EXPLORING POSSIBLE ORGANIZATIONS FOR VISUALLY-GUIDED ORIENTATION

Zoltán Biró

Submitted by Zoltán Biró to the University of Skövde as a dissertation towards the degree of M.Sc. by examination and dissertation in the Department of Computer Science.

September 1997

I hereby certify that all material in this dissertation which is not my own work has been identified and that no work is included for which a degree has already been conferred on me.

Zoltán Biró

Abstract

This dissertation is concerned with exploring the behaviour and the internal properties of agent controllers which employ recurrent artificial neural networks, whose internal dynamics have been evolved by an evolutionary algorithm for coping with a visually-guided orientation task. Guided by a very simple fitness function, defined at behaviour level, the evolution allowed for emergence of different strategies, two of which are analyzed in detail using traditional connectionist analysis methods. The examined strategies succeed in the orientation task by inherently handling a simple form of the object persistence problem in the limited environment, and by employing behaviours which include passive and active search, object tracking, obstacle avoidance, and a simple form of discrimination, each of which were observed to correspond to hidden unit activation subspaces.

To my brother, László Biró

Contents

List of figures

Chapter 1

Introduction

Common to all animals inhabiting Earth is that they exhibit adaptive behaviour in their environments allowing them to survive long enough to breed and pass on their genetic material to their offspring. In their drive to remain viable, they exhibit sometimes very complex behaviours while interacting with their environment.

Braitenberg (1984) showed that it is very easy for an external observer to ascribe complexity to simple internal mechanisms, provided that the agents (e.g. animals; defined below) employing these mechanisms exhibit certain behaviours. Through series of thought experiments, he demonstrated, that vehicles using just two sensors wired to two motors showed behaviours that appeared sophisticated and complex enough for ascribing terms like avoiding, attacking, scurrying, basking, having anger, and knowing. Braitenberg's vehicles could even exhibit behaviours that were associated with fear, aggression, love, values, taste, free will, foresight, egoism, and optimism. (Sharkey & Heemskerk, 1997; Cliff, 1994).

Similarly, it appears to be very easy to ascribe cognitive phenomena such as intentionality to an object by observing its behaviour alone. According to Eysenck & Keane (1995), Dittrich and Lea (1994) identified three factors, each of which contributes to the perception of intentional motion in an object: it is moving in a direct fashion, it moves faster than other objects, and the goal towards which it is moving is visible. Subsequently, they concluded that "the perception of intentionality can be a relatively immediate, bottom-up process, probably occurring quite early in the visual processing" (Eysenck & Keane, 1995, p. 93).

The lesson drawn from these observations is thus that one cannot take for granted that there are internal cognitive or psychological structures corresponding to all phenomena that can be ascribed to the object, no matter how complex the exhibited behaviours are.

The cause of the behaviours has traditionally been explained as the result of the internal functioning alone, a view supported by classical computationalism (classical cognitive science). This view has recently been challenged by both the enactive (or embodied) approach to cognition (Varela, 1991) and the dynamical approach to cognition (van Gelder, 1992, 1997; Beer, 1995). Both approaches consider cognition as emergent from agent-environment interaction; in fact agent and environment are considered to form an indivisible pair.

> ... the approach of embodied cognition maintains that the mind or the cognitive system arises as a direct result of the *interaction* between organism and environment. It is not simply that the mind has pointers to the world; it is the interplay between beings and their world that creates minds. (Sharkey & Heemskerk, 1997, original emphasis)

Proponents of both the enactive and the dynamical approach would therefore claim, that while the internal functioning does play an important role in behaviour generation, the external world in which an agent is situated plays an equally important role, and that ultimately, behaviours are caused by the interaction between agent and environment. Similar arguments can be traced in the following passage:

> ... ecological psychologists understand visually guided locomotion as change in a dynamical system which includes aspects of both the organism and the environment (e.g., the optic flow; Warren, 1995). (van Gelder, 1997)

These approaches thus propose an alternative to the classical computationalist theory of behaviour (as causal relation of purely internal processing) for explaining why Braitenberg vehicles and organisms with simple nervous systems exhibit so complex behaviours[1]. According to Sharkey & Heemskerk (1997), “with appropriate environment, a simple mechanism may exhibit an apparently complex behavioural repertoire”.

