Cooper [88] introduced in his work to explain the adaptive behaviour of lizards — which is effectively about when they make the decision and what information they use to make the decision (about when) to "flee" from a *perceived* predator (or, to borrow another term from ethology, how our artificial prey's *flight initiation distance* might be varied to generate more adaptive behaviours in response to the dynamics of its environment).

## 7.3 Gradient Perception and Response

Using our implementation of the H3RP, to consider the potential of our emotion-based architecture further the rest of this chapter will report the results of two different sets of experiments. The first (outlined in the first part of this chapter) were actually an extension to those of the previous chapter. However, in our second (and final) experiments, we aimed to build on the results obtained so as to consider how we might create more dynamic predator-prey relationships. That is, more in line with biological examples — using our hormone-like mechanism to vary the *intensity* of prey response. In common with our other experiments, however, we still primarily aim to consider our results primarily in terms of action selection, emotion and brain-body-environment interactions in the predator-prey relationship.

## 7.4 I. Fixed versus Gradient Hormone Release (Varying Sensory Signals)

Using the same H3RP as detailed in the last chapter, in the first experiments examined here we varied parameters of our prey's hormone-like mechanism, testing the effects of using a new "gradient" hormone release for our distal prey in order to compare it with our otherwise previously "fixed" hormone release (in our new gradient mechanism, the amount of hormone released by the prey's mechanism is calculated as a function of how far away the predator is perceived to be). Testing the same four combinations of predator (proximal/distal) and prey (proximal/distal) as before, we thereby created a further six new conditions for testing our new type of hormone-like mechanism, which demonstrates a more *gradient* response than the original one — again also varying prey and predator perceptual distance and hormone strength. For comparison, and focusing on our distal prey, we looked at a total of twelve experimental conditions for this data set by also including the results we had already obtained from the six distal prey conditions described in the previous chapter. The purpose of this was to better enable (more informative) comparisons to be made between these conditions and the previously-studied response of our prey. Our focus was especially on the difference, if any, that the design of a more fine-grained (or integrated) hormonal mechanism might make to the prey's survival.

## 7.5 Prey

Because our first experiments here were more an extension of our previous set, the only animat to change was the prey (the different predators and arena having been already described in the last chapter).

|  | | | 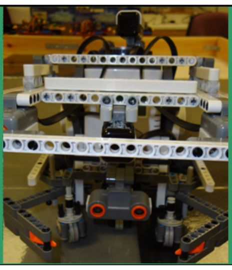 | |
|---|---|---|---|---|
| *Variations in Prey "Body" [Hormone Release and Strength]** | | | *Variations in Predator "Body" [Perceptual Distance]** | |
| **Perceptual Distance** | **Hormone Release** | **Hormone Strength** | **Proximal Predator** | **Distal Predator** |
| **Distal Prey** | **Fixed Release** | **H25** | Distal Prey (Fixed H Release)<br>vs<br>Prox Pred | Distal Prey (Fixed H Release)<br>vs<br>Distal Pred |
| | | **H50** | | |
| | | **H100** | | |
| | **Gradient Release** | **H25** | Distal Prey (Gradient H Release)<br>vs<br>Prox Pred | Distal Prey (Gradient H Release)<br>vs<br>Distal Pred |
| | | **H50** | | |
| | | **H100** | | |

Figure 7.1: Summary of experimental conditions: In these experiments a factorial 2x3x2 design (resulting in twelve experimental conditions) was used to investigate further how the adaptive value of our emotion-based architecture varies when we both vary both the strength and release type of our hormone-like mechanism and the perceptual abilities of our prey *and* predator. This design enables us to identify the effects of a) the "emotion-like" substrate represented by our hormone-like mechanism, using different hormone release rates and types (fixed/gradient release) to represent further variation in our prey's "brain" and "body" and b) our predator's perceptual distance; proximal/distal detection representing both a variation in our predator's "body" and prey's "environment". *Variations in the body of one agent potentially leading to a variation in the other's environment and vice versa.

### 7.5.1 Interoceptive Modulation (Graded Hormone-Release of 0, 25, 50 and 100)

Using the same basic distal prey as before, for these experiments we created a new, "gradient" hormone-like mechanism by integrating information from our distal sensor further. This "new" mechanism calculated the intensity of the external stimulus of the predator as a function of the distance away the predator is perceived. If the predator is perceived via the ultrasonic sensor more than 30cm away less hormone is released than if the predator is perceived via an actual attack (where the maximum level of hormone is released). Similarly to the previous experiments conducted using our old "fixed" hormone-release rate, we tested our new "gradient" hormone-release rate at four levels: 0, 25, 50 and 100. To recap, a level of 25 in the "fixed" condition means a release of 25 percent of the maximum concentration of hormone, regardless of whether the predator is detected via an attack or the ultrasonic. Meanwhile, a level of 25 in the "gradient" condition means a release of hormone as a function of the distance away the predator is detected: full release of 25 percent of the maximum concentration occurring after an attack, otherwise a percentage of that 25 percent.

## 7.6 Results

For each condition, data was collected from fifteen runs, collecting the same type of data and using the same procedure as reported in the previous chapter (resulting in a total of a further 90 new runs for our analysis). When combined with the runs obtained in the last set of experiments, our resulting data set consists of 180 runs (across twelve conditions). Figures 7.2 to 7.6 show how each type of prey performed as measured by Life Span, Overall Comfort and Physiological Balance. Again, to provide a baseline, we include the results collected from the prey alone as well. That is, performing the less dynamic and underlying two resource problem (in the absence of the predator).

### 7.6.1 Measures of Viability

From the data collected, we again first calculated and looked at the results in terms of our chosen measures of viability (previously outlined along with our basic implementation). Figures 7.2-7.6 show how each type of prey performed as measured by Life Span, Overall Comfort and Physiological Balance. Again, for purposes of comparison, we have also included the results obtained from testing our prey alone, without a predator agent. The purpose of this was to again check that our predators were actually having an impact on the prey in terms of representing a hazardous resource.

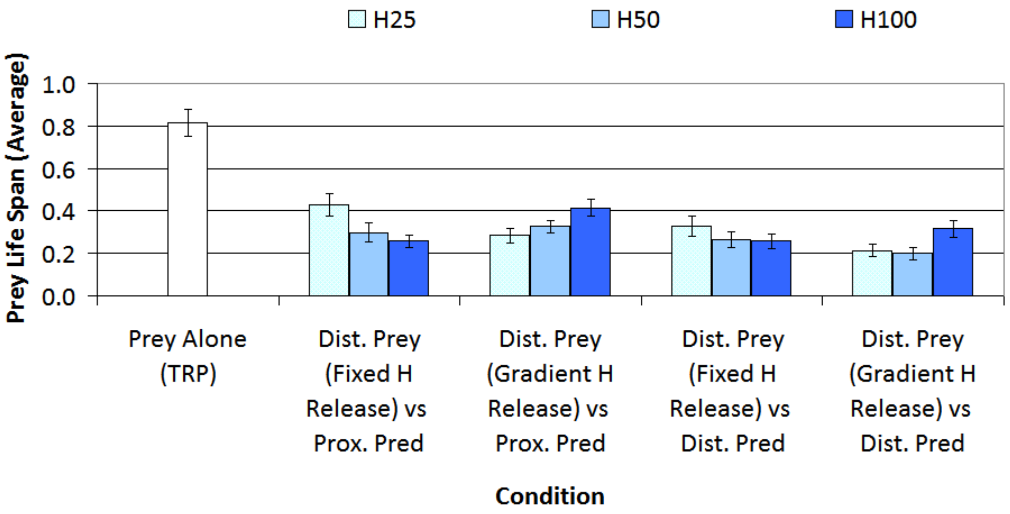

Figure 7.2: Summary of results showing average Life Span (LS) of our prey across experimental conditions. Error Bars show the Standard Error of the Mean (SEM). This shows the average for each condition time-steps. Interestingly, we see a mirror image in terms of the trends shown here for the types of prey tested using the “fixed” hormone release, as opposed to those tested using the newer “gradient” hormone release. That is to say, we observe a tendency for LS to decrease with an increase in hormone strength combined with fixed hormone release. However, we see the opposite tendency with an increase in hormone strength when combined with a gradient hormone release. A significant main effect for predator perceptual distance was found on LS. While no main effects were found for our other independent variables, an interaction was suggested between hormone type and strength.

#### 7.6.1.1 Life Span, Overall Comfort and Physiological Balance

In terms of Life Span, in Figure 7.2 we again show the average for each condition in time-steps. Interestingly, but perhaps not surprisingly we see a mirror image in terms of Life Span for the types of prey tested using the “fixed” hormone release, as opposed to those tested using the newer “gradient” hormone release.

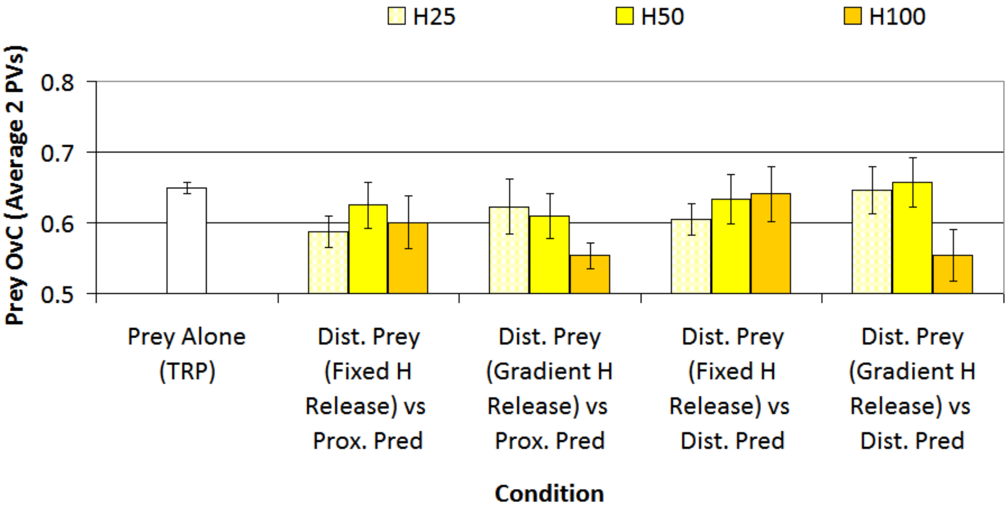

Figure 7.3: Summary of results showing average Overall Comfort (OvC) of our prey across experimental conditions. Error Bars show the Standard Error of the Mean (SEM). In this case, calculations of OvC exclude the third variable of integrity. In comparison to the patterns shown for Life Span (LS) in Figure 7.2 from this it seems that, although we can trace a mirror image in terms of “fixed” hormone release as opposed to the newer “gradient” hormone release here too, the trends are generally in the opposite direction. That is to say, the longer our prey’s average LS, the lower their OvC (excluding the third variable of integrity). The condition that does not show this quite so clearly is that in which our distal prey was given a fixed hormone release and paired with our proximal predator.

Performing a multi-way ANOVA (2x2x3) we looked for main effects of our independent variables of predator perceptual distance (proximal/distal) hormone type (fixed/graded) and hormone strength (25, 50, 100). Again, as our original results also showed, after arcsine transformation, a significant main effect for predator perceptual distance was found (p=0.002) for Life Span in time-steps. No main effects were found for our other independent variables. However, an interaction was suggested between hormone type (fixed/graded) and strength (25, 50, 100) (p=0.000). Likewise, we found the same effects for Life span in minutes for both predator perceptual distance (p=0.000) and the same interaction between hormone type and strength (p=0.001). In

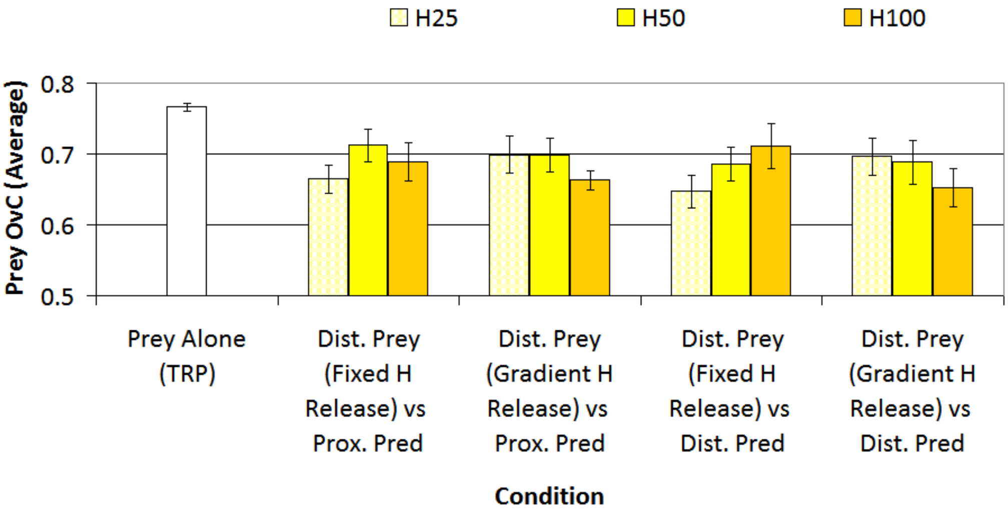

Figure 7.4: Summary of results showing average Overall Comfort (OvC) of our prey across experimental conditions. Error Bars show the Standard Error of the Mean (SEM). In this case, calculations of OvC include the third variable of integrity. Here we can trace the same general trends per condition shown in Figure 7.3 excluding the third variable. However, the conditions which stand out to us are those for our distal prey with gradient hormone release of 100. For, together, these results appear to suggest greater differences between calculations of OvC using 2 PVs or 3 PVs in relation to the same calculations for conditions with lower hormone releases, particularly against our distal predator. To us, this implies that this type of prey was more likely to sacrifice/lower the OvC of our two original variables so as to increase the average level of integrity than its fixed hormone release counterpart.

terms of Overall Comfort and Physiological Balance, no significant differences were found for Overall Comfort. Nevertheless, an interesting result was found when testing the interaction between hormone type and strength (p=0.056). A main effect was found for predator perceptual distance for Physiological Balance (p=0.027). However, no significant interactions were found between variables.

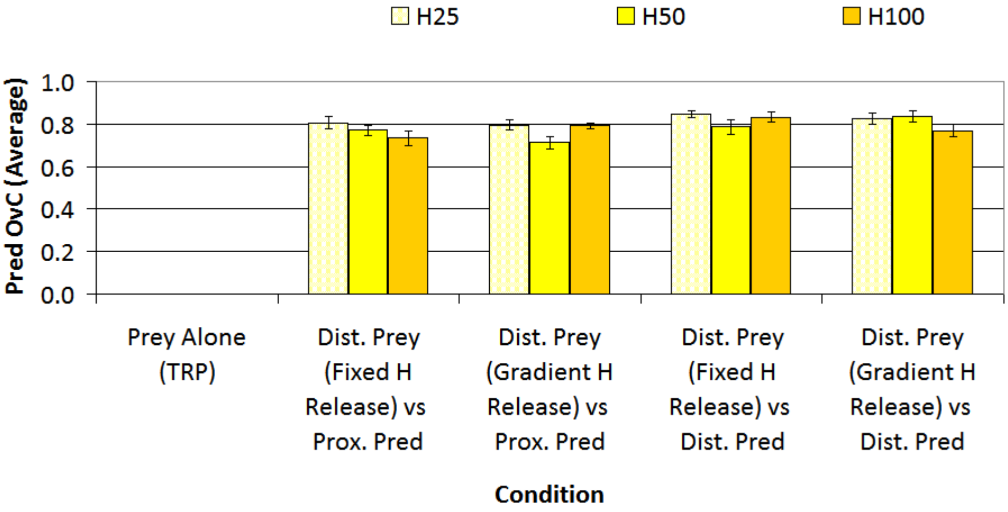

Figure 7.5: Summary of results showing average Overall Comfort (OvC) of our predator across experimental conditions. Error Bars show the Standard Error of the Mean (SEM). Though it might be expected that the OvC of our predator would be the opposite to that shown for our prey, these results show that this was not necessarily the case.

#### 7.6.1.2 Risk of Death Management

In terms of Management of Death Risk (as previously defined) Figures 7.7 to 7.10 show these results in terms of Risk of Death (RoD). Again, each bar represents the average percentage of time our prey was in each “zone”. As before, a RoD of 1 means the prey is almost, if not *the* most “healthy” it can be, whereas a RoD of 10 means the prey is “near death”. In terms of our “gradient” hormone-release conditions, these results are interesting in that, for these conditions, our calculations seems to show the least (relative) variability when combined with the potential hormone-release strength of 100. Perhaps this reflects the case that in the “fixed release” conditions, hormone-release does not incorporate the additional information that the “gradient release” conditions do in terms of the distance away the predator is perceived. This would then further support the conclusions of the previous chapters (hormone-release having a role as both

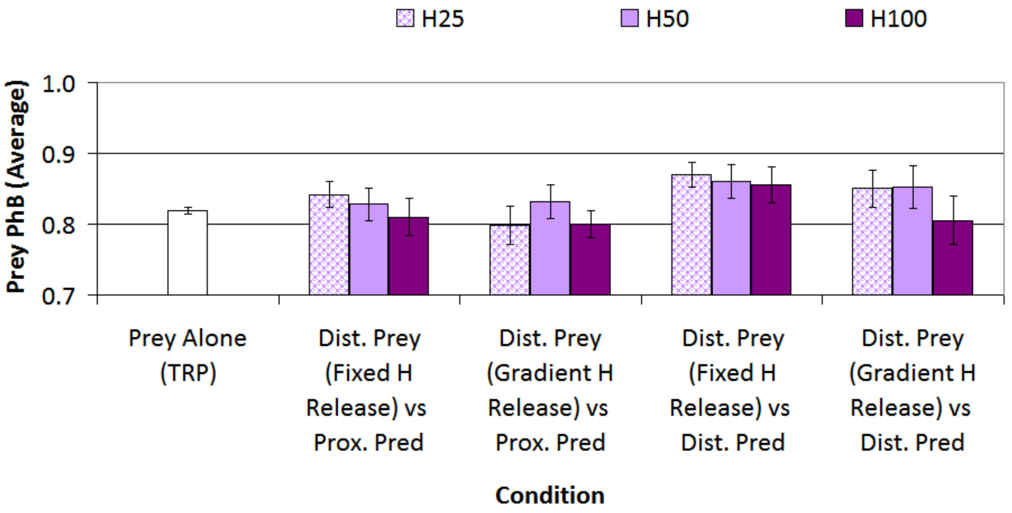

Figure 7.6: Summary of results showing average Physiological Balance (PhB) of our prey across experimental conditions. Error Bars show the Standard Error of the Mean (SEM). In this case, calculations of PhB include the third variable of integrity. The condition showing the greatest average PhB was that for our distal prey with a fixed hormone release when paired with our distal predator. The conditions showing lowest average PhB tend to be those for prey with a hormone release of 100.

a predictor of attack and assessment of risk). To see how this relates to our predator's experience, Figures 7.11-7.14 show calculations of Risk of Death for our predator.

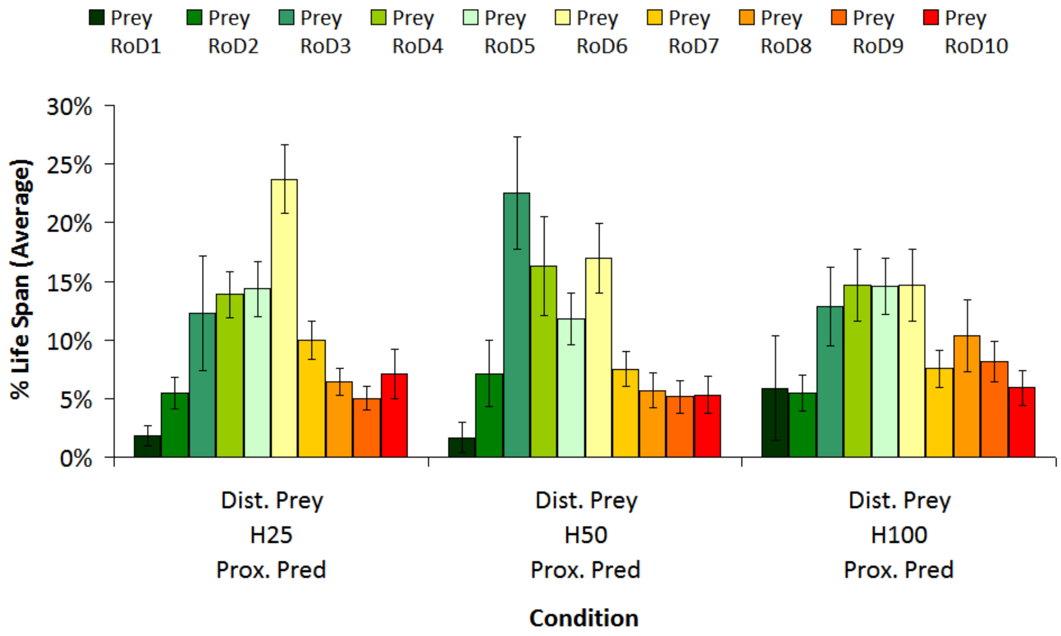

Figure 7.7: Summary of results showing Risk of Death (RoD) Management of our prey across conditions where distal prey versus proximal predator and the hormone release of our emotion-based architecture is fixed. Each bar represents the average percentage of time our prey was in each RoD "zone". Error Bars show the Standard Error of the Mean (SEM). A RoD of 1 (green) means that the prey is almost if not exactly the most "healthy" it can be, whereas a RoD of 10 (red) means the prey is "near death". Here we see a more even distribution for our prey with a fixed hormone release of 100. From this, our prey with a fixed hormone release of 50 would seem to have tended to spend a greater proportion of time in a more "healthy" zone. However, considering Life Span, it was our prey with a fixed hormone release of 25 that lived longest.

At the risk of introducing a little anthropomorphism to better illustrate our findings, from these results and our own observations it did seem that, in the "fixed" hormone-release condition our prey suffered from something akin to "anxiety", unnecessarily sacrificing opportunities to feed and drink when detecting the predator in front of it. With the prey in the "gradient" hormone-release condition though, we saw something more akin an actual "assessment" taking place (a result of the gradual build-up of hormone in the prey's system with the approach of its predator). In some cases, the prey therefore actually appeared to "wait" for the predator to approach further before moving. The emergence of such behaviour is particularly interesting in light of the fact that, as in nature, instant "flight" does not always prove optimal for our prey: whether

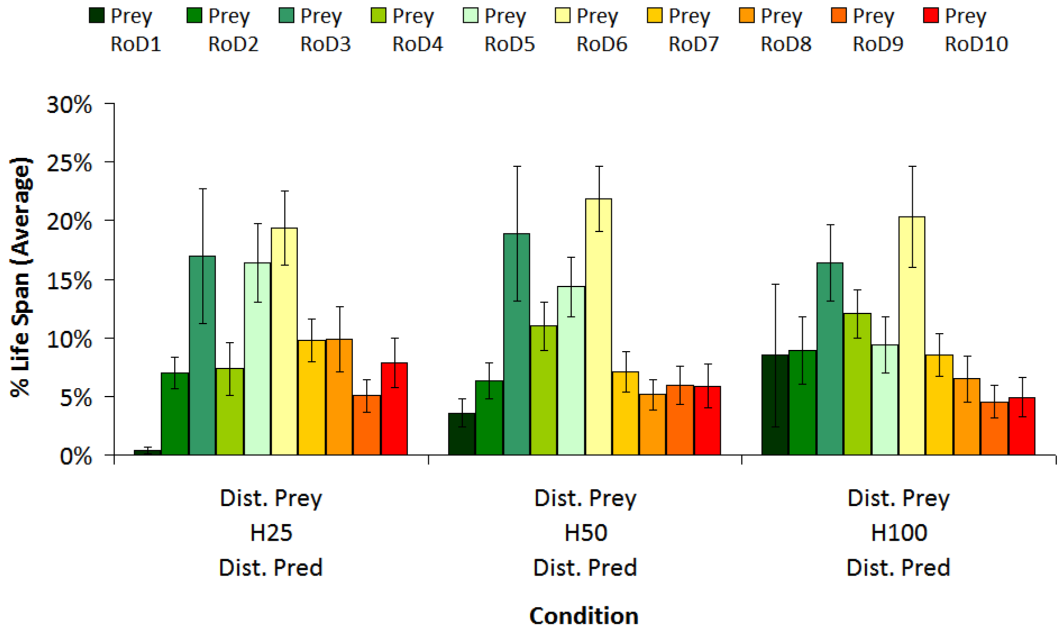

Figure 7.8: Summary of results showing Risk of Death (RoD) Management of our prey across conditions where distal prey versus distal predator and the hormone release of our emotion-based architecture is fixed. Each bar represents the average percentage of time our prey was in each RoD “zone”. Error Bars show the Standard Error of the Mean (SEM). A RoD of 1 (green) means that the prey is almost if not exactly the most “healthy” it can be, whereas a RoD of 10 (red) means the prey is “near death”. Here we see quite similar results for our prey with a fixed hormone release. That is to say, the greatest proportion of their time was spent in “zones” 3 and 6.

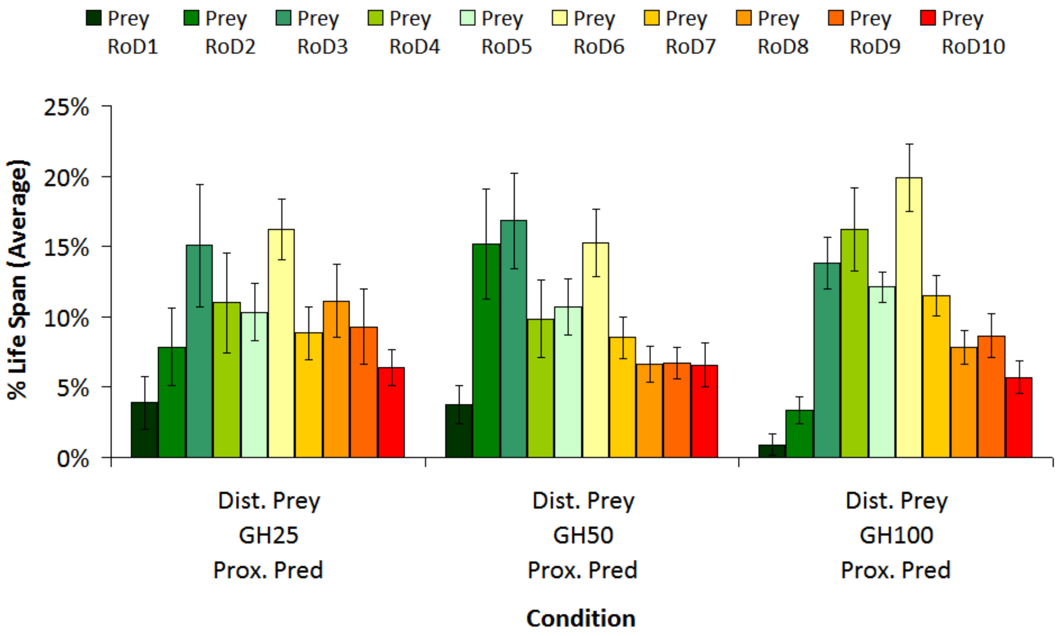

Figure 7.9: Summary of results showing Risk of Death (RoD) Management of our prey across conditions where distal prey versus proximal predator and the hormone release of our emotion-based architecture follows a gradient release (the amount of hormone released being a function of the distance at which the predator is perceived). Each bar represents the average percentage of time our prey was in each RoD "zone". Error Bars show the Standard Error of the Mean (SEM). A RoD of 1 (green) means that the prey is almost if not exactly the most "healthy" it can be, whereas a RoD of 10 (red) means the prey is "near death". The highest peaked distribution is shown here for a hormone release of 100. This hormone release further displayed the highest average life span out of these three conditions.

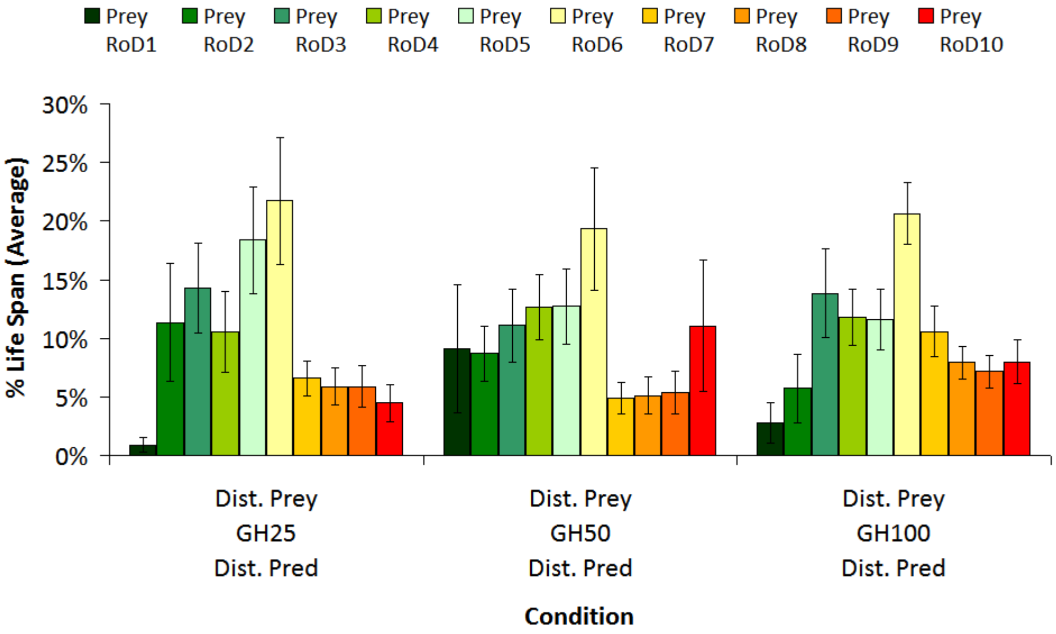

Figure 7.10: Summary of results showing Risk of Death (RoD) Management of our prey across conditions where distal prey versus distal predator and the hormone release of our emotion-based architecture follows a gradient release (the amount of hormone released being a function of the distance at which the predator is perceived). Each bar represents the average percentage of time our prey was in each RoD “zone”. Error Bars show the Standard Error of the Mean (SEM). A RoD of 1 (green) means that the prey is almost if not exactly the most “healthy” it can be, whereas a RoD of 10 (red) means the prey is “near death”. Here the prey with a hormone release of 100 appeared to spend, on average, a higher proportion of its life span in “zones” 7-10. However, this was also the prey with the longest life span of the three.

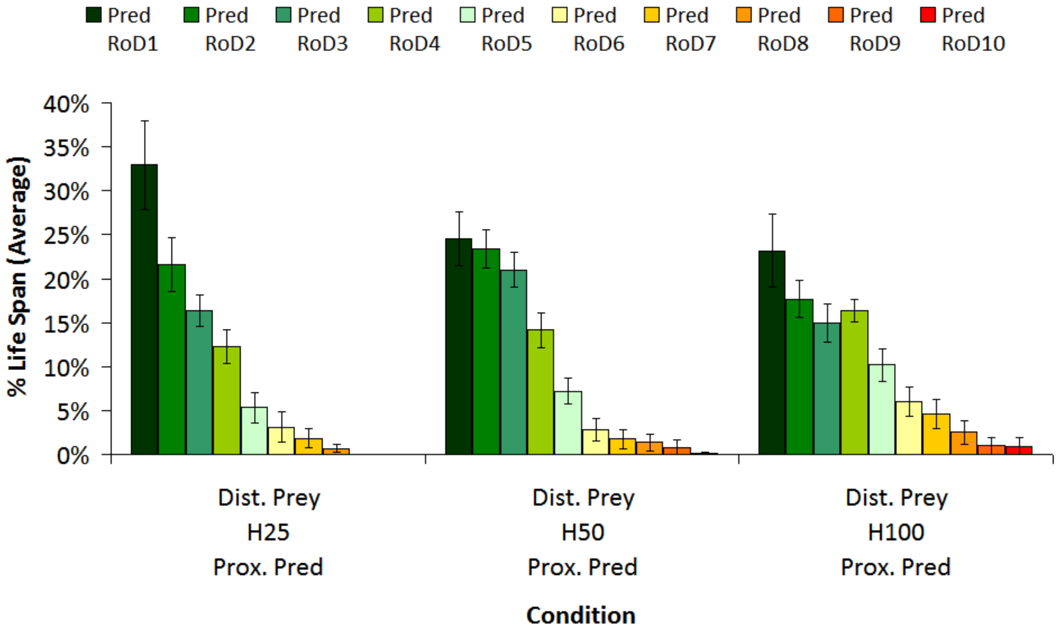

Figure 7.11: Summary of results showing Risk of Death (RoD) Management for our predator across conditions where distal prey versus proximal predator and the hormone release of our emotion-based architecture is fixed. Each bar represents the average percentage of time our predator was in each RoD "zone". Error Bars show the Standard Error of the Mean (SEM). A RoD of 1 (green) means that the predator is almost if not exactly the most "healthy" it can be, whereas a RoD of 10 (red) means the predator is "near death". Here we see that our predator spent most of its life in a "healthy" state.

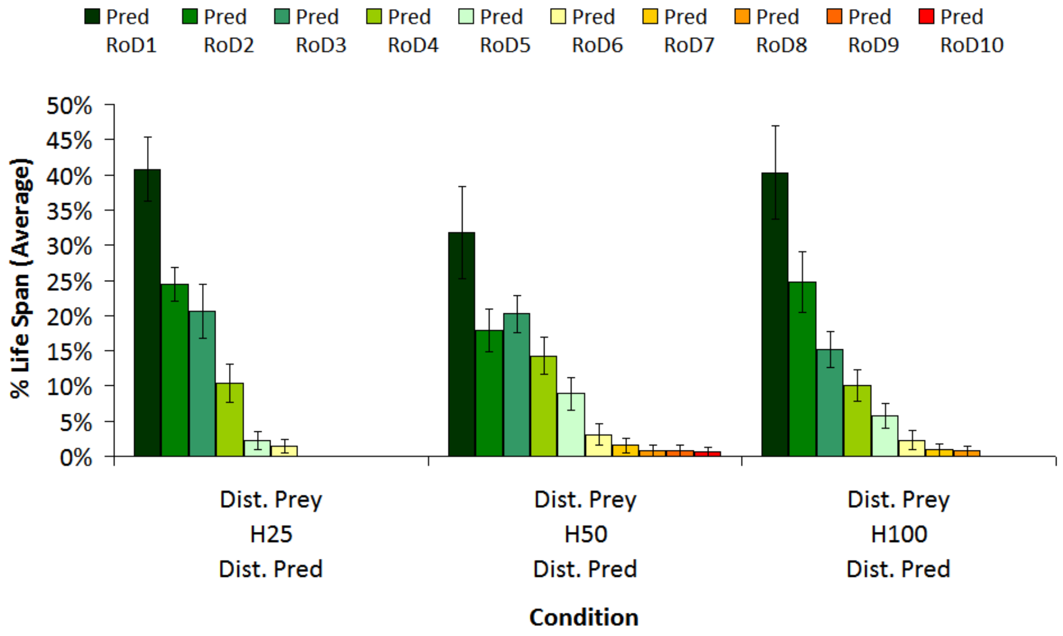

Figure 7.12: Summary of results showing Risk of Death (RoD) Management for our predator across conditions where distal prey versus distal predator and the hormone release of our emotion-based architecture is fixed. Each bar represents the average percentage of time our predator was in each RoD “zone”. Error Bars show the Standard Error of the Mean (SEM). A RoD of 1 (green) means that the predator is almost if not exactly the most “healthy” it can be, whereas a RoD of 10 (red) means the predator is “near death”. Here we see that our predator spent most of its life in a “healthy” state. However, the peaks shown here do not quite reach the same levels as those in Figure 7.11.

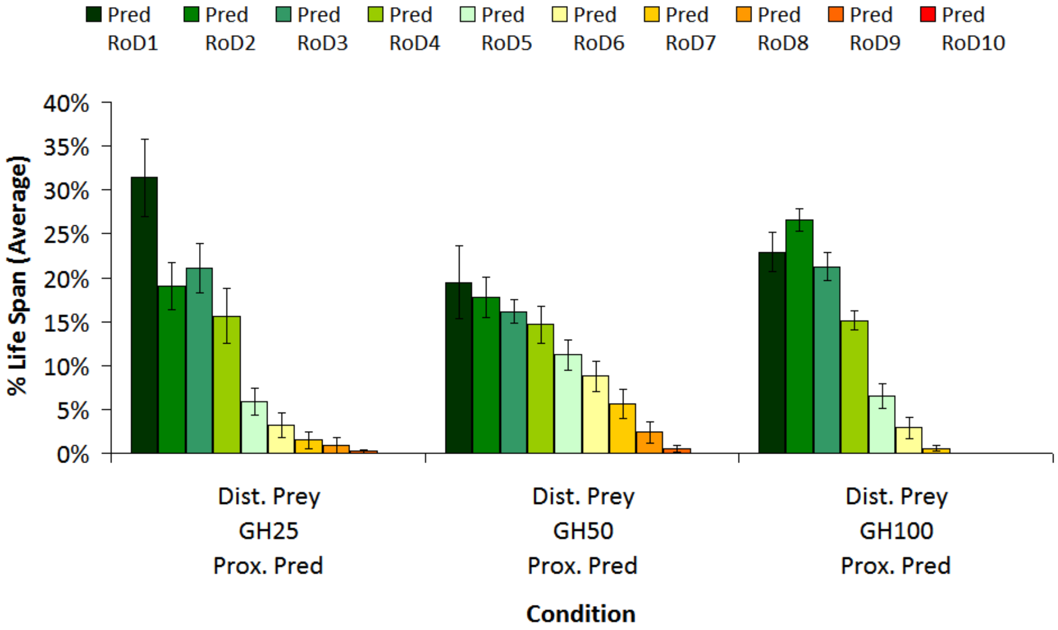

Figure 7.13: Summary of results showing Risk of Death (RoD) Management for our predator across conditions where distal prey versus proximal predator and the hormone release of our emotion-based architecture follows a gradient release (the amount of hormone released being a function of the distance at which the predator is perceived). Each bar represents the average percentage of time our predator was in each RoD "zone". Error Bars show the Standard Error of the Mean (SEM). A RoD of 1 (green) means that the predator is almost if not exactly the most "healthy" it can be, whereas a RoD of 10 (red) means the predator is "near death". Here we see the prey with a release rate of 25 allowed our predator to live a life spending the greatest proportion of that time in the most "healthy" state possible. Not surprisingly, as shown in Figure 7.2 this led to a lower average life span for the prey.