A shift away from employing a central processor with deliberate reasoning as starting point for studying the nature of adaptive behaviour is also supported by Beer, who argues that “... from an evolutionary perspective, the human capacity for language and abstract thought is a relatively recent elaboration of a much more basic capacity for situated action that is universal among animals” (Beer, 1995, p. 175). Supporting this stance is research in developmental psychology, where for example Piaget argued that in infants sensorimotor intelligence is the basis of higher level intellect, such as abstract reasoning, thus deliberate reasoning emerges from simpler sensorimotor coordinations (Ginsburg & Opper, 1988).

Arguments like these have been raised to motivate the omission of deliberate reasoning as long as it can be done without impinging on the adaptiveness in behaviour. Leaving out deliberate reasoning implies that much of the phenomena, often descriptively defined in vague terms, can be set aside to focus on the problem of the nature of adaptive behaviour.

This is the outset of adaptive behaviour research[2], which approaches the nature of adaptive behaviour by studying autonomous agents and their interactions with their environments. According to Wheeler

1. “[(Simon, 1969)] pointed out that the complexity of the behaviour of [an ant walking along a beach] is more a reflection of the complexity of its environment than its own internal complexity.“ (Brooks, 1991, p. 584)

2. Also referred to as the animat approach to AI (Wilson, 1985).

> [a]n *autonomous agent*[3] is a fully integrated, self-controlling, adaptive system which, while in continuous long-term interaction with its environment, actively behaves so as to achieve certain goals. So, for a system to be an autonomous agent, it must exhibit *adaptive behaviour*[4], behaviour which increases the chances that the system can survive in a noisy, dynamic, uncertain environment. (Wheeler, 1994, p. 3, original emphasis)

Any natural or artificial agent can thus be considered as an autonomous agent, as long as it meets the criteria above. The specific branch of adaptive behaviour research that resides in artificial intelligence (AI) is primarily focused on studying artificial autonomous agents, while occasionally incorporating ideas from animal studies conducted by for example ethologists and neuroethologists. While the constituents of artificial agents are often extremely simplified models of those in their animate counterparts (the former often employs for instance just a couple of distance sensors for "vision") the advocates for studying adaptive behaviour using artificial agents argue against the necessity of physical realism by claiming that as long as these agents confront the similar environment as animals do, they raise similar issues about the nature of the adaptive behaviour. As Beer notes,

> [t]he agent's "vision" is certainly not intended as a serious model of the actual physics of light or photoreception. However, it does raise the similar issues in the perception of objects using a spatially-structured array of distal sensors. (Beer, 1996, p. 422)

---

3. There are numerous more or less restrictive definitions of the concept "autonomous agent" in the AI literature (Franklin, 1996 for an overview), which resulted in a redefinition of autonomy so as to refer to "a multidimensional continuum, not an all-or-nothing property (Boden, 1994, ...)." (Ziemke, 1997, p. 3)

The concept of "autonomous agent" used in this dissertation subscribes to a weak definition of autonomy, where the agent responses to external stimuli is based on inputs and internal states (A2) with the restrictions that the internal controller mechanisms are not selectively activated at any moment for behaviour generation (not A3). Furthermore, the controller mechanism is externally imposed rather than self-generated (not A1), and acquired through phylogenetic (species-level) rather than ontogenetic (individual level) learning during a separate training phase (no on-line adaptation).

(A1, A2, and A3 refer to Ziemke's three aspects of behavioural control which are "crucial" for autonomy.)

4. As Cliff points out, "Adaptation or plasticity may itself give rise to new or improved adaptive behaviors, but there are many cases of adaptive behaviors which are genetically determined (e.g. "hard-wired" behaviors such as reflexes and instincts)." (Cliff, 1994, pp. 1-2)

This view implies that new tools can be adopted from, for example, AI for studying the nature of adaptive behaviour, while avoiding ethical issues that are associated with animal experiments. Simulations of adaptive behaviour allow controllers employing for instance artificial neural networks (ANNs), evolved through simulated evolution with internal and external noise to exhibit certain behaviours, to be observed and analysed with great accuracy, at both behavioural and neural level, without the interference of internal (neural) noise; an advantage which is not available to neuroethologists (Cliff, 1992). This approach of studying the neural mechanisms responsible for generating adaptive behaviour in artificial agents (originally models of simple animals (Beer, 1990)) using computational modelling is sometimes referred to as computational neuroethology (Cliff, 1994).