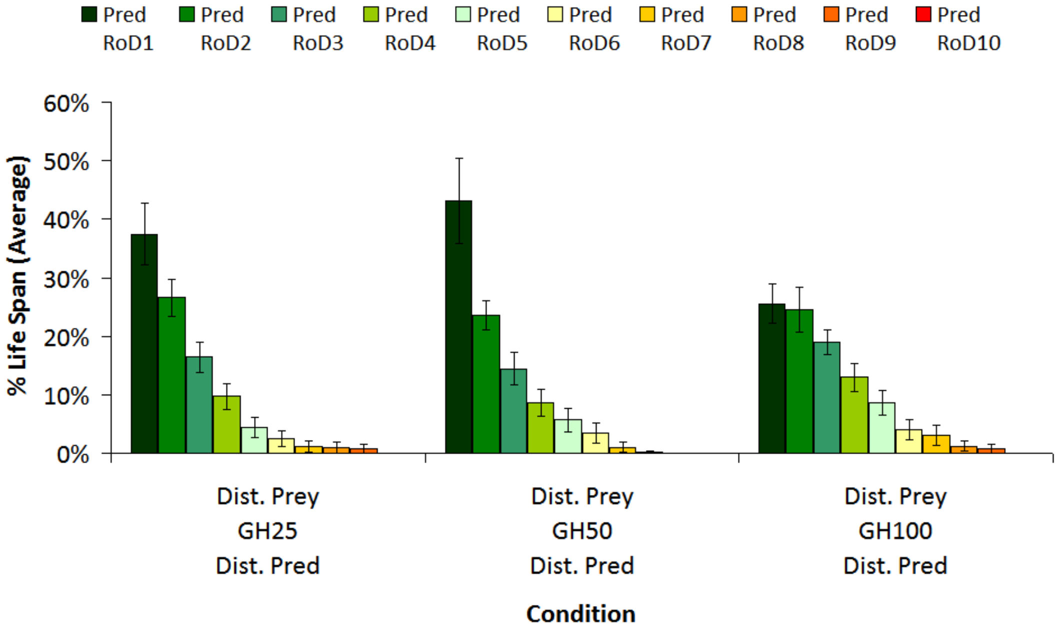

Figure 7.14: Summary of results showing Risk of Death (RoD) Management for our predator across conditions where distal prey versus distal predator and the hormone release of our emotion-based architecture follows a gradient release (the amount of hormone released being a function of the distance at which the predator is perceived). Each bar represents the average percentage of time our predator was in each RoD "zone". Error Bars show the Standard Error of the Mean (SEM). A RoD of 1 (green) means that the predator is almost if not exactly the most "healthy" it can be, whereas a RoD of 10 (red) means the predator is "near death". Here we see the prey with a release rate of 100 meant our predator spent a lower proportion of its time in the most "healthy" states possible. Not surprisingly, as shown in Figure 7.2 this was the condition with the highest average life span for the prey out of these conditions.

this be because it leads our prey to ignore all opportunities to feed in the presence of a resource to search for the nest (something which, depending on how long our prey takes to find the nest, can lead to death anyway) or because it causes our prey to inadvertently attract the predator's attention (as the infra-red ball is mounted on the "front" of our prey's body, there were instances in these experiments where the decision to "flee" to the nest by our prey led it to "attract" the attention of our predator — whereas, if it had stayed where it was, it might have survived). This latter might also indicate a need for further development to our architecture in terms of the actions selected. For example, in terms of introducing a "freezing" ability or more sophisticated way of assessing predator movement.

## 7.6.2 Activity Cycles and Action Patterns

### 7.6.2.1 States and State Transitions

Looking next at how this performance translates to a pattern of action, Figure 7.15 shows some averages for different types of state transition across conditions. Here we see that, on average, all of our prey demonstrated more opportunism than persistence.

Looking further, in terms of an activity profile, Figures 7.16 to 7.19 show the average time spent on each action per condition, as a percentage of overall life span. Together with Life Span, these graphs allow us to explore differences in terms of both the absolute number of time-steps that each type of prey spent on each activity, as well as the percentage of their "lives" that this made up. In terms of our "fixed" hormone release conditions, we can see that when our prey was up against our proximal predator, our prey spent more time resting or searching for its nest than anything else. This was also true when the same prey was up against our distal predator. However, in the latter conditions, we tend to see our prey spending a much greater proportion of its lifetime searching for its nest than actually resting within it. In contrast, the proportion of time divided between searching for the nest/resting in the former conditions appears more evenly split. This may be because our distal predator (due to its more active "hunting" ability) is naturally coming within the perceptual distance of our prey more often, thereby triggering hormone-release more often. Likewise, our "gradient" hormone release conditions show similar trends.

## 7.6.3 Brain-Body-Environment Maps

Figure 7.20 shows a selection of Brain-Body-Environment maps constructed from different conditions. Here, we observe similarities between our Distal Prey (G100) vs Prox Pred condition and Distal Prey (H25) vs Distal Pred condition.

Looking further at the movement of our different animats and time spent on different resources, Figures 7.22 to 7.25 show the average time spent (as a proportion of

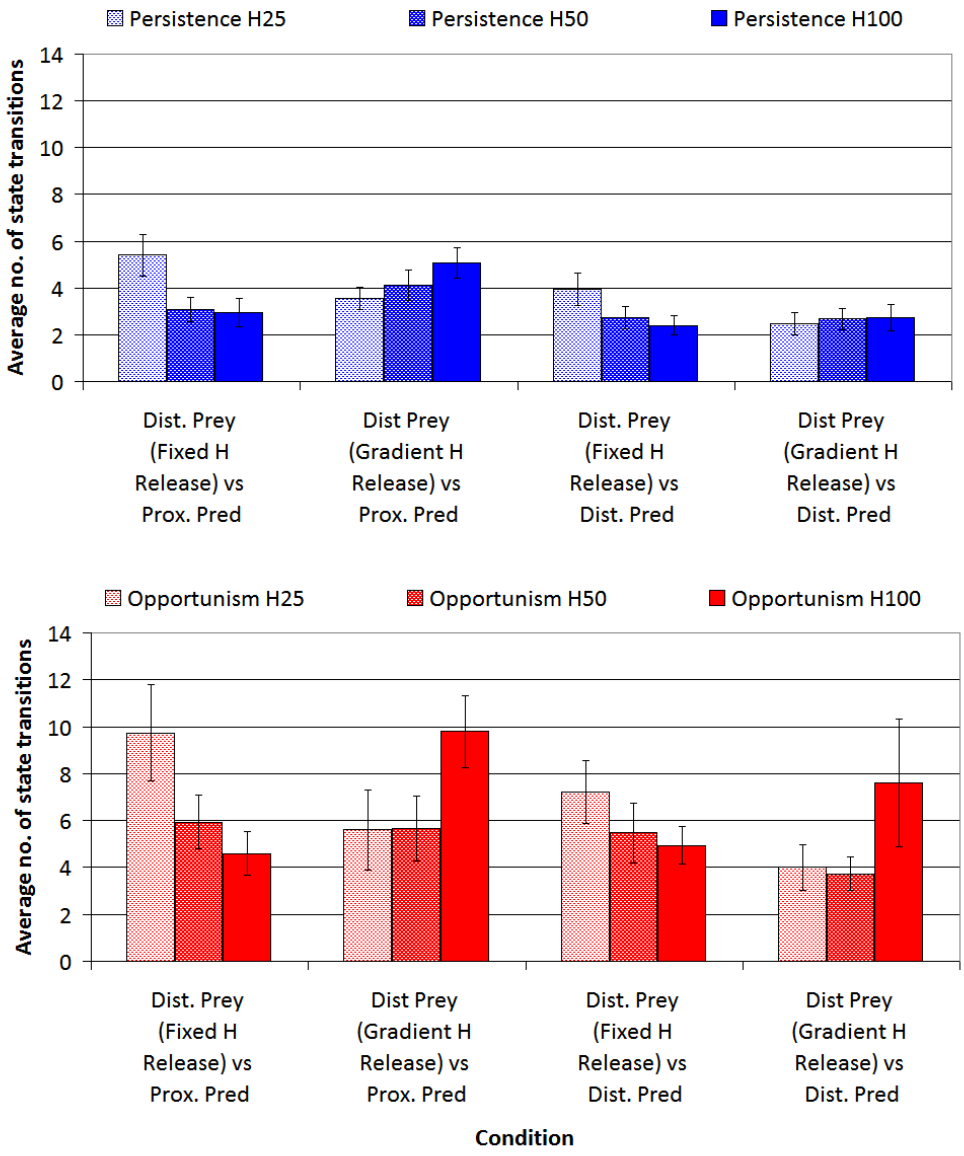

Figure 7.15: Summary of results showing the average number of different types of “state transitions” of our prey across conditions. Error bars show the Standard Error of the Mean (SEM). Exploring the idea of action selection as a cycle of activities, rather than a series of individual decisions, state transitions occur whenever our agent switches from one of the six available actions to another. Having constructed a state transition matrix for each individual run, here we define Opportunism as the number of times our prey consumed one resource while searching for another. Persistence represents the number of times our prey ended up consuming the resource it was searching for. Here we see that, on average, all of our prey demonstrate more opportunism than persistence.

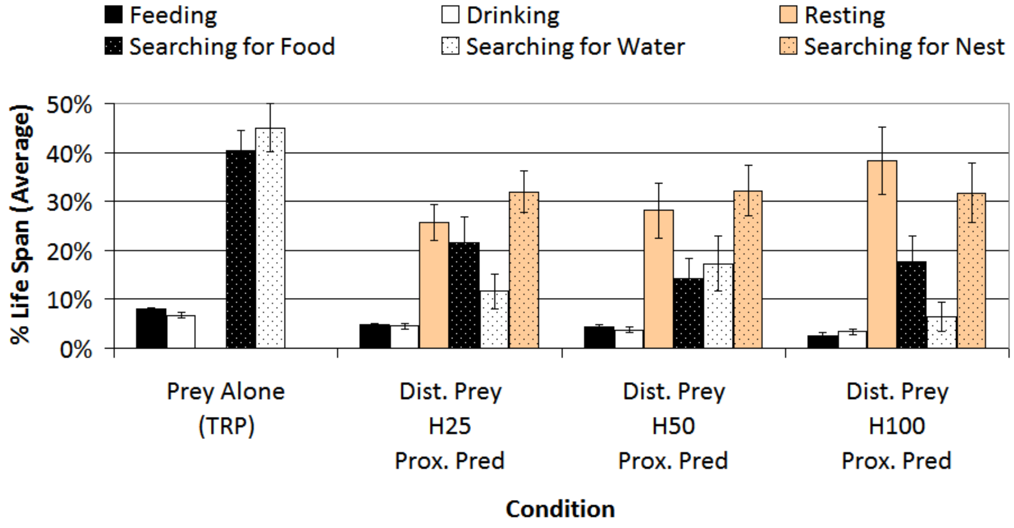

Figure 7.16: Summary of results in terms of "state sinks" across conditions where distal prey versus proximal predator and the hormone release of our emotion-based architecture is fixed. In this case, the average time (as a percentage of overall life span) spent by our different prey performing each possible activity/state out of the six available for our architecture to select from. Error Bars show the Standard Error of the Mean (SEM). As the activities that our emotion-based architecture is most likely to influence directly, we can focus our attention here on the proportion of time spent by our prey either resting or searching for the nest.

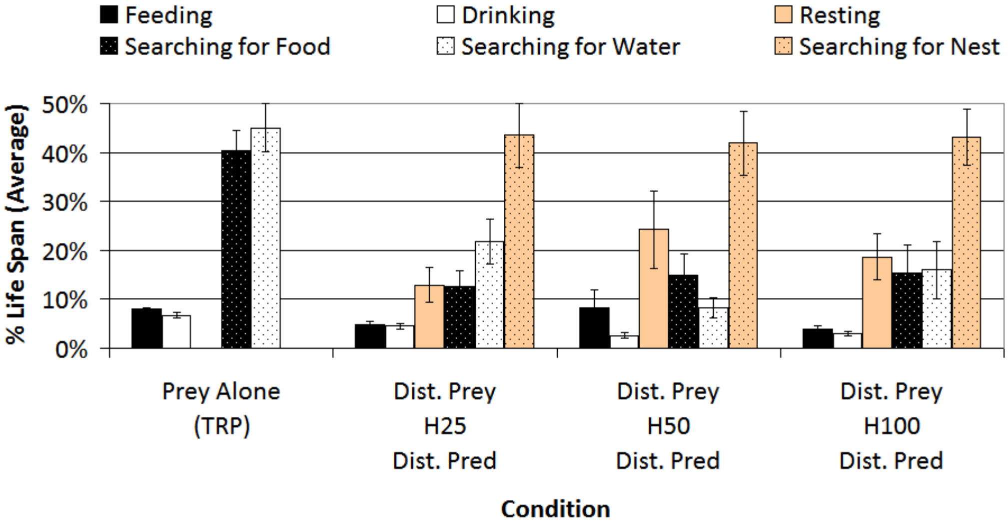

Figure 7.17: Summary of results in terms of "state sinks" across conditions where distal prey versus distal predator and the hormone release of our emotion-based architecture is fixed. In this case, the average time (as a percentage of overall life span) spent by our different prey performing each possible activity/state out of the six available for our architecture to select from. Error Bars show the Standard Error of the Mean (SEM). As the activities that our emotion-based architecture is most likely to influence directly, we can focus our attention here on the proportion of time spent by our prey either resting or searching for the nest.

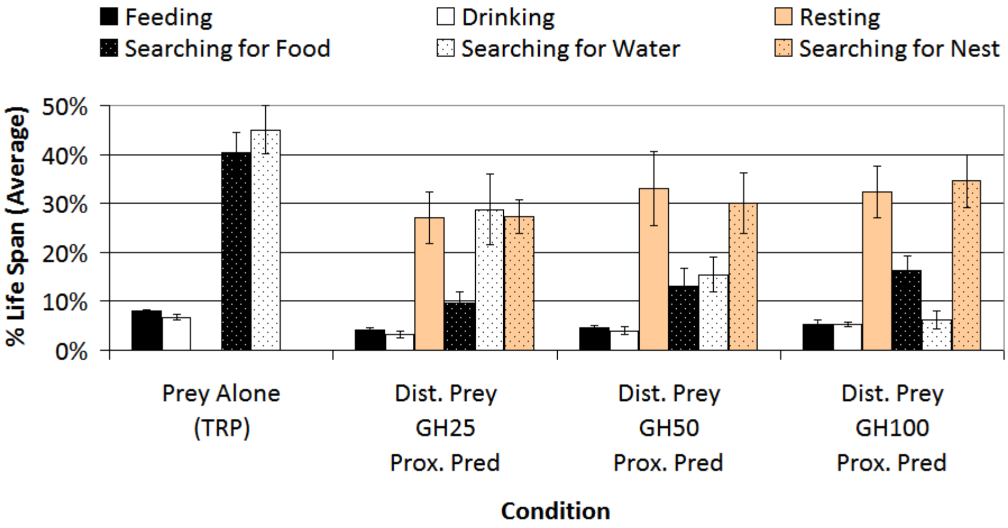

Figure 7.18: Summary of results in terms of "state sinks" across conditions where distal prey versus proximal predator and the hormone release of our emotion-based architecture follows a gradient release (the amount of hormone released being a function of the distance at which the predator is perceived). In this case, the average time (as a percentage of overall life span) spent by our different prey performing each possible activity/state out of the six available for our architecture to select from. Error Bars show the Standard Error of the Mean (SEM). As the activities that our emotion-based architecture is most likely to influence directly, we can focus our attention here on the proportion of time spent by our prey either resting or searching for the nest.

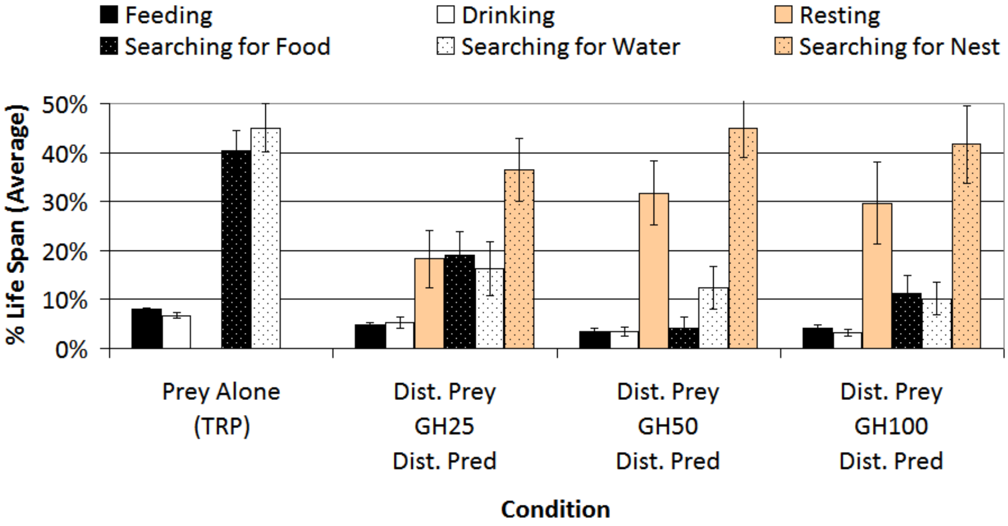

Figure 7.19: Summary of results in terms of “state sinks” across conditions where distal prey versus distal predator and the hormone release of our emotion-based architecture follows a gradient release (the amount of hormone released being a function of the distance at which the predator is perceived). In this case, the average time (as a percentage of overall life span) spent by our different prey performing each possible activity/state out of the six available for our architecture to select from. Error Bars show the Standard Error of the Mean (SEM). As the activities that our emotion-based architecture is most likely to influence directly, we can focus our attention here on the proportion of time spent by our prey either resting or searching for the nest.

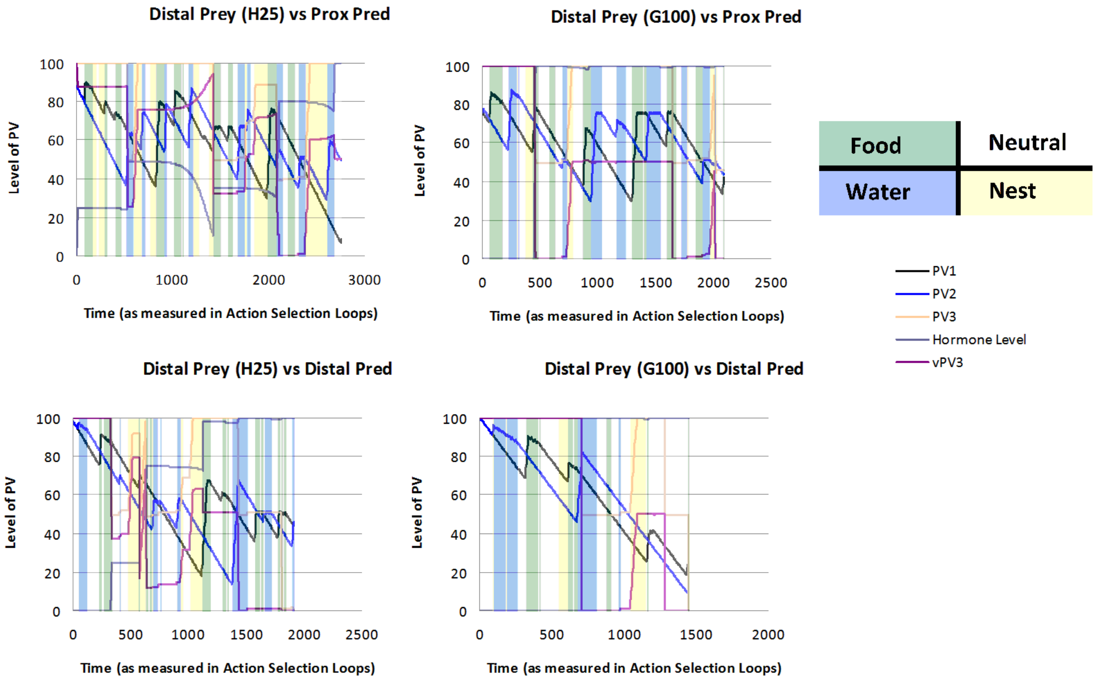

Figure 7.20: Illustration of differences (relating to aspects of prey Brain, Body and Environment over time) as observed across conditions between individual runs and shown by our BBE Maps. The background bars here represent our prey's position within its environment (green bars show where our prey detected "food", blue bars show "water", the absence of any bar indicates the prey detected no resources present in its immediate environment at that time and yellow bars indicate the prey was in its "nest") while each line represents the internal variables of our prey, including each of its physiological variables (black for blood sugar, blue for vascular volume, yellow for integrity) and hormone levels/effects. Additionally, here vPV3 represents the prey's *perceived* level of integrity, which in the modulated conditions differed from the *actual* level of integrity when concentrations of hormone were released.

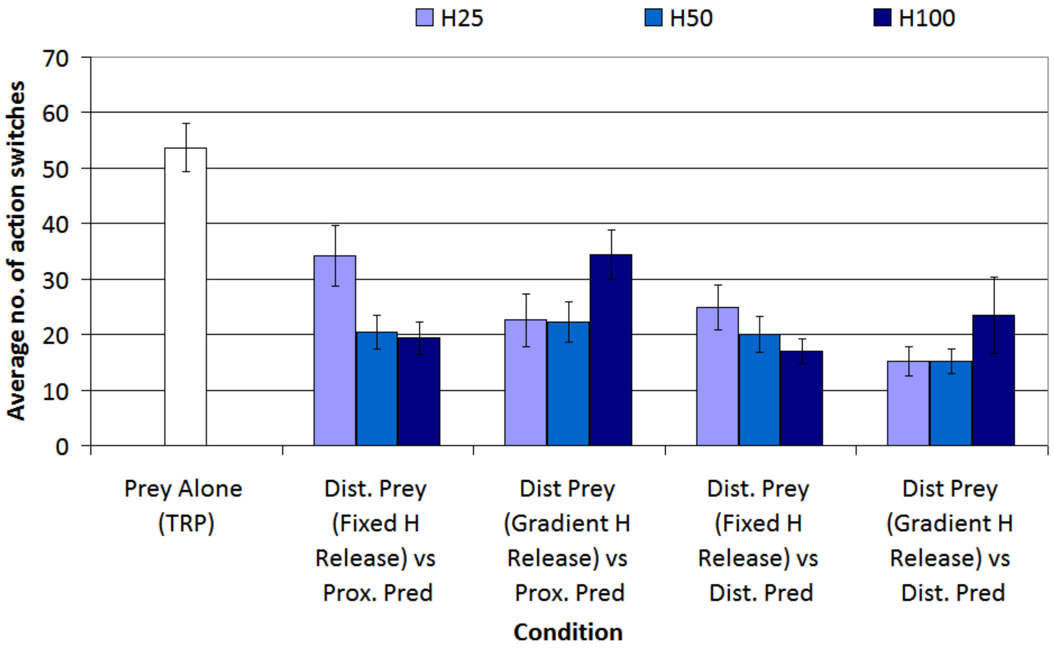

Figure 7.21: Summary of results showing the average number of “action switches” of our prey across conditions. Error bars show the Standard Error of the Mean (SEM). Exploring the idea of action selection as a cycle of activities, rather than a series of individual decisions, we use this simple measure to compare our prey. An “action switch” occurs whenever our prey makes the decision to switch from one of the six actions available to it to another. Here we can see that, on average, for our prey with a fixed hormone release a hormone release of 25 led to a greater number of action switches. Conversely, for our prey with a gradient hormone release a hormone release of 100 led to a greater number of switches.

their overall life span) by our prey on different resources/areas in the arena. Compared to the results of previous chapters, we see that the prey we are looking at in these conditions are consistently spending a greater proportion of their time in the nest.

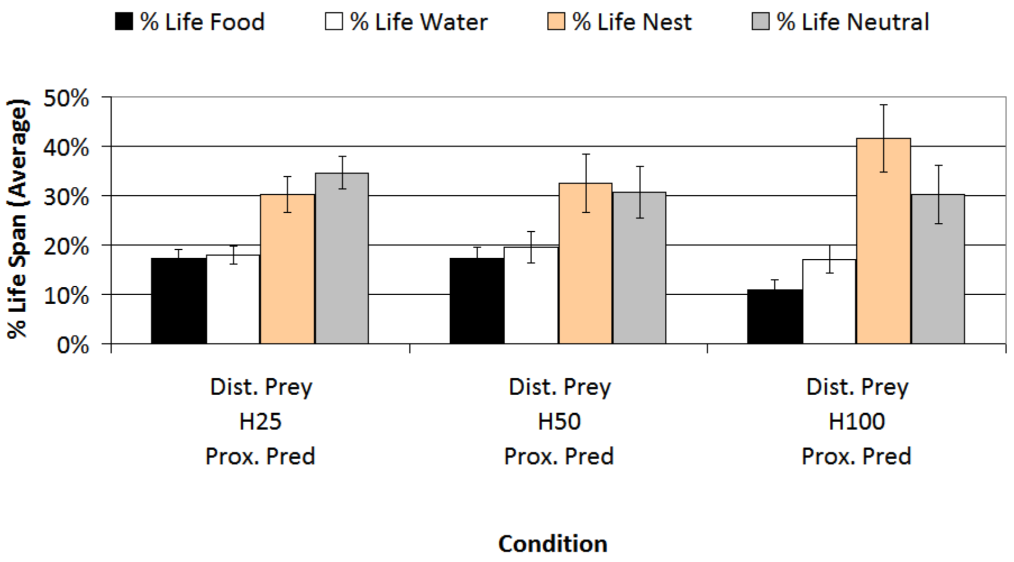

Figure 7.22: Summary of results showing time (as the average proportion of our prey's life span) spent on different resources/areas in its environment/arena by our prey across conditions where distal prey versus proximal predator and the hormone release of our emotion-based architecture is fixed. Error Bars show the Standard Error of the Mean (SEM).

## 7.7 II. Fixed versus Gradient Response for Prey (Varying Motor Response)

In designing the second set of experiments reported in this chapter, we once again took our inspiration from ideas surrounding the fight-or-flight response in the biological sciences. Since its initial christening as the fight-or-flight response, researchers of this phenomenon have also referred to it by many other names, reflecting the knowledge collected about this response. Among its aliases are the fight-flight-or-freeze response, hyperarousal and acute stress response. As a theory, it states that animals react to threats with a general discharge of the sympathetic nervous system, priming the animal for fighting or fleeing. This response, first described by Walter Bradford Cannon [110; 111] was later recognised as the first stage of a general adaptation syndrome regulating stress responses in many organisms.

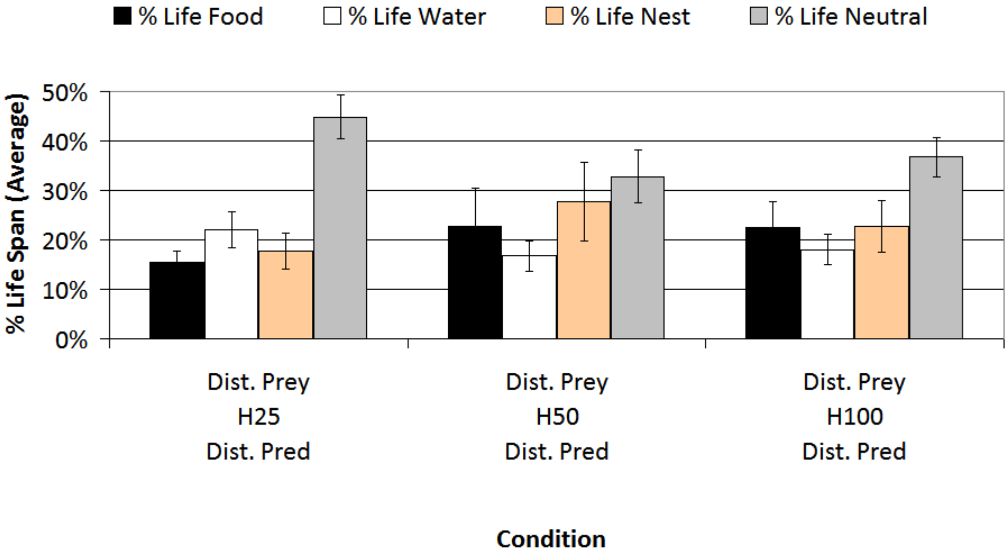

Figure 7.23: Summary of results showing time (as the average proportion of our prey's life span) spent on different resources/areas in its environment/arena by our prey across conditions where distal prey versus distal predator and the hormone release of our emotion-based architecture is fixed. Error Bars show the Standard Error of the Mean (SEM).

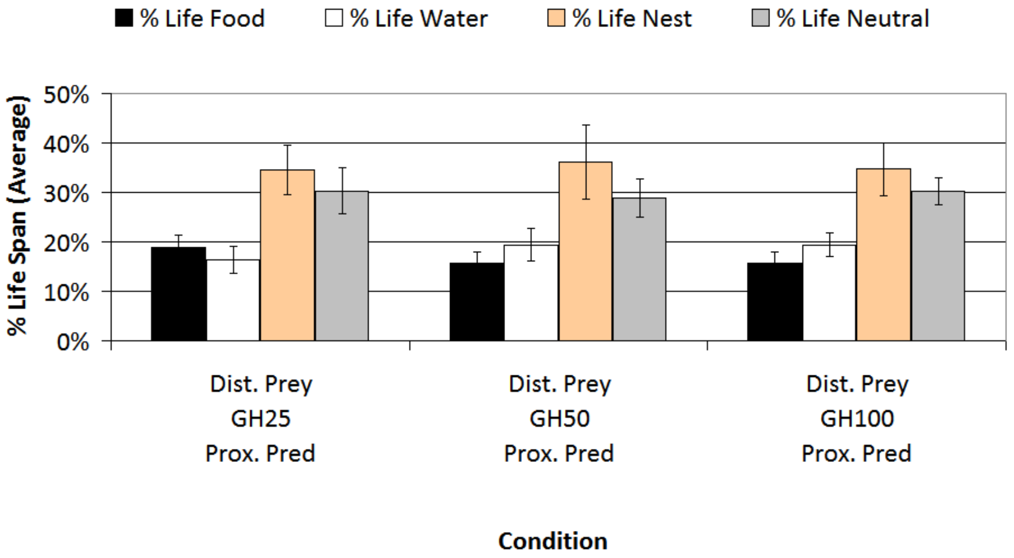

Figure 7.24: Summary of results showing time (as the average proportion of our prey's life span) spent on different resources/areas in its environment/arena by our prey across conditions where distal prey versus proximal predator and the hormone release of our emotion-based architecture follows a gradient release (the amount of hormone released being a function of the distance at which the predator is perceived). Error Bars show the Standard Error of the Mean (SEM).

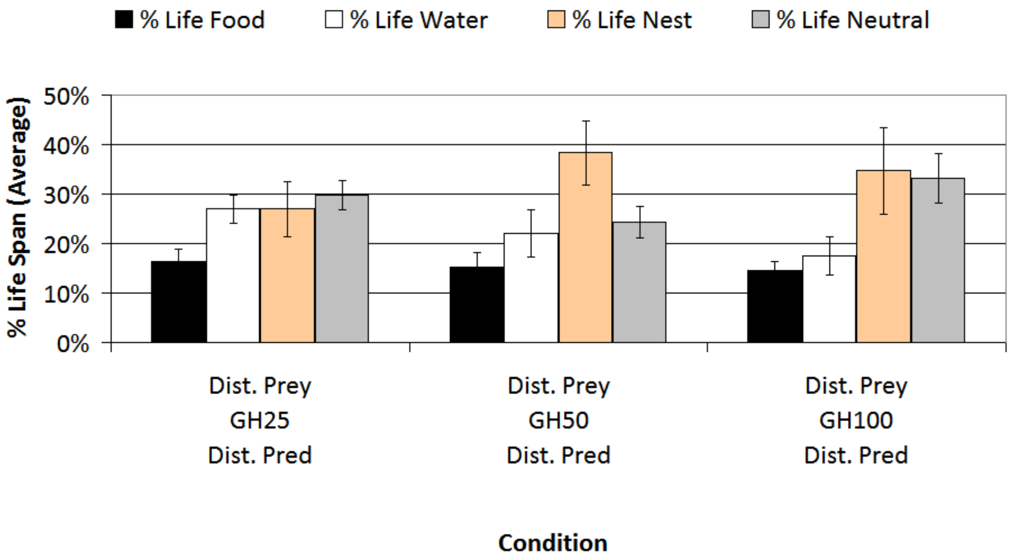

Figure 7.25: Summary of results showing time (as the average proportion of our prey's life span) spent on different resources/areas in its environment/arena by our prey across conditions where distal prey versus distal predator and the hormone release of our emotion-based architecture follows a gradient release (the amount of hormone released being a function of the distance at which the predator is perceived). Error Bars show the Standard Error of the Mean (SEM).

| 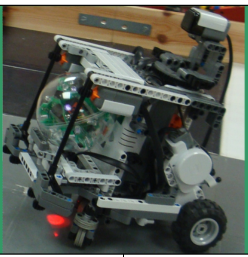 | | *Variations in Prey “Brain-Body”*<br>*[Emotion-based architecture]* | |
|---|---|---|---|
| | | **Fixed Response: Hormone affects action selection only (speed stays constant)** | **Gradient Response: Hormone affects action selection *and* intensity (“speedy prey”)** |
| *Variations in Prey "Body" [Visual Field]* | **Small Visual Field** | Small Visual Field Fixed H Response | Small Visual Field Gradient H Response |
| | **Large Visual Field** | Large Visual Field Fixed H Response | Large Visual Field Gradient H Response |

Figure 7.26: Summary of experimental conditions: In these experiments a factorial 2x2 design (resulting in four experimental conditions) was used to investigate how modifications to our existing emotion-based architecture, namely allowing our hormone-like mechanism to influence the *intensity* (speed) of prey response as well as the action selected, will interact with the perceptual abilities (body) of our prey so as to generate more adaptive behaviour.

Taking inspiration from this theory for improving our prey's performance, in our final experiments, we were particularly interested in the idea that we could use our hormone-like mechanism in a similar way to "prepare" our animat prey. That is, in terms of our emotion-based architecture, to allow the hormone-like mechanism to react and adapt to threats not just by changing perception, but also by change the actions selected in a similarly physical way (that is, similar to the observed changes in the sympathetic nervous system, priming the animal for fight or flight). This thereby has the added benefit of allowing us to study the importance of other aspects of "body" using our H3RP. (Here, we look at the importance of our prey's visual field.) But simultaneously, it allows us to see how we might integrate our hormone-like mechanism with other components of the underlying motivation-based architecture. Perhaps more to the point, it generates a situation where we can consider the possible costs and benefits of such integration.

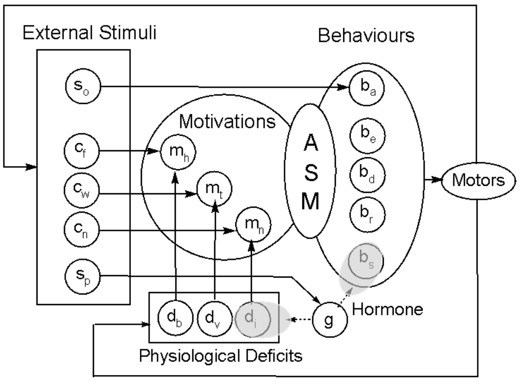

Figure 7.27: Overview of our developed architecture ("brain") for a prey agent: internal "body" is represented through physiological variables, deficits of which act as drives which, combined with the presence/absence of external stimuli, are used to calculate motivational and behavioural intensity. For example, calculations of motivational intensity for a motivation representing hunger will take into account both physiological deficits such as blood sugar and the presence/absence of food in the environment. In our experiments we vary external "body" using different physical sensors. Emotion-like states are modelled by the addition of a gland (g); releasing a "hormone" in the presence of a specific stimulus (in this case the predator) which affects both perception of internal physiological deficits, increasing calculations of motivational intensity, and the behaviour selected in terms of physical response (speed or tempo of behaviour is increased if hormone is present)

## 7.8 Prey

### 7.8.1 Interoceptive Modulation (Graded Hormone Release and Response Intensity)

To create four different experimental conditions, we first created a new (distal) prey by changing the ultrasonic sensor of our prey to an infra-red seeker (matching its sensor-type to that of our distal predator). Following this, we varied the actions of our architecture so as to incorporate infra-red avoidance into our prey, mounting an infra-red ball on our predator. We then created a second prey by varying the interoceptive modulation of our emotion-based architecture. This was achieved by connecting the hormonal concentration of our prey to calculations of motor speed, triggered by the presence of the predator detected via the infra-red detector. That is, the speed of this prey when performing each action was a function of the level of total hormone concentration at any given time. We also changed the decay rate of our mechanism back to the original exponential one (having tested both in trials beforehand so as to see which might be the better one for these experiments). In this way, we created a modification of our emotion-based architecture, further illustrated in Figure 7.27.

### 7.8.2 Visual Field (Small/Large)

However, being interested in thinking about the possible functions of our hormone-like mechanism in terms of attention, and still interested in other aspects of body for their importance, in these experiments we also chose to vary another aspect of our prey's embodiment. This time, we selected what might be thought of as our prey's visual/perceptual field, using our new prey's infra-red seeker to vary this either to allow hormone to be released via signals from the sides of the prey. That is, to include peripheral vision (using the full visual field/abilities of the infra-red sensor) or giving our prey something more like "tunnel vision"; where hormone is likely to be released only when the predator is detected in front of the prey.

## 7.9 Predator

The predator used in these experiments was the Distal Predator described in the preceding chapters. However, one additional change was incorporated in its design in as much as we mounted an infra-red ball on top of this animat (see Figure 7.28).

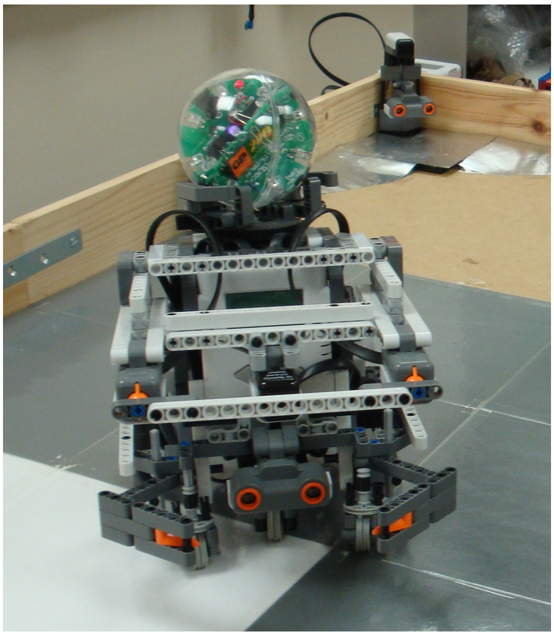

Figure 7.28: Illustration of our modified predator, on which we mounted an infra-red ball to enable our modified prey (now equipped with an infra-red sensor) to detect its presence.

## 7.10 A Robotic Nest

Finally, another change made to the basic arena in these experiments was the addition of a more “dynamic” nest resource. While the other resources in our H3RP were still represented in the same way, we used a LEGO NXT to effectively act as a deterrent to the predator (making it less likely to enter the nest) and at the same time, allow the prey to (via Bluetooth) find and enter the nest from a short distance away (see Figure 7.29).