The advantages of employing simulated evolution for evolving agent controllers for adaptive behaviour rather than designing controllers by hand is that a priori assumptions implying subsequent constraints, on controller design can be to a great extent avoided. Simulated evolution is guided for faster (than just random) convergence of controller organization for desired behaviours by a fitness function. As the fitness function can be specified at the level of behavioural constraints, the evolved organization of the controller may exhibit emergent behaviours that arise without explicit reference to the factors underlying these behaviours. This observation is supported by Cliff *et al.* (1992) who conducted experiments on evolving controllers and visual morphologies for an agent with two photoreceptors for vision, where the task of the agent was to reach to the centre of a simulated circular arena.

> As was demonstrated in [(Cliff *et al.*, 1993)], visual guidance emerges without explicit reference to vision in the [fitness function]. (Cliff *et al.*, 1992, p. 4)

Adaptive behaviour researchers employing evolutionary methods (evolutionary roboticists) generally agree that the study of the nature of adaptive behaviour should start by investigating the properties of minimal systems (simple controllers evolved for displaying simple behaviours; cf. Beer, 1996; Cliff *et al*., 1992). The investigation of how the internal properties of evolved minimal systems correlate to its behaviours is often complicated by the close intertwining of environment and agent. An example from neuroethology is mentioned by Brooks, who states that

> [f]or one very simple animal *Caenorhabditis elegans*, a nematode, we have a complete wiring diagram of its nervous system, including its development stages [(Wood, 1988)]. ... Even though the anatomy and behavior of this creature are well studied, and the neuronal activity is well probed, the way in which the circuits control the animal's behavior is not understood very well at all. (Brooks, 1991, p. 582, original emphasis)

As simple controllers only take us that far, there are also advocates for incremental evolution, who consider scaling up from simpler to more complex controllers by incrementally increasing the difficulty of the tasks while as starting point at every new task using the evolved controllers from agents that exhibited the most successful behaviours in the previous task (Harvey *et al*., 1996).

We have now reached the point which establishes the outset of this dissertation, namely the study of artificially evolved internal organization of controllers employed by artificial agents that exhibit adaptive behaviour.

1.1 Aim

The aim of the dissertation is to investigate the properties of internal organizations of artificially evolved recurrent connectionist controllers for visually-guided behaviour. The analysis of the underlying mechanisms would provide insights of internal structures and the strategies the agent employs for successful handling of the task.

1.2 Delimitation

The focus of the study is on organizations for visually-guided behaviour, where the task has been limited to simple visually-guided orientation.

The issues raised in visually-guided orientation can be approached differently depending on the initial assumptions made. This dissertation focuses on evolutionary approaches where the internal dynamics of recurrent connectionist controllers are evolved for visually-guided orientation.

There are several examples of evolutionary approaches to evolving controllers for visually-guided behaviour (e.g. Cliff *et al*., 1992, Harvey *et al*., 1994, Floreano & Mondana, 1994, Beer, 1996, Salomon, 1996). The most relevant of these are Beer's work (Beer, 1996), and the work done by at COGS, University of Sussex (Harvey *et al*., 1996).

The work conducted by Beer (Beer, 1996) forms the basis for the current work. Whereas Beer focused (mostly) on qualitative assertion of the evolved controllers at the level of behaviours, this project focuses on the internal controller mechanisms underlying the behaviours, where the internal dynamics of the controller is realized by a recurrent artificial neural network. The experiments are variations of the visually-guided orientation experiments carried out by Beer (1996), but have been extended to allow two dimensional agent motion rather than one dimensional motion as employed in the Beer's original experiments.

The analysis will be performed by employing standard connectionist analysis techniques. No dynamical analysis (e.g. by using dynamical systems theory) of the dynamical properties of the ANNs will be performed. No classical computational approaches will be considered in the dissertation for the reasons discussed in the next chapter.

1.3 Target group

This dissertation was written with a target group in mind. It requires knowledge corresponding to MSc level in artificial intelligence and especially connectionism and evolutionary algorithms, and preferably also some basic knowledge in cognitive science.