## 7.11 Results

For each run, both predator and prey were tested following the standard procedures used in the previous chapters. Four experimental conditions were created by varying both our prey’s interoceptive modulation (in terms of fixed/graded response) and visual field (small/large). For each of these conditions, the data from five runs were collected testing each condition (a total of twenty runs). The reason for this sample size was because we also wanted to study our H3RP from yet another different perspective. After each run, data again collected from both prey and predator robots with regards to sensory and internal (physiological) data. However, in this set of experiments we

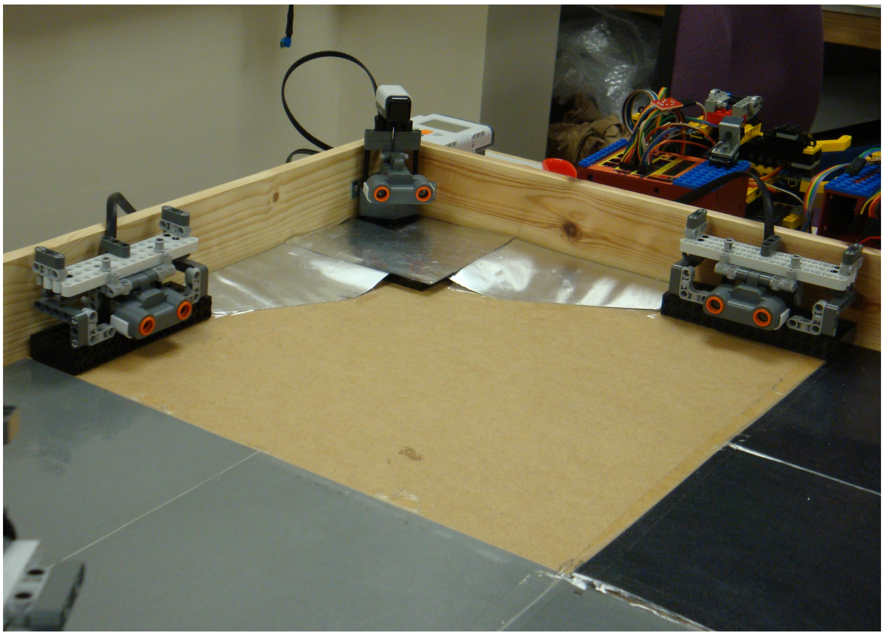

Figure 7.29: Our robotic nest: designed to add a further dynamic to the environment of our developed animats. This consisted of a LEGO NXT, with ultrasonic sensors fitted to deter the predator from the area, as well as to allow the prey to detect and “home in” to its nest from a (small) distance away.

wanted to compare data from an external source. That is, via observation of our animats as though they were animals. Thus, we collected additional data by recording our animats for each run (mounting a video camera above the arena). We use this data here so as to present findings from an alternative perspective. In addition to which, we also present the results of a fifth condition alongside our four main experimental conditions. In this condition (Large Visual Field Gradient H Response with Cost) we introduce a "cost" of increasing speed for our prey (where the decrease of vascular volume and blood sugar increases where our prey goes fast) allowing us to consider whether putting a "price" on this action will negate or otherwise affect its adaptive value.

### 7.11.1 Measures of Viability

From the data collected, we first calculated and looked at the results in terms of our chosen measures of viability. These included calculations of Life Span, Overall Comfort and Physiological Balance as well as calculations of our prey's Risk of Death Management. Figure 7.30 shows how each type of prey performed as measured by our three measures of viability of Life Span, Overall Comfort and Physiological Balance.

#### 7.11.1.1 Life Span, Overall Comfort and Physiological Balance

In terms of Life Span, at first glance Figure 7.30 appears to show the most noticeable difference between the condition which manipulated both prey perceptual field and the hormone-like mechanism's connection to action (gradient response). Again, similar trends are seen between the average Life Span calculated in time-steps and that measured in minutes. From this, it seems that the gradient response conditions may result, on average, in longer Life Spans.

Performing a multi-way ANOVA (2x2) on our data, we looked for main effects of our independent variables of both prey interoceptive modulation (fixed/graded response) and visual field (small/large). In terms of Life Span as calculated in time-steps, after arcsine transformation we find a significant main effect for prey visual field (p=0.030) but not interoceptive modulation (p=0.181). No interaction between the two variables was suggested (p=0.190). However, the results in terms of real-time Life Span (measured in minutes) showed both an effect for prey visual field (p=0.006) and hormone response (p=0.034). Similarly to the previous measure, no interaction was found between the two variables. Turning to the results for Overall Comfort and Physiological Balance, the trends for these measures seem to differ from that of Life Span. For both these measures, it appears that the "gradient-response" or "speedy prey" conditions resulted in higher average levels of these values than their "fixed-response" counterparts. This may reflect a "penalty", incurred on these measures for the apparently increased Life Span for these conditions. However, due to the hunting abilities of

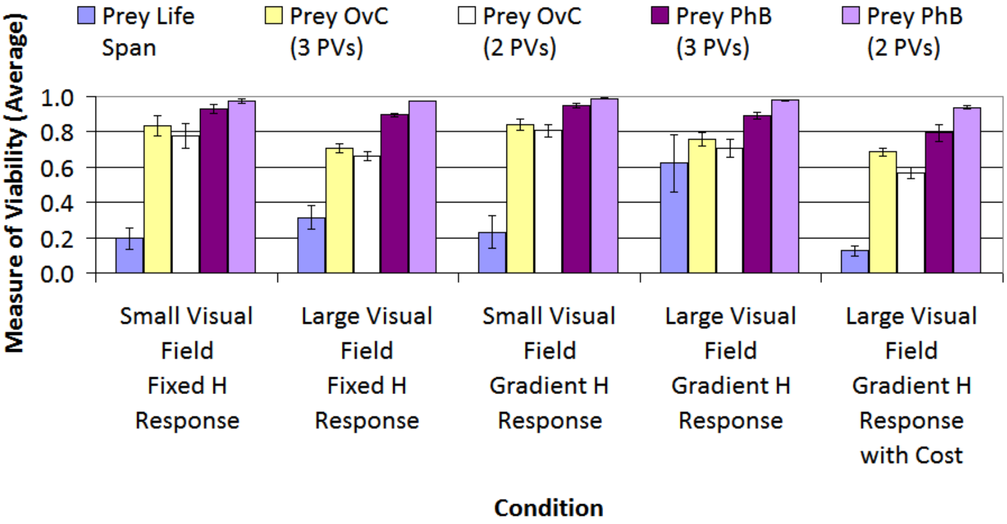

Figure 7.30: Summary of results showing average Life Span, Overall Comfort and Physiological Balance of our prey across conditions. Error Bars show the Standard Error of the Mean (SEM).

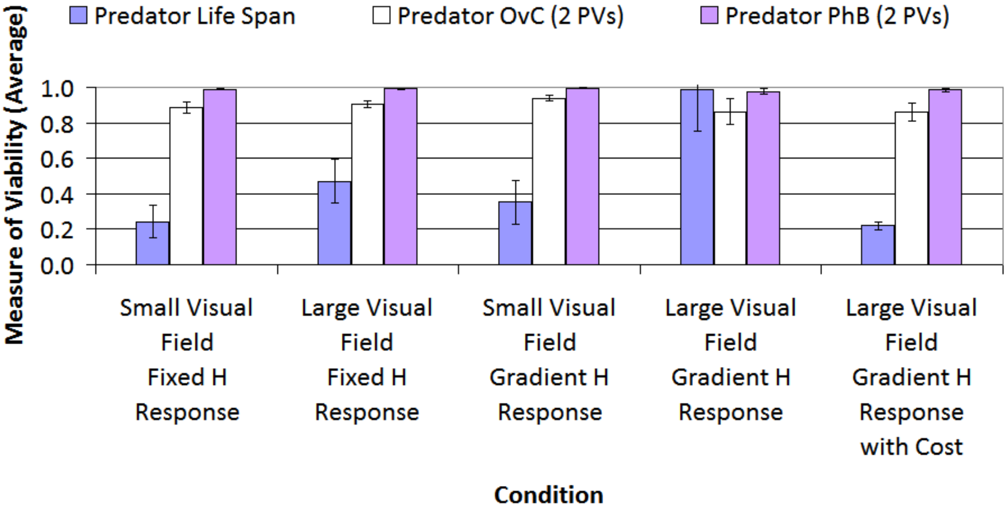

Figure 7.31: Summary of results showing average Life Span, Overall Comfort and Physiological Balance of our predator across conditions. Error Bars show the Standard Error of the Mean (SEM).

the predator, and taken in conjunction with our observations it seems more likely that the other conditions were those in which the predator had a perceptual "advantage" and literally found the agents in these conditions "easy prey". Having performed separate multi-way ANOVAs for these measures, looking at the results for each of these together, our tests found a significant main effect of visual field on Overall Comfort (p=0.017) and Physiological Balance (p=0.016). However, no main effects were found for our "gradient-response" conditions.

#### 7.11.1.2 Risk of Death Management

With regards to Risk of Death (RoD) Figure 7.32 shows these results for each condition. As before, each bar represents the average percentage of time our prey was in each "zone". What is interesting about these results, complementing the information given by our previous measures, is that, across the conditions for our "speedy prey" (that is, the ones which we implemented and tested with our modified architecture) we see similar trends across the runs within these "speedy" conditions: the majority of these agents' lives seeming to have been spent in the sixth zone. To see how this relates to our predator's experience, Figure 7.33 shows calculations of Risk of Death for our predator.

### 7.11.2 Activity Cycles and Action Patterns

#### 7.11.2.1 States and State Transitions

Figure 6.11 shows some averages for different types of state transition across conditions. Here we see that, on average, all of our prey continue to demonstrate more opportunism than persistence. The prey with large visual field and gradient hormone response more often demonstrated persistence in searching for and successfully reaching the nest. Here and in our observations from having constructed state transition diagrams from individual runs, we see more similarities between our prey with the same visual field (those prey with larger visual fields demonstrating more connections between states).

Looking further, in terms of an activity profile, Figures 7.35 and 7.36 show the average time spent on each action for each animat. Figure 7.35 shows differences in terms of the average percentage of their "lives" spent on each activity by the prey, while Figure 7.36 shows the absolute number of time-steps that the predator spent on each activity. From these graphs, we again see the greatest similarities between those prey with large visual fields. However, comparing the absolute number of time-steps the predator spent hunting the prey for each of these conditions, the prey with a gradient response is seen to have lived to be pursued for longer.

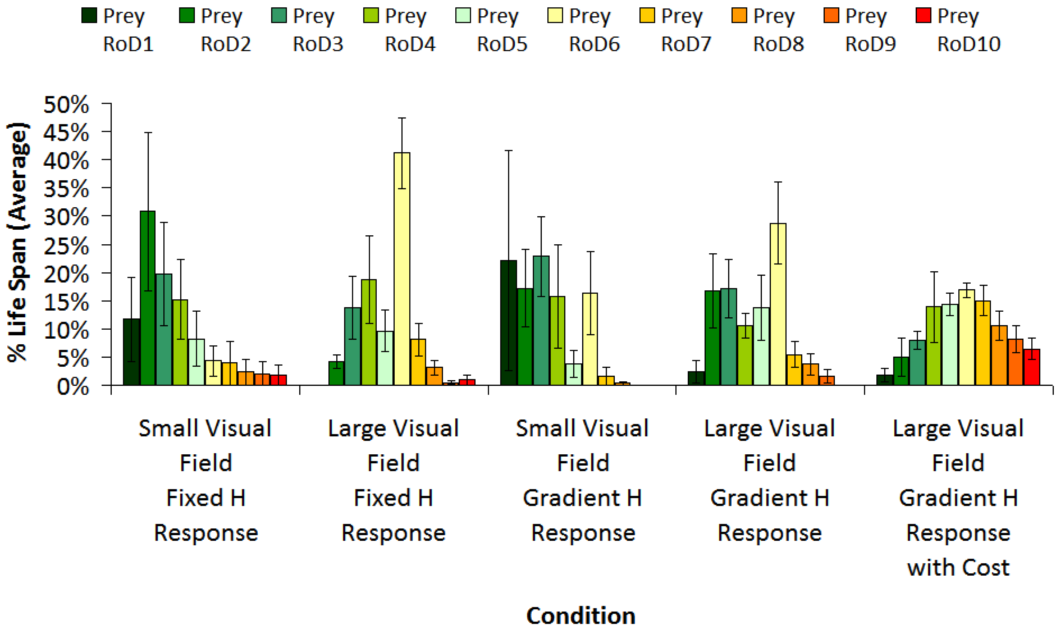

Figure 7.32: Summary of results showing Risk of Death (RoD) Management across conditions for our prey. Each bar represents the average percentage of time our prey was in each RoD "zone". Error Bars show the Standard Error of the Mean (SEM). A RoD of 1 (green) means that the prey is almost if not exactly the most "healthy" it can be, whereas a RoD of 10 (red) means the prey is "near death".

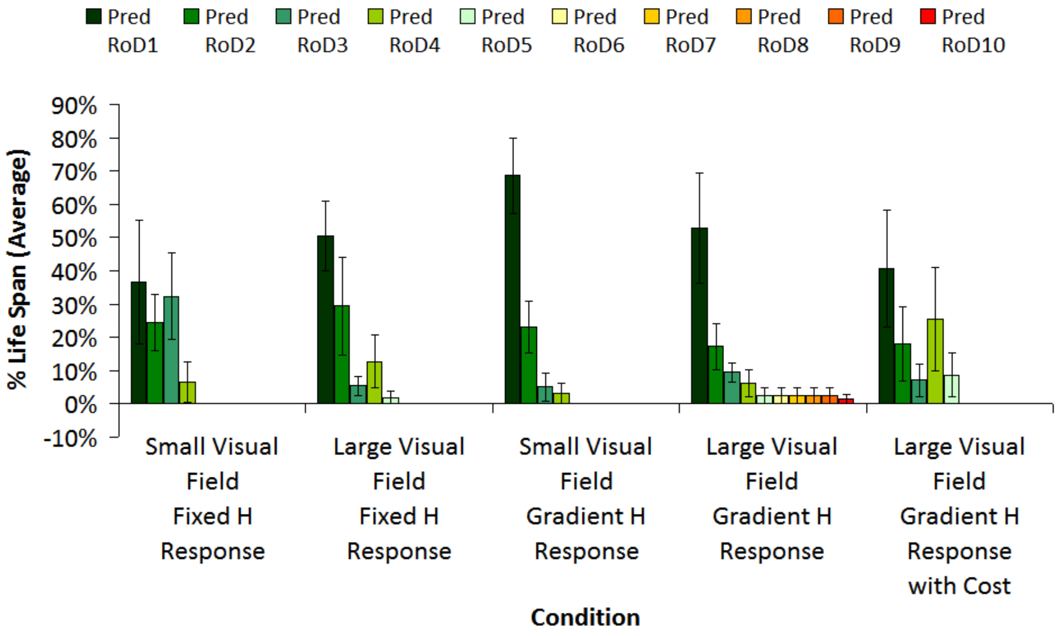

Figure 7.33: Summary of results showing Risk of Death (RoD) Management across conditions for our Predator Animat. Each bar represents the average percentage of time our predator was in each RoD "zone". Error Bars show the Standard Error of the Mean (SEM). A RoD of 1 (green) means that the predator is almost if not exactly the most "healthy" it can be, whereas a RoD of 10 (red) means the predator is "near death".

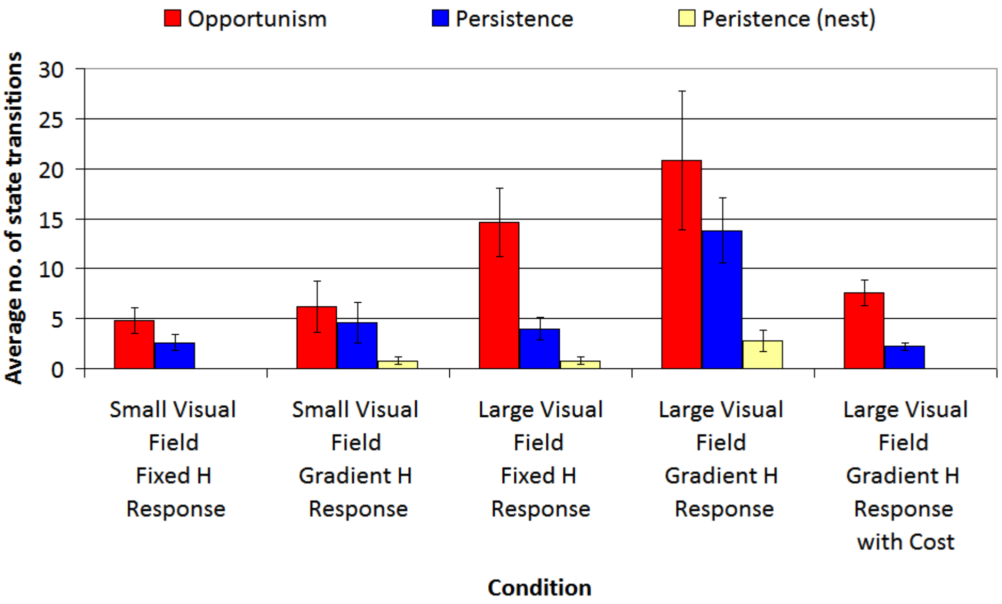

Figure 7.34: Summary of results showing the average number of different types of "state transitions" of our prey across conditions. Error bars show the Standard Error of the Mean (SEM). Exploring the idea of action selection as a cycle of activities, rather than a series of individual decisions, state transitions occur whenever our agent switches from one of the six available actions to another. Having constructed a state transition matrix for each individual run, here we define Opportunism as the number of times our prey consumed one resource while searching for another. Persistence represents the number of times our prey ended up consuming the resource it was searching for and we add a further category to this, Persistence (nest), which is actually a subset of Persistence (focusing on those instances where our prey both searched for the nest and ended up resting in its nest). Here we see that, on average, all of our prey demonstrate more opportunism than persistence. The prey with large visual field and gradient hormone response more often demonstrated persistence in searching for and successfully reaching the nest.

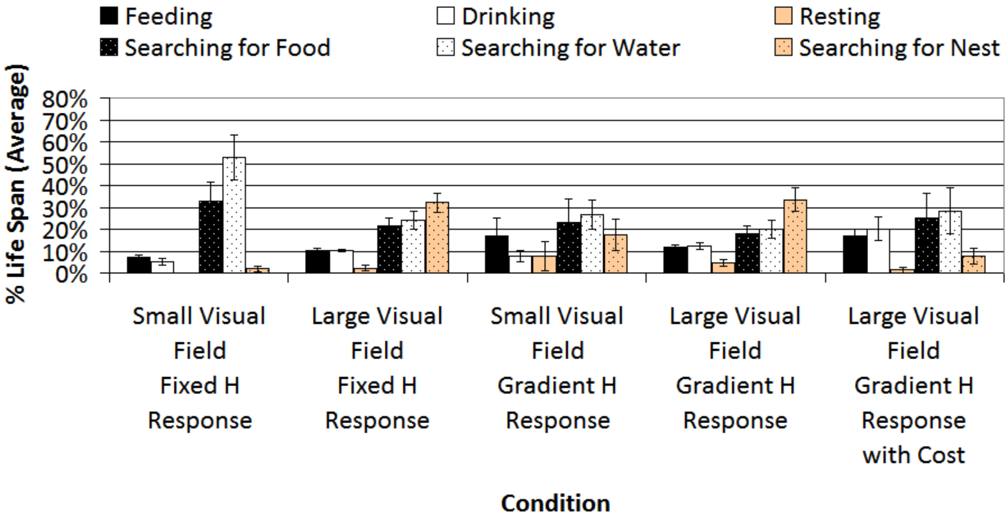

Figure 7.35: Summary of results in terms of "state sinks"". In this case, the average time (as a percentage of overall life span) spent by our different prey performing each possible activity/state out of the six available for our architecture to select from. Error Bars show the Standard Error of the Mean (SEM). As the activities that our emotion-based architecture is most likely to influence directly, we can focus our attention here on the proportion of time spent by our prey either resting or searching for the nest.

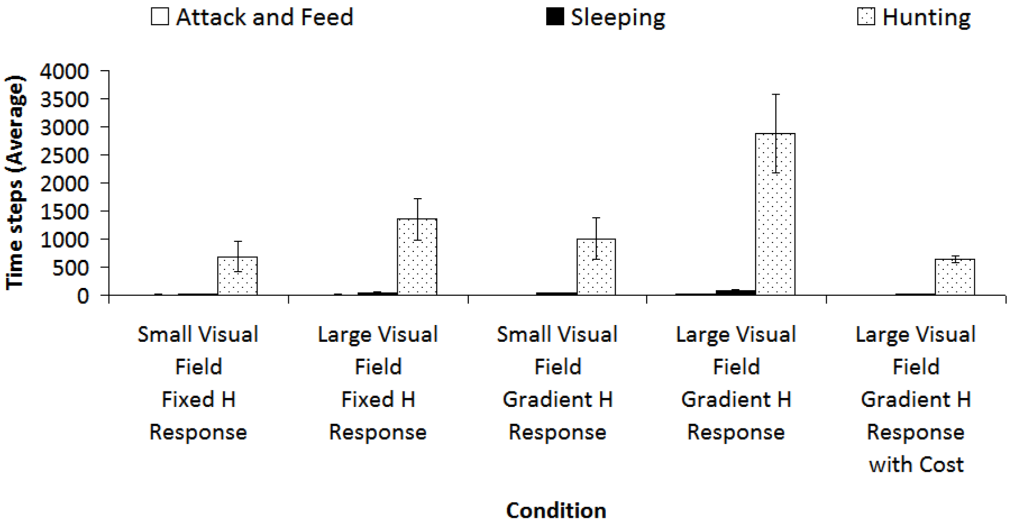

Figure 7.36: Summary of results in terms of "state sinks". In this case, the average time (in terms of action selection loops) spent by our predator performing each possible activity/state out of the three available for its architecture to select from. Error Bars show the Standard Error of the Mean (SEM).

To complement this data, and turning now to consider some of the results of our observational analyses, Figure 7.37 allows us to compare our models of state-transitions constructed from our internally-collected data with those constructed from externally-collected data. Using a structural ethogram to record our robots' behaviour over time, the idea was that this would also allow us to consider how ethologists might classify our robots' behaviour. One of the reasons for doing this, for example, was to see whether we might be able to identify fixed action patterns between/within conditions. In this case, by describing our agent's performance at a more abstract level, it provides us with a neat summary of our most successful prey's action patterns[1].

### 7.11.3 Brain-Body-Environment Maps

Figures 7.38 to 7.40 show selected runs from our experimental conditions in the form of our Brain-Body-Environment maps. Figure 7.38 shows these across our four main experimental conditions for our prey, while Figure 7.40 shows these across our four main experimental conditions for our predator. Figure 7.41 presents a sample map

---

[1] In tests of inter-observer reliability, when taking into account expected agreements by chance, from the resulting value of kappa, the strength of agreement between observers was judged to be good/substantial (between 0.61 and 0.8).

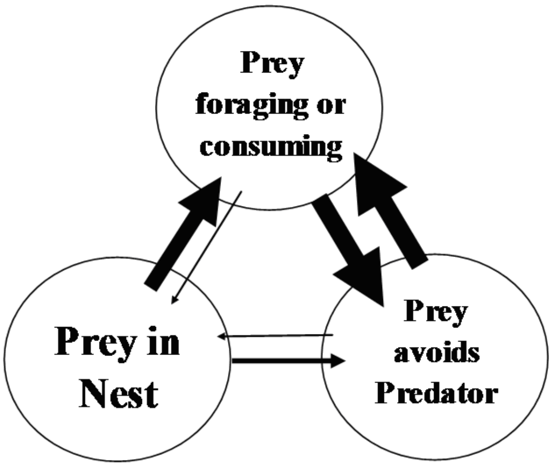

Figure 7.37: Summary of results from *external* observation of our prey animat. Here, we have constructed a Markov state space diagram using the results of observational analysis for our modified architecture (constructed from an example run using our most successful prey — full visual field, gradient hormone response). Specifically, this has further enabled us to compare the models constructed from data obtained both internally and externally from our implemented system. That is, to explore differences between recorded *actions selected* by our animats (internal state transitions) to the *behaviour observed.* In this case, this has also enabled us to identify more general patterns of behaviour, allowing us to group actions together and describe the *behaviour* observed using a method much more in tune with those used by traditional ethologists.

for both animats from our “additional” condition, where we can see the effects of introducing a physiological “cost” of speeding up movement through the environment (physiological variables now being seen to decrease with physical effort as well as time, creating a further *dynamic* and more intricately entwining elements of brain, body and environment).

As shown, we see the “success” of our modified architecture is dramatically improved when given to a prey embodied with a large visual field, allowing the emergence of a more “stable” cycle of activity. Within our Large Visual Field Gradient H Response condition, this also greatly affected the emergent behaviour observed (becoming much more “emotion-like” to us as observers, reminiscent of the “flight” response noted in nature). To provide more of an overview of this data, Figure 7.39 shows the average number of state transitions or “action switches” of our prey across conditions. Here we can see that, on average, our prey with both a large visual field and gradient hormone response performed the greatest number of action switches.

Looking further at the movement of our different animats and time spent on different resources, Figure 7.42 shows the average time spent (as a proportion of their overall life span) on different resources/areas in the arena. As before, we see greater similarities in distribution between those animats with the same visual fields.

## 7.12 Discussion and Conclusion

In this chapter we have reported the results of two sets of experiments, both investigating how our emotion-based architecture might benefit from (and have a role in) sensory integration. The results of our first set of experiments (Figures 7.2 and 7.20 providing the majority of our evidence here) suggest a role for the body in producing a gradient response from our prey. Again, supporting the results of our previous chapters, this suggests a possible emergent function or property of our emotion-based architecture in terms of risk assessment.

The results of our second set of experiments also suggest a role for the body in producing a gradient response from our prey (Figures 7.30, 7.34, 7.38 and 7.39 supporting this most). However, these results seem to go further in order to indicate a possible function of our emotion-based architecture in terms of resource allocation. For, although not simulated in our variation of the H3RP, in the real world there is usually a cost associated with physical exertion such as speed. Our emotion-based architecture might have a role in such situations. Indeed, this shows the potential of our architecture for future robotic research (one of the goals of many researchers being to design agents that are able to manage/allocate their own finite resources — a good example being that of battery power).

Combining the results of our experiments, what is also interesting is what our latter experiments appear to suggest about the former. A gradient hormone-release for our

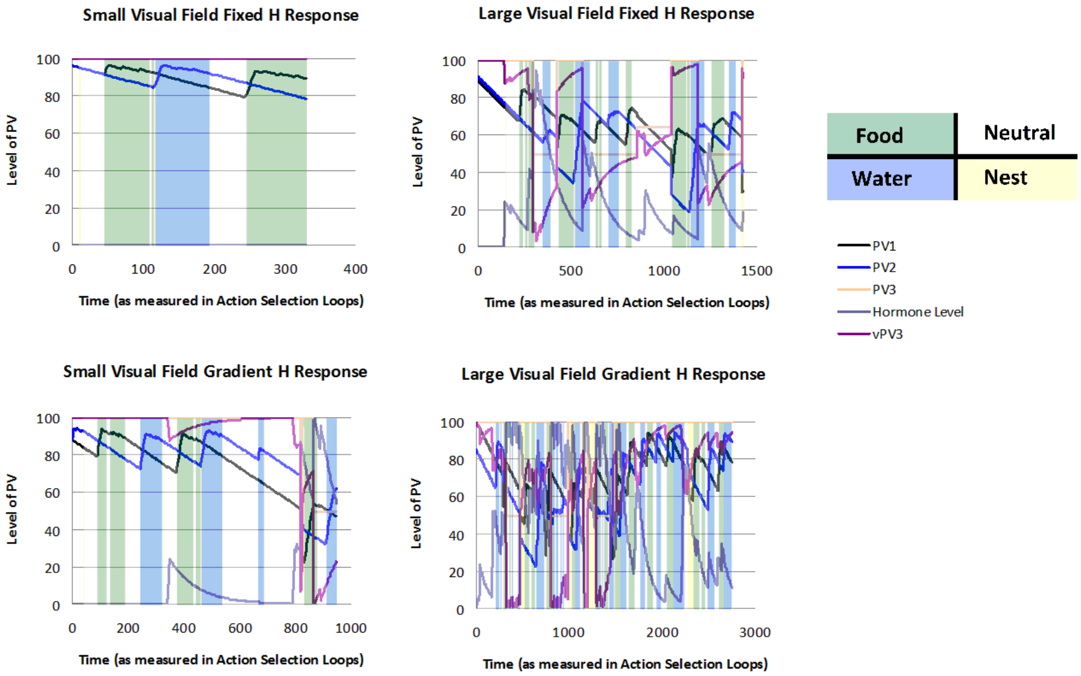

Figure 7.38: Illustration of differences (relating to aspects of prey Brain, Body and Environment over time) as observed across conditions between individual runs and shown by our BBE Maps. The background bars here represent our prey's position within its environment (green bars show where our prey detected "food", blue bars show "water", the absence of any bar indicates the prey detected no resources present in its immediate environment at that time and yellow bars indicate the prey was in its "nest") while each line represents the internal variables of our prey, including each of its physiological variables (black for blood sugar, blue for vascular volume, yellow for integrity) and hormone levels/effects. Additionally, here vPV3 represents the prey's *perceived* level of integrity, which in the modulated conditions differed from the *actual* level of integrity when concentrations of hormone were released.

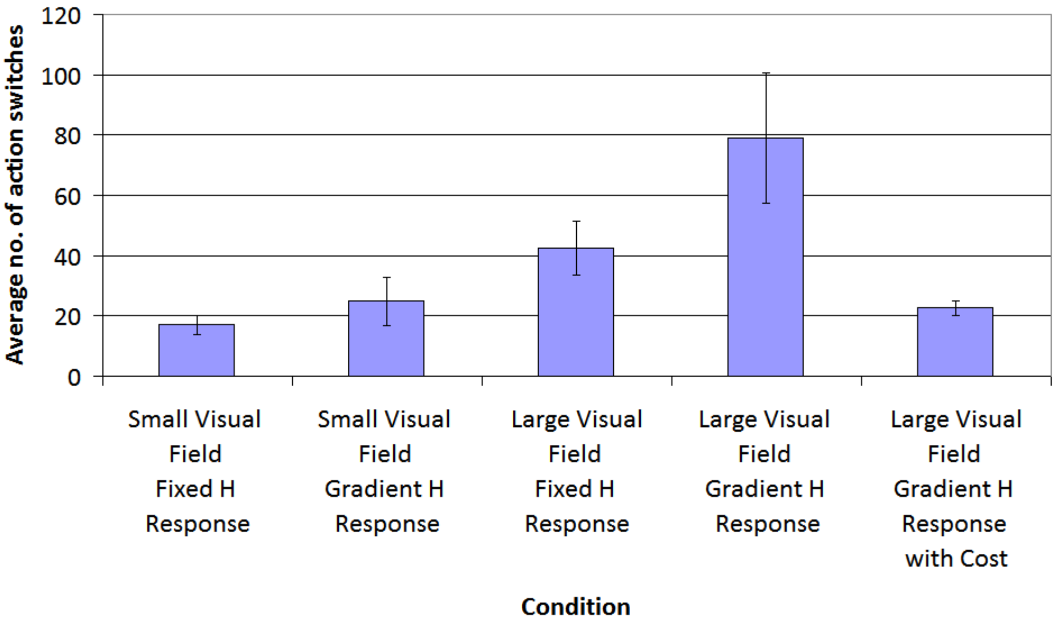

Figure 7.39: Summary of results showing the average number of "action switches" of our prey across conditions. Error bars show the Standard Error of the Mean (SEM). Exploring the idea of action selection as a cycle of activities, rather than a series of individual decisions, we use this simple measure to compare our prey. An "action switch" occurs whenever our prey makes the decision to switch from one of the six actions available to it to another. Here we can see that, on average, our prey with both a large visual field and gradient hormone response performed the greatest number of action switches.

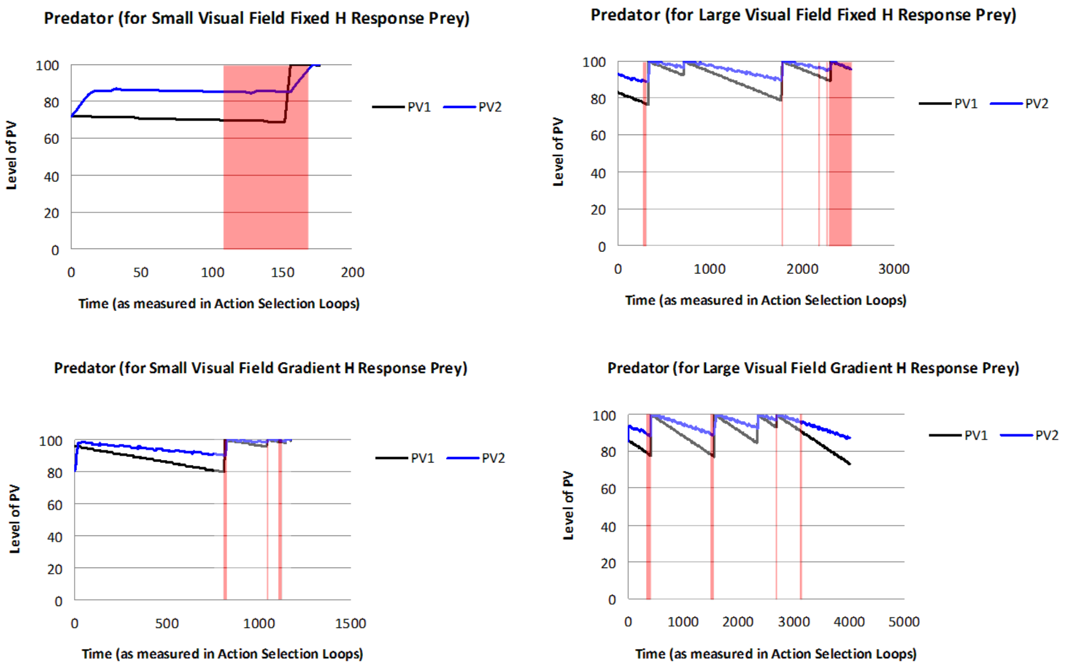

Figure 7.40: Illustration of differences (relating to aspects of *predator* Brain, Body and Environment over time) as observed across conditions between individual runs and shown by our BBE Maps. The background bars here represent our predator's detection of the prey in its environment (red bars show where our predator has detected the prey via its sensors) while each line represents the internal variables of our predator.

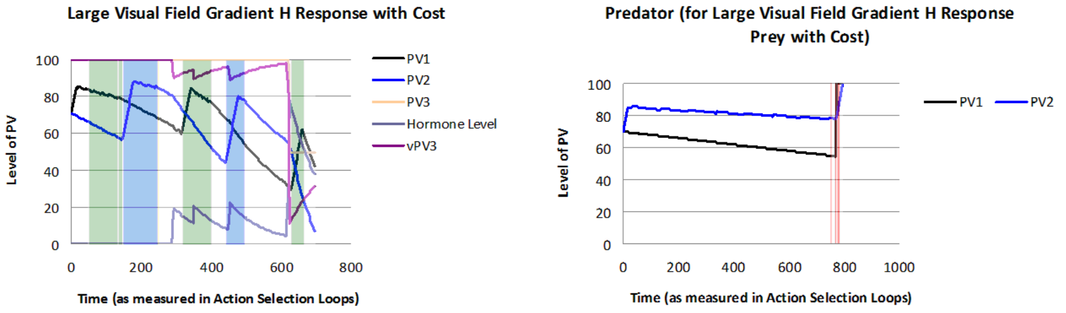

Figure 7.41: Example BBE Map constructed from one of the runs of our prey in our "Large Visual Field Gradient H Response with Cost" condition (not generally a part of our analysis here, but included for interest).

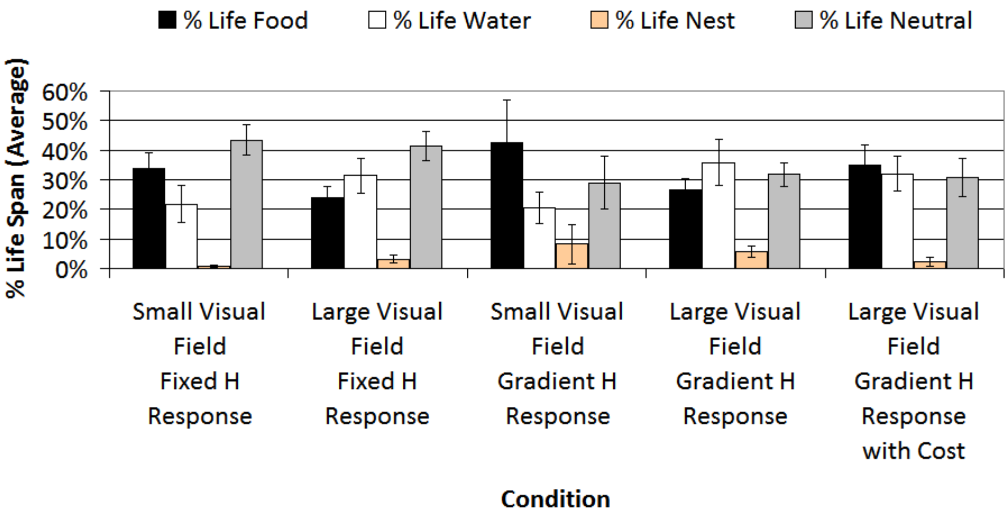

Figure 7.42: Summary of results showing time (as the average proportion of our prey's life span) spent on different resources/areas in its environment/arena across conditions where distal prey versus proximal predator and the hormone release of our emotion-based architecture is fixed. Error Bars show the Standard Error of the Mean (SEM).

prey indicates less "cost" in terms of the hormone used. With regards to "costs" of physical actions in particular, we conclude our emotion-based architecture could have more potential than has been shown in the H3RP. From an observer's point of view, it also seems to result in a prey that is more like a prey animal, and less like the abstracted prey agents found in simulation. For our prey with "gradient" hormone-release rates (release being a function of the distance away the predator is perceived) are not as ready to give up their sometimes hard-found resources. Indeed, for them, the emergent functions that we may associate with our architecture is not so much related to notions of "panic" as it is for our previously-studied prey with "fixed" hormone-release. We conclude that varying the integration (and gradience) of our mechanism may result in more emergent functions than have yet been seen. That is, beyond a simple "flight" response from the prey. Moreover, we conclude that our emotion-based architecture is likely to benefit from (and have a role in) additional sensory integration.

Through our latest experiments, especially the introduction of our own novel architecture, we suggest that our work will advance the research area in relation to emotion-based architectures by suggesting a further way in which we might develop architectures such as Avila-García's [29] and Mendao's [84] using what we have discovered from a more systematic and incremental investigation of a robotic predator-prey relationship. Together with our previous results, these results may also enable us to further advance the research area from another direction by providing a further means for biological researchers (including ethologists such as Cooper [88]) to test theories and provide further proof of concepts for the potential role of brain, body and environment in shaping the predator-prey relationship.

# Chapter 8

# Conclusions and Future Work

*“The only source of knowledge is experience.”*

**— Albert Einstein**

The goals and ideas of this thesis resonate strongly with the ideas of embodied, embedded cognition. For, we suggest, if artificial emotion is to be adaptive for artificial agents, however this may be represented computationally, we need to understand the conditions making the presence of such emotion “just right”. In our case, we focus on not abstracting away aspects of body and environment, further incorporating them into the design and study of our animat robots instead. To do so, we take inspiration from biological ideas about neuromodulation using a hormone-like mechanism in our experiments to simulate neuromodulation (and fear-like states) at a more abstract level: looking to explore the role and importance of various aspects of brain, body and environment in the generation of adaptive action selection/behaviour.

In this chapter we will first draw together, for a brief overview, the results obtained across all of our experiments during the study of our developed system (H3RP). As our concluding chapter we will therefore consider and provide a review of the conclusions drawn from all our experiments, individually as well as collectively. Following this, we bring together our observations from these experiments, as researchers of action selection and emotion and explorers of brain-body-environment interactions, to discuss the relative merits of our chosen measures of adaptive value (as defined in earlier chapters of this thesis and used to analyse the results of our experiments). By reiterating the contributions of this thesis, we will suggest how our work might be extended, using our research experience to suggest further contributions towards an analysis of the predator-prey relationship in terms of costs and benefits. Finally, we suggest how our model of embodied action selection, emotion and the predator-prey relationship might now be useful to others.

## 8.1 Results Summary

In this thesis, we first presented the results of a simple set of experiments exploring the effect of our implemented prey robot's perceptual distance and presence/absence of an emotion-based architecture. We used these results to suggest that the hormone-like component of our architecture will be adaptive in conditions where this mechanism is somehow *predictive* of environmental dynamics (that is to say, the predator's behaviour). From both the experimental results and our own observations, we suggest that one of the emergent functions of the hormone-like mechanism in our emotion-based architecture is that it acts as a predictor of environmental dynamics. Our mechanism can further be connected to ideas of anticipation in that our "seeing" prey, when implemented with our emotion-based architecture, was able to perform at least a functional equivalent of "anticipation" of future attacks and, through interoception, was seen to be distributing its time between activities in a way more conducive to survival (that is to say, adapting to the environmental dynamic introduced by the presence of the predator). In terms of brain-body-environment interactions the results of this chapter further support the idea that the body can be important for action selection, further suggesting a possible role for emotion-based architectures in adjusting connections between brain and environment adaptively.

From these results, we then went further to systematically vary aspects of both prey and predator embodiment (in terms of perceptual distance) as well as exploring the parameters of our hormone-like mechanism (in terms of strength and adopting a sigmoidal decay rate). Taking a more in-depth look at the predator-prey relationship, as well as our hormone-like mechanism, we varied both environmental dynamics of our prey (in terms of predator perceptual distance) and hormone strength. Due to observations made whilst conducting the previous experiments, a sigmoidal decay rate was introduced to our basic hormone-like mechanism. This changed the boundaries where one action switches to another by allowing our simulated hormone to stay in the prey's system for longer. This allowed the hormone-like mechanism to influence our action selection for longer, encouraging further persistence in seeking refuge after an attack.

From our results, we conclude that the mechanism we are studying will be adaptive in conditions where this mechanism is an accurate reflection of how *dangerous* the predator is. Experimenting with the measures used, we also see that calculations using time-steps sometimes provide us with a different look at the data than more "real-time calculations" in minutes and seconds. In terms of brain-body-environment interactions the results of this chapter further support the idea that the emergent functions of our emotion-based architecture will be affected by prey brain-body combinations (proximal/distal and hormone released) as well as environmental dynamics in the form of predator brain-body combinations (proximal/distal).

Finally, in our last set of experiments, we have demonstrated how our emotion-based architecture might benefit from (and have a role in) additional sensory integration. For, together, these experiments show that the mechanism we are studying emerges further functions from brain-body-environment interactions when integrated further with sensory abilities. Firstly, we found that a gradient hormone-release for our prey indicates less "cost" in terms of the hormone used. From an observer's point of view, it also seems to result in a prey that is more like a prey animal, and less like the abstracted prey agents found in simulation. For our prey with "gradient" hormone-release rates (release being a function of the distance away the predator is perceived) are not as ready to give up their sometimes hard-found resources. Indeed, for them, the emergent functions that we may associate with our architecture is not so much related to notions of "panic" as it is for our previously-studied prey with "fixed" hormone-release.

Secondly, we found that the mechanism we are studying can have additional functionality when integrated with behavioural response of our architecture as well. That is, the intensity of the action selected by the architecture. This can be thought of as akin to the flight/fight physiological response (readying body for action) and it has interesting connections to the ethological literature. We thus conclude that varying the integration (and gradience) of our mechanism may result in more emergent functions than have yet been seen. That is, moving beyond the elicitation of a simple and reflexive "flight" response from the prey, and moving further towards the development of architectures which allow our agents to function more as adaptive assessors of and responders to risk.

Turning back to our initial research question, our results thus provide evidence to show the ways in which the perceptual abilities of predator and prey agents can interact to affect the dynamics of their relationship. Specifically, in the context of our model, we can see that the perceptual abilities of *both* agents affect the performance (and thus adaptive value) of our chosen emotion-based architecture for action selection. For such artificial emotion to be adaptive for a prey agent, our work suggests this agent will benefit if the combination of perceptual abilities and our emotion-based architecture:

- Provides an accurate prediction of environmental dynamics. In our work, this refers to the chances of an attack by the predator.

- Provides an accurate assessment of the threat involved or represented by such a dynamic. That is, the degree of risk or hazard to agent survival it represents. In our work, this refers to how *dangerous* the detected presence of the predator is (or is otherwise considered to be) by the prey.

- Integrates the artificial emotion and perceptual abilities in such a way as to manage situations where it may be beneficial for an agent to take risks or otherwise balance potential opportunity costs, which might be caused if the agent responds too hastily to the hazardous aspect of its environment. For example, a starving prey detecting a predator at a distance may benefit more from staying on the food resource and continuing to feed than instigating a “flee” response automatically.
- Integrates information obtained from the combination of perceptual abilities and emotion-based architecture with the response. Specifically, in our experiments, this has appeared true of the *intensity* of prey response.

## 8.2 Levels of Analysis and Measures of Adaptive Value

With regards to different levels of analyses and measures of adaptive value, we conclude our work demonstrates that measures such as physiological balance and overall comfort are not necessarily, by themselves, informative indicators of adaptive behaviour per se. For, drawing further on our observations and more qualitative measures (including our own maps) we have sometimes been able to identify those “flashes” of adaptive behaviour, or “life-death” moments characteristic of the predator-prey relationship where a split-second decision can mean the difference between surviving or not, that these more general measures (looking at averages across multiple lifetimes, rather than the specific patterns in an individual’s lifetime) simply do not register. Having said this, by combining information obtained from using multiple levels of analyses, we have been able to build up a more detailed picture of our system in action. Thus, we consider all the measures we have studied to be most usefully thought of as tools which complement one another, rather than as tools which offer competing alternatives where one will always out-perform or offers more than another.

## 8.3 Three is the magic number?

Throughout this thesis, it has been interesting to see how often we can find examples where models that make concepts easier to understand make use of triads. Indeed, this is true both internally and externally to the area of science. For a florist, working in threes is quite often the route to the most aesthetically-pleasing arrangement. For a writer, it can be a key way to make an impact, whether this be in poetry or headlines. For a scientist, the number three is evidenced to us in the composition of an atom (consisting of protons, neutrons and electrons) or, perhaps more relevantly to our research, MacLean's early ideas of the triune brain (as composed of the reptilian complex, paleomammalian complex and neomammalian complex). Indeed, it is a number that often combines well with our own abilities, making summaries of key elements easier to remember and think about. The "three Rs" of education illustrate this point quite well (Reading, Writing and Arithmetic). For, while it is not the most accurate, it is certainly more memorable. However, that is not to say that we believe our model has proven itself only to be a memorable way of summarising key elements (brain, body and environment). Rather, from the results of our experiments, we suggest that models such as ours will prove more useful in generating more useful knowledge and understanding about action selection architectures at work in our own world. That is, action selection and the possible role for artificial emotion in real-time, dynamic environments.

## 8.4 Optimal Escape Theory and FID

According to Meyer the "benefits of using... truly rich sensori-motor equipment can be investigated only through real-world experiments, because no simulation is accurate enough to correctly reproduce the inner workings of, for example, a whisker system". Inspired by this argument, in summarising the conclusions of this thesis, we have further considered whether and how it may be possible to extract further insights from our work and draw broader conclusions by exploring the proven potential, with regards to the "richness", of our implementation so far.

To do this, we can think about how the results of our experiments compare to real-world experiments on/studies of biological agents. In particular, in our work we have identified such experiments in the field of ethology. That is, the study of animals. We thus believe our work might fit into the very same concepts and framework that researchers in ethology have proposed with regards to the flight initiation distance (FID) of various animals. As previously defined, the flight initiation distance point of a given animal (or in our case robot) can be described as that point at which a given prey animal (or robot) "decides" to flee in the presence of a given stimulus. Researchers have introduced various theories to explain the adaptive behaviour of animals. However, Cooper suggests that "decisions regarding flight initiation distance have received scant theoretical attention". With regards to this particular metric, in the course of our research we have developed a particular interest in it as it not only seems to be associated with the ethologist's functional equivalents (or emergent function) of "fear"[1] but much of the evidence, with regards to the information it is suggested different animals use in determining this distance, seems to complement our own studies of artificial

---

[1] This is said with an awareness that many ethologists, as observers of behaviour, are likely to cringe at the cognitive implications of the term were we to attribute it directly to the animals they are studying

prey agents. This includes the presence of food, distance to refuge, intruder starting distance, intruder speed and individual differences.

Indeed, reflecting on the results collected in this thesis, from an interdisciplinary perspective, as well as an animat, bottom-up one, we now hope that our findings may have an appeal for and even act as inspiration for robotic models to complement existing literature on biological predator-prey relationships. Furthermore, our model of embodied action selection shows a possible application for our model across disciplines. Namely, to test theories and models of flight initiation distance. According to Cooper, it is a model first proposed by Ydenberg and Dill (1986. The economics of fleeing from predators) that has "guided research for the past 20 years". According to the literature, this model "specifies that escape begins when the prey detects a predator, monitors its approach until the costs of escape and of remaining are equal, then flees". Flight initiation is thus supposed to occur where the costs of remaining and costs of fleeing intersect. However, Cooper's research has concluded that optimality models might be preferable to such a break-even model for a number of reasons. Regardless of model, however, the performance of our emotion-based architecture would seem to encourage further collaboration between disciplines, which could prove conducive to increasing our understanding of both the biological and the artificial.

## 8.5 Future Work

Thus future work could include a more interdisciplinary look at the predator-prey relationship. However, more specifically to the literature of artificial intelligence and artificial life, future work might now include the further development of mechanisms such as ours to continue to study further the emotion-like reactions resulting from use of this mechanism. In particular, we see one of the next steps as trying to find other ways to measure and quantify the differences in our system across conditions. Taking inspiration from the re-conceptualisation of action selection as activity cycles, for example, this might be in the form of adopting measures that are more ethologically-inspired. In future, we would like to see how researchers might further integrate different types of hormone-modulated sensory ability using the same basic architecture to study the dynamics of "fight-or-flight" emotional reactions in more complex prey-predator scenarios, including scenarios where the predator's abilities are also modulated. More broadly, however, it is hoped that some of the ideas within this thesis will encourage and be adopted or adapted by others so as to contribute towards acquiring knowledge as to the intricate relations likely to exist between artificial brains, bodies and environments.

# Publications

1. **“The Importance of the Body in Affect-Modulated Action Selection: A Case Study comparing Proximal versus Distal Perception in a Prey-Predator Scenario.”** by Claire O’Bryne, Lola Cañamero and John C. Murray.

   *submitted to and accepted for presentation at the 3rd International Conference on Affective Computing and Intelligent Interaction, Amsterdam, the Netherlands, 2009. (IEEE)*

2. **“Emotion in Decisions of Life and Death - Its Role in Brain-Body-Environment Interactions for Predator and Prey.”** by Claire O’Bryne and Lola Cañamero.

   *submitted to and accepted for presentation at the ALife XII Conference, Odense, Denmark, 2010. (MIT)*

# Glossary

## Action Selection

Thought of as a way of characterising the most basic problem of intelligent systems: what to do next.

## Adaptive value

The ability of our developed agents and/or their architectures to produce adaptive behaviour in response to and ensures survival of given environmental dynamics. For example, specific variations of our motivation-based architecture have been found to have greater adaptive value in certain circumstances over others.

## Animat Approach

An approach to addressing problems including perception, categorisation and sensorimotor control that focuses on characteristics neglected by standard approaches to AI. This approach takes inspiration from what is known about mechanisms underlying biological intelligence and approaches AI in a bottom-up manner.

## Animat

A simulated animal or robot whose structures and functionalities are as much inspired from current biological knowledge as possible, in the hope it will exhibit at least some of the survival capacities of real animals. The term “animat” was introduced by Wilson in 1985.

---

## Arbitration mechanism

In the context of action selection or control architectures, the arbitration mechanism/action policy refers to the way in which conflicts between competing behavioural alternatives are resolved. In the research of this thesis specifically, we use a "winner-take-all" arbitration mechanism.

## Autonomous

Having self-government, acting independently and/or having the freedom to do so. That is to say, the state of functioning independently, without extraneous influence. This further includes the ability of an agent to make its own decisions.

## Bottom-up Approach

In the context of the animat approach, this refers to an approach which goes from simple, but "complete" creatures to more complex ones. This is as opposed to a top-down approach.

## Body

In biological terms, this can be defined as the physical structure of a person or animal, including bones, flesh and organs. In the context of our research, as with "brain" we define "body" that bit more abstractly: focusing on "body" more as that which embodies, or gives concrete reality, to an agent.

## Brain-Body-Environment Relationship

Refers to the collective relationship and inter-connectedness of brain, body and environment.

## Brain

In biological terms, this can be defined as an organ of soft nervous tissue contained in the skull of vertebrates, functioning as the coordinating center of sensation and intellect. In the context of our own work, however, we take a more abstract definition of what a “brain” is: focusing less on the physical substrate and more on the functions of the brain as a *control mechanism*.

## Competitive Two Resource Problem (CTRP)

An extension of the Two Resource Problem which introduces two robots in the same environment simultaneously performing their own two-resource problem. This introduces competition for resources, as both robots might need access to the same resource at the same time.

## Distal Perception

The ability to sense something situated away from the center of the body or from the point of attachment.

## Embodied, Embedded Cognition

A philosophical theoretical position in cognitive science, closely related to situated cognition, embodied cognition, embodied cognitive science and dynamical systems theory. The theory states intelligent behaviour emerges out of the interplay between brain, body and world. The world is not just that on which the brain is acting. Rather, brain, body and world are equally important factors in the explanation of how particular intelligent behaviours come about in practice.

## Embodiment

A tangible or visible form of an idea, quality or feeling, along with the representation or expression of something in such a form. In our research, this refers to the idea that internal and external features of physical “body” (including homeostatic and hormonal states) heavily influence the higher “cognitive” processes in the brain, presumably via the emotional system (which also follows theories such as Damasio’s somatic marker theory).

## Emergent

The way in which complex systems and patterns arise out of a multiplicity of relatively simple interactions.

### Emotion-based architecture

A control architecture which takes inspiration in some manner from biological theories of emotion (and its role in action selection problems confronted by animals and/or humans).

### Emotion

A natural, instinctive state of mind deriving from circumstances, mood or relationships with others and/or any of the particular feelings that define such a state of mind, including anger, fear, love, hate etc. Thought to be related to certain activities in brain areas that direct our attention, motivate our behaviour and determine the significance of what is going on around us.

### Environment

The surroundings or conditions in which a person, animal or plant lives or operates and/or the setting or conditions in which a particular activity is carried on.

### Ethology

The science of animal behaviour or the study of human behaviour/social organisation from a biological perspective.

### Exteroception

Sensitivity to stimuli originating outside of the body.

### Hazardous Three Resource Problem (H3RP)

An extension of the Two Resource Problem (and in contrast to the Competitive Two Resource Problem) that introduces a hazardous environmental dynamic through via an additional resource as well as a "predator" robot. In this scenario, the robot of the original TRP has to effectively deal with the dynamics introduced by the "predator" by choosing between these three resources.

## Hormone-like mechanism

A mechanism which includes, is modelled on or takes inspiration from biological hormones (regulatory substances produced in an organism and transported in tissue fluids such as blood or sap to stimulate specific cells or tissues into action).

## Internal Robotics

An area of research that concentrates on reproducing in robots the inside of the body of organisms and studies the interactions of a robot's control system with what is inside the body. This is in contrast to external robotics, the more often-studied area, which concentrates on the interactions of an organism's nervous system with the external environment.

## Interoception

Sensitivity to stimuli originating inside of the body.

## Motivation-based architecture

A type of control mechanism/action selection architecture that takes inspiration from biological ideas about motivations and their role in action selection.

## (Neuro)modulation

In neuromodulation several classes of neurotransmitters regulate diverse populations of central nervous system neurons (one neuron uses different neurotransmitters to connect to several neurons).

## Predator-Prey Relationship

The interaction between types of agents (predators) that survive by taking resources from another type of agents (prey).

## Predator

An agent that survives by preying on other agents.

## Prey

An agent hunted by another because it is a source for its survival.

## Proximal Perception

The ability to sense something situated nearer to the center of the body or point of attachment.

## Situatedness

Sometimes used interchangeably with the idea of embeddedness, this refers to the idea that the physical interaction between body and environment strongly constrains the possible behaviours of an agent, in turn influencing (and in fact partly constituting) the cognitive processes that emerge out of the interaction between agent and environment

## Top-down approach

The inverse of the bottom-up approach, this method involves looking at the highest conceptual level, before working down to the details.

## Two Resource Problem (TRP)

A scenario known in the action selection literature, providing a standard that allows comparison of results between researchers, where a self-sufficient (biological or artificial) creature must continuously decide which of its two survival-related needs to satisfy by choosing between two resources available in the environment. This can be seen as the minimal scenario to test action selection mechanisms.

## Viability

Capable of success or continuing effectiveness.

## Voting-based arbitration mechanism

In the context of action selection, this is where resolution of conflicts between competing behavioural alternatives allow for “compromise actions” i.e. actions that are not best to satisfy any active goal in isolation, but rather compromise between multiple goals.

## Winner-take-all arbitration mechanism

In the context of action selection, this is where resolution of conflicts allow only the behavioural alternative with the highest activation to stay active.

# References

[1] D. McFarland and T. Bösser. *Intelligent Behavior in Animals and Robots*. Complex Adaptive Systems. MIT Press, 1993. 2

[2] O. Holland and D. McFarland. *Artificial Ethology*. Oxford University Press, 2001. 2, 15

[3] Dario Floreano and Claudio Mattiussi. *Bio-Inspired Artificial Intelligence: Theories, Methods, and Technologies*. The MIT Press, 2008. 2

[4] G Buason and Tom Ziemke. Competitive co-evolution of predator and prey sensory-motor systems. In Raidl et al, editor, *Applications of evolutionary Computing, EvoWorkshops*, pages 605–615, 2003. 2, 26

[5] Tom Stafford and Kevin N Gurney. Biologically constrained action selection improves cognitive control in a model of the stroop task. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1485):1671–1684, 2007. 2

[6] Christopher G. Langton. *Artificial Life: An Overview*. MIT Press, Cambridge, MA, USA, 1st edition, 1995. 2

[7] Mark Bedau. The scientific and philosophical scope of artificial life. *Leonardo*, 35(4):395–400, 2002. 2

[8] Artificial intelligence, artificial life and alternatives to computationalism and objectivism, 1995. 2

[9] M. Negnevitsky. *Artificial intelligence: a guide to intelligent systems*. Addison-Wesley Longman, 2005. 3

[10] J.J. Bryson. Mechanisms of action selection: Introduction to the special issue. *Adaptive Behavior - Animals, Animats, Software Agents, Robots Adaptive Systems*, 15(1):5–8, 2007. 3

[11] AK Seth. The ecology of action selection: Insights from a-life. *Philosophical Transactions of the Royal Society Series B*, 362(1485):1545–58, 2007. 3, 21

[12] B. Girard, N. Tabareau, Q. C. Pham, A. Berthoz, and J. J. Slotine. 2008 special issue: Where neuroscience and dynamic system theory meet autonomous robotics: A contracting basal ganglia model for action selection. *Neural Netw.*, 21(4):628–641, May 2008. 3

[13] R. Brooks. Asynchronous distributed control system for a mobile robot. In *SPIE Conference on Mobile Robots*, pages 77–84, 1986. 3, 13

[14] Toby Tyrrell. *Computational mechanisms for action selection*. PhD thesis, University of Edinburgh, Edinburgh, 1993. 3, 17, 19, 20, 22, 25, 26

[15] B. Blumberg. Action-selection in hamsterdam: Lessons from ethology. In *Proceedings of the third international conference on Simulation of adaptive behavior: from animals to animats*, volume 3, pages 108–117, 1994. 3

[16] Joanna Bryson and Jessica Flack. Emotions and action selection in an artificial life model of social behavior in non-human primates. In *In Charlotte Hemelrijk, editor, Proceedings of the International Workshop on Self-Organization and Evolution of Social Behaviour, Monte Verita*, 2002. 3

[17] E Tanguy, P Willis, and JJ Bryson. Emotions as durative dynamic state for action selection. In *Proceedings of the 20th international joint conference on Artificial Intelligence*, pages 1537–1542, 2007. 3, 24

[18] Rodrigo Ventura, Carlos Pinto-ferreira, and Rua Rovisco Pais. Emotion-based agents: Three approaches to implementation (preliminary report), 1999. 3

[19] JM Fellous. Neuromodulatory basis of emotion. *The Neuroscientist*, pages 283–294, 1999. 3, 24, 34, 39

[20] J.M. Fellous and M.A. Arbib. *Who Needs Emotions?: The Brain Meets the Robot*. Series in Affective Science. Oxford University Press, 2005. 3, 34, 66

[21] Hillel J. Chiel and Randall D. Beer. The brain has a body: adaptive behavior emerges from interactions of nervous system, body and environment. *Trends in Neuroscience*, 20:553–557, 1997. 3, 18, 29

[22] A. Clark. *Being There: Putting Brain, Body, and World Together Again*. Bradford Book. Mit Press, 1998. 4

[23] M. Wilson. Six views of embodied cognition. *Psychonomic Bulletin & Review*, 9(4):625–636, 2002. 4

[24] O. Miglino, D. Parisi, and M. Ponticorvo. The role of internal stimuli in building up an 'inner world'. a study with evolved 'blind' robots. In G. Pezullo, MV. Butz, O. Sigaud, and G. Baldassarre, editors, *Proceedings of the fourth workshop on anticipatory behavior in adaptive learning systems (ABiALS)*, June 2008. 4, 18, 29

[25] C. Breazeal. A motivational system for regulating human-robot interaction. In *Proceedings of the National Conference on Artificial Intelligence*, pages 54–61. John Wiley & Sons Ltd, 1998. 4

[26] M. Bernard, S. NGuyen, P. Pirim, B. Gas, and J.A. Meyer. Phonotaxis behavior in the artificial rat psikharpax. In *International Symposium on Robotics and Intelligent Sensors, IRIS*, 2010. 4, 17

[27] G Buason, N Bergfeldt, and T Ziemke. Brains, bodies and beyond: competitive co-evolution of robot controllers, morphologies and environments. *Genetic Programming and Evolvable Machines*, 6(1):25–51, 2005. 4, 21, 22

[28] A. Ohman and S. Wiens. On the automaticity of autonomic responses in emotion: An evolutionary perspective. *Handbook of affective sciences*, pages 256–275, 2003. 4

[29] Orlando Avila-García. *Towards emotional modulation of action selection in motivated autonomous robots*. PhD thesis, University of Hertfordshire, Hertfordshire, 2004. 6, 16, 17, 19, 21, 22, 24, 26, 32, 37, 38, 41, 63, 92, 172

[30] Jeffrey Dean. Animats and what they can tell us. *Trends in Cognitive Sciences*, 2:60–67, February 1998. 13, 14, 15

[31] Rodney A Brooks. Achieving artificial intelligence through building robots. MIT AI Lab Memo No. 899, May 1986. 13, 15

[32] R. Brooks. A robust layered control system for a mobile robot. *Robotics and Automation, IEEE Journal of*, 2(1):14–23, 1986. 13

[33] Rodney A Brooks. Intelligence without representation. *Artificial Intelligence Journal*, 47:139–159, 1991. 13

[34] Rodney A Brooks. Intelligence without reason. MIT AI Lab Memo No. 1293, April 1991. 13

[35] Stewart Wilson. Classifier systems and the animat problem. *Machine Learning*, 2(3), November 1987. 13, 19

[36] Jean-Arcady Meyer and Agnès Guillot. Simulation of adaptive behavior in animats: review and prospect. In *Proceedings of the first international conference on simulation of adaptive behavior on From animals to animats*, pages 2–14, Cambridge, MA, USA, 1990. MIT Press. 13, 16

[37] A. Guillot and J.A. Meyer. The animat contribution to cognitive systems research. *Cognitive Systems Research*, 2(2):157–165, 2001. 14

[38] T. Ziemke and N.E. Sharkey. A stroll through the worlds of robots and animals: Applying jakob von uexkull's theory of meaning to adaptive robots and artificial life. *SEMIOTICA-LA HAYE THEN BERLIN-*, 134(1/4):701–746, 2001. 14

[39] R. Pfeifer. Building fungus eaters: Design principles of autonomous agents. *From animals to animats*, 4:3–12, 1996. 14

[40] Jean-Arcady Meyer. The animat approach: simulation of adaptive behaviour in animals and robots. In *In Animals and Robots, NSI98*, 1998. 15, 16, 29

[41] B. Webb. Animals versus animats: Or why not model the real iguana? *Adaptive Behavior*, 17(4):269–286, 2009. 15

[42] L. Steels. When are robots intelligent autonomous agents? *Journal of Robotics and Autonomous Systems*, 15:3–9, July 1995. 15, 16, 27

[43] Luc Steels, Peter Stuer, and Danny Vereertbrugghen. Issues in the physical realisation of autonomous robotic agents. reports on the physical realisation of a robotic ecosystem, 1996. 15, 16, 27

[44] Luc Steels. A case study in the behavior-oriented design of autonomous agents. In *SAB94: Proceedings of the third international conference on simulation of adaptive behavior on From animals to animats 3*, Cambridge, MA, USA, 1994. MIT Press. 15, 16, 27

[45] P. Maes. Modeling adaptive autonomous agents. In C Langton, editor, *Artificial Life: an overview*, pages 135–162. MIT Press, Cambridge, MA, USA, 1994. 15, 16, 20, 29

[46] Jean-Arcady Meyer and Agnès Guillot. From sab90 to sab94: four years of animat research. In *SAB94: Proceedings of the third international conference on simulation of adaptive behavior on From animals to animats 3*, Cambridge, MA, USA, 1994. MIT Press. 15, 16

[47] Tom Ziemke. Embodied ai as science: models of embodied cognition, embodied models of cognition or both? *Embodied Artificial Intelligence*, pages 27–36, 2004. 15

[48] Barbara Webb. Can robots make good models of biological behaviour? *Behavioral and Brain Sciences*, 24:1033–1050, 2001. 15

[49] Randall D. Beer. A dynamical systems perspective on autonomous agents. Technical Report CES-92-11, Department of Computer Engineering and Science, Case Western Reserve University, 1992. 16

[50] George Kiss. Autonomous agents, ai and chaos theory. In *Proceedings of the first international conference on simulation of adaptive behavior on From animals to animats*, pages 518–524, Cambridge, MA, USA, 1990. MIT Press. 16

[51] Luc Steels. Towards a theory of emergent functionality. In *Proceedings of the first international conference on simulation of adaptive behavior on From animals to animats*, pages 451–461, Cambridge, MA, USA, 1990. MIT Press. 16

[52] Joanna Bryson. Cross-paradigm analysis of autonomous agent architecture. *Journal of Experimental and Theoretical Artificial Intelligence, JETAI*, 12(2):165–189, 2000. 17

[53] Agnès Guillot and Jean-Arcady Meyer. The animat contribution to cognitive systems research. *Journal of Cognitive Systems Research*, 2:157–165, 2001. 17

[54] Paolo Pirjanian. *Multiple objective action selection and behavior fusion using voting*. PhD thesis, Department of Medical Informatics and Image Analysis, Institute of Electronic Systems, Aalborg University, Aalborg, Denmark, 1998. online at: http://www.vision.auc.dk/ paolo/phd.html. 17

[55] B. Blumberg. Action-selection in hamsterdam: Lessons from ethology. In *Proceedings of the third international conference on Simulation of adaptive behavior: from animals to animats*, volume 3, pages 108–117, 1994. 17

[56] J.A. Meyer. From natural to artificial life: Biomimetic mechanisms in animat designs. *Robotics and Autonomous Systems*, 22(1):3–21, 1997. 17

[57] R.C. Arkin, K. Ali, A. Weitzenfeld, and F. Cervantes-Pérez. Behavioral models of the praying mantis as a basis for robotic behavior. *Robotics and Autonomous Systems*, 32(1):39–60, 2000. 17

[58] K. Gurney, T.J. Prescott, and P. Redgrave. A computational model of action selection in the basal ganglia. i. a new functional anatomy. *Biological cybernetics*, 84(6):401–410, 2001. 17

[59] T.J. Prescott. Comparing a brain-inspired robot action selection mechanism with'winner-takes-all'. In *From animals to animats 7: proceedings of the Seventh International Conference on Simulation of Adaptive Behavior*, volume 7, page 75. The MIT Press, 2002. 17

[60] O. Avila-García and L. Cañamero. Hormonal modulation of perception in motivation-based action selection architectures. In *Proc. Agents that Want and Like: Motivational and Emotional Roots of Cognition and Action, Symposium of the AISB*, volume 5, pages 14–15, 2005. 17, 18

[61] S. Franklin. Autonomous agents as embodied ai. *Cybernetics & Systems*, 28(6):499–520, 1997. 17

[62] Rolf Pfeifer and Josh C. Bongard. *How the body shapes the way we think: a new view of intelligence*. MIT Press, Cambridge, MA, USA, November 2006. 18

[63] Rolf Pfeifer and Christian Scheier. *Understanding intelligence*. MIT Press, Cambridge, MA, USA, 1999. 18

[64] Domenico Parisi. Internal robotics. *Connection Science*, 16(4):325–338, 2004. 18, 22

[65] Neural darwinism, 2004. Interview with G. Edelman online at http://www.digitalnpq.org/archive/2004_summer/edelman.html. 19

[66] Peter M. Todd and Stewart W. Wilson. Environment structure and adaptive behavior from the ground up. In Jean-Arcady Meyer, HL Roitblat, and Stewart W. Wilson, editors, *Proceedings of the second international conference on From animals to animats 2: simulation of adaptive behavior*, pages 11–20, Cambridge, MA, USA, 1993. MIT Press. 19, 20

[67] Lola Cañamero. Designing emotions for activity selection in autonomous agents. In Robert Trappl, Paolo Petta, and Sabine Payr, editors, *Emotions in humans and artifacts*, chapter 4, pages 115–148. MIT Press, Cambridge, MA, USA, 2003. 19, 20, 21, 24, 26, 35

[68] Richard S Sutton. On the significance of markov decision processes. In W. Gerstner, A. Germond, M. Hasler, and J-D Nicoud, editors, *Artificial Neural Networks ICANN*, pages 273–282. Springer, 1997. 19

[69] A Robert, F Chantemargue, and M Courant. Grounding agents in emud artificial worlds. In *Proceedings of the first international conference on Virtual Worlds*, 1998. 19

[70] Michael L. Littman. Memoryless policies: theoretical limitations and practical results. In *SAB94: Proceedings of the third international conference on Simulation of adaptive behavior : from animals to animats 3*, pages 238–245, Cambridge, MA, USA, 1994. MIT Press. 19

[71] Orlando Avila-García and Lola Cañamero. Using hormonal feedback to modulate action selection in a competitive scenario. In S Schaal, AJ Ijspeert, A Billard, S Vijayakumar, J Hallam, and J-A Meyer, editors, *Proceedings of the 8th international conference on simulation of adaptive behavior (SAB04): From animals to animats 8*, pages 243–252, 2004. 19

[72] R.A. Brooks. Elephants don't play chess. *Robotics and autonomous systems*, 6(1-2):3–15, 1990. 19

[73] O. Miglino, H.H. Lund, and S. Nolfi. Evolving mobile robots in simulated and real environments. *Artificial Life*, 2(4):417–434, 1995. 19

[74] P Todd, S Wilson, Somayaji A, and Yanco H. The blind breeding the blind: adaptive behavior without looking. In D Cliff, P Husbands, J-A Meyer, and S Wilson, editors, *Proceedings of the third international conference on the simulation of adaptive behavior: From animals to animats 3*, pages 228–237, 1994. 20

[75] D Floreano. Patterns of interactions in shared environments. In *Second European conference on Artificial Life, ECAL*, 1993. 20

[76] L. Yao, W.A. Sethares, and D.C. Kammer. Sensor placement for on-orbit modal identification via a genetic algorithm. *AIAA journal*, 31(10):1922–1928, 1993. 22

[77] SY Chen and YF Li. Automatic sensor placement for model-based robot vision. *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, 34(1):393–408, 2004. 22

[78] L. Cañamero, O. Avila-García, and E. Hafner. First experiments relating behavior selection architectures to environmental complexity. In *Intelligent Robots and Systems, 2002. IEEE/RSJ International Conference on*, volume 3, pages 3024–3029. IEEE, 2002. 22

[79] F.J. Odling-Smee, K.N. Laland, and M.W. Feldman. *Niche Construction: The Neglected Process in Evolution*. Monographs in Population Biology. Princeton University Press, 2003. 22, 25

[80] Timothy Rumbell, John Barnden, Susan Denham, and Thomas Wennekers. Emotions in autonomous agents: comparative analysis of mechanisms and functions. *Autonomous Agents and Multi-Agent Systems*, 25(1):1–45, 2012. 23

[81] Sandra Clara Gadanho and John Hallam. Robot learning driven by emotions. *Adaptive Behavior*, 9(1):42–64, 2001. 24

[82] Carlos Delgado-Mata and Ruth S Aylett. Fear and the behaviour of virtual flocking animals. In *Advances in Artificial Life*, pages 655–664. Springer, 2007. 24, 38, 39

[83] Ernesto Burattini and Silvia Rossi. Periodic activations of behaviours and emotional adaptation in behaviour-based robotics. *Connection Science*, 22(3):197–213, 2010. 24

[84] Miguel Mendao. Hormonally moderated neural control. In *Proceedings of the 2004 AAAI Spring Symposium. AAAI Press, Menlo Park, CA*, 2004. 24, 38, 40, 92, 172

[85] R.S. Lazarus. *Emotion and adaptation*. Oxford University Press, 1994. 24

[86] A. Damasio. *Descartes' Error: Emotion, Reason and the Human Brain*. Random House, 2008. 24, 34

[87] I. Kelly, O. Holland, M. Scull, and D. McFarland. Artificial autonomy in the natural world: building a robot predator. In *Proceedings 5th European conference on Artificial Life, ECAL*, pages 289–93, 1999. 25, 26

[88] William E. Cooper, Valentn Prez-Mellado, Teresa Baird, Troy A. Baird, Janalee P. Caldwell, and Laurie J. Vitt. Effects of risk, cost, and their interaction on optimal escape by nonrefuging bonaire whiptail lizards, cnemidophorus murinus. *Behavioral Ecology*, 14(2):288–293, 2003. 25, 124, 126, 173

[89] D. Floreano and S. Nolfi. Adaptive behavior in competing co-evolving species. In P Husbands and I Harvey, editors, *Proceedings of the fourth European conference on Artificial Life*, pages 378–387. MIT Press, 1997. 26

[90] Lola Cañamero. Emotion understanding from the perspective of autonomous robots research. *Neural Networks*, 18(4):445–455, 2005. 29, 66

[91] Robert Lowe, Carlos Herrera, Anthony Morse, and Tom Ziemke. The embodied dynamics of emotion, appraisal and attention. In Lucas Paletta and Erich Rome, editors, *Attention in Cognitive Systems. Theories and Systems from an Interdisciplinary Viewpoint*, pages 1–20. Springer-Verlag, Berlin, Heidelberg, 2008. 29

[92] C. Breazeal and B. Scassellati. Infant-like social interactions between a robot and a human caregiver. *Adaptive Behavior*, 8(1):49, 2000. 35, 38, 39

[93] L. Moshkina, R.C. Arkin, J.K. Lee, and H.R. Jung. Time-varying affective response for humanoid robots. *Progress in Robotics*, 44:1–9, 2009. 38, 40

[94] R.C. Arkin. Moving up the food chain: Motivation and emotion in behavior-based robots. *Who Needs Emotions: The Brain Meets the Robot*, pages 245–270, 2005. 45

[95] B. Bagnall. *Maximum LEGO NXT: Building Robots with Java Brains*. Variant Press, 2007. 53

[96] L. Cañamero and O. Avila-García. A bottom-up investigation of emotional modulation in competitive scenarios. In *Proceedings of the 2nd international conference on Affective Computing and Intelligent Interaction*, pages 398–409. Springer-Verlag, 2007. 58, 59

[97] M. Neal and J. Timmis. Timidity: A useful emotional mechanism for robot control. *Informatica*, 27(4):197–204, 2003. 58

[98] M. Neal and J. Timmis. Once more unto the breach: Towards artificial homeostasis. *Recent developments in biologically inspired computing*, pages 340–365, 2005. 58

[99] J. Panksepp. *Affective Neuroscience: The Foundations of Human and Animal Emotions*. Oxford University Press, 1998. 66, 67

[100] D. Evans and P. Cruse. *Emotion, evolution, and rationality*. Oxford University Press, 2004. 66

[101] J. Nadel and D. Muir. *Emotional Development: Recent Research Advances*. Oxford University Press, USA, 2005. 66

[102] S. Wilson and N. Haslam. Is the future more or less human? differing views of humanness in the posthumanism debate. *Journal for the Theory of Social Behaviour*, 39(2):247–266, 2009. 66

[103] P.A. Spikins, H.E. Rutherford, and A.P. Needham. From hominininity to humanity: Compassion from the earliest archaics to modern humans. *Time and Mind*, 3(3):303–325, 2010. 67

[104] D.M. Buss. *The Handbook of Evolutionary Psychology*. John Wiley & Sons, 2005. 67

[105] A. Ohman. Of snakes and faces: An evolutionary perspective on the psychology of fear. *Scandinavian Journal of Psychology*, 50(6):543–552, 2009. 67

[106] Carroll E. Izard. The many meanings/aspects of emotion: Definitions, functions, activation, and regulation. *Emotion Review*, 2(4):363–370, 2010. 67

[107] Paul Ekman. An argument for basic emotions. *Cognition & Emotion*, 6(3-4):169–200, 1992. 67

[108] J.C. Borod. *The Neuropsychology of Emotion*. Series in Affective Science. Oxford University Press, USA, 2000. 67

[109] William E Cooper and William G Frederick. Predator lethality, optimal escape behavior, and autotomy. *Behavioral Ecology*, 21(1):91–96, 2010. 124

[110] P. Cobos, M. Sanchez, C. García, M. Nieves Vera, and J. Vila. Revisiting the james versus cannon debate on emotion: startle and autonomic modulation in patients with spinal cord injuries. *Biological Psychology*, 61(3):251 – 269, 2002. 151

[111] H. D. Critchley. Neural mechanisms of autonomic, affective, and cognitive integration. *The Journal of Comparative Neurology*, 493(1):154–166, 2005. 151