1.4 Dissertation outline

Chapter 2 will cover the background for the current work by providing a general overview of the foundations related to the approach taken in this dissertation. It will also cover some previous work in visually-guided behaviour, where especially evolutionary approaches are discussed. Chapter 3 gives the methodological and technical background to the design decisions and strategies by covering evolution mechanisms, experiment conditions and data analysis techniques. It will also discuss the design and condition repercussions on the possible organizations of controllers. Chapter 4 presents the results obtained from the experiments. Chapter 5 will provide a discussion and conclusions based on the results from the previous chapter, and outline possible future work.

Chapter 2

Background

Over the past several years, adaptive behaviour research has gained several new insights at methodological and philosophical level of inquiry. The reasons for these new insights, amongst which a fundamental change of view on the nature of cognition can be traced, are results of many contributing factors. Most of these factors will be covered in some detail in this section, because they in a radical way motivate this project.

2.1 Foundations

In his influential paper “What might cognition be if not computation”, van Gelder gives the following definition of cognition and cognitive science:

> Human behavior is, in the first instance, a matter of subtle interaction with a constantly changing environment. Cognition is the internal processing which underlies that interaction, and cognitive science is the study of that processing. (van Gelder, 1992, p. 10)

The orthodox view of cognition, that dominated the cognitive scientific endeavour, is known as classical computationalism (Language of Thought Hypothesis (Fodor, 1975), Physical Symbol System Hypothesis (Newell, 1980), Cognitivism (Haugeland, 1981)).

The Physical Symbol System Hypothesis (PSSH) states that

> The necessary and sufficient condition for a physical system to exhibit general intelligent action is that it be a physical symbol system.
>
> *Necessary* means that any physical symbol system that exhibits general intelligence will be an instance of a physical symbol system.
>
> *Sufficient* means that any physical symbol system can be organized further to exhibit general intelligent action.
>
> *General intelligent action* means the same scope of intelligence seen in human action: that in real situations behavior appropriate to the ends of the system and adaptive to the demands of the environment can occur, within some physical limits.
>
> (Newell, 1980, p. 170, original emphasis)

Physical Symbol Systems (PSS) refer to the broad class of systems physically realizable and capable of having and manipulating symbols. PSS belong to the class of universal machines, and is therefore equivalent to a Turing Machine. Representation[5], being another term to refer to a structure that designates, is central to PSSH, which in (Newell, 1980) is stated as

> The most fundamental concept for a symbol system is that which gives symbols their symbolic character, i.e., which lets them stand in for some entity. We call this concept *designation* ... (Newell, 1980, p. 156)

Thus, according to classical computationalists, cognition necessitates representation and can only be explained in the context of representations. Computation occurs at the symbolic level, i.e., in the context of representations.

5. This definition of representation is the one used in the PSSH. There are numerous definitions of representation, cf. Palmer (1978), Dorffner (1997), Sharkey (1997)

In classical computationalism, the mind is presupposed to have independent, disembodied[6], and self-contained existence. Critiques have been raised against this disembodiment, and the limits of purely symbolic models of mind, and it is in this context that the symbol grounding problem was formulated (Harnad, 1990). The symbol grounding problem is concerned with the following questions:

> How can the semantic interpretation of a formal symbol system be made *intrinsic* to the system, rather than just parasitic on the meanings in our heads? How can the meanings of the meaningless symbol tokens, manipulated solely on the basis of their (arbitrary) shapes, be grounded in anything but other meaningless symbols? (Harnad, 1990, p. 335, original emphasis)

Since PSS manipulate symbols found in internal representations, one of the most important contributions of embodiment[7] to AI is that it renders grounding of internal representations possible. Symbol grounding makes semantic interpretations of formal symbol systems intrinsic to the system, rather than parasitic on the meanings in our heads. This avoids the infinite regress which arises when meaningless symbols are grounded in other meaningless symbols. Symbol grounding also implies intrinsic environmental constraints based on non-symbolic (sensory) representations, besides the existing syntactical constraints. Embodiment[8] is further motivated by the following statement:

> The expectation has often been voiced that “top down” (symbolic) approaches of modeling cognition will somehow meet “bottom-up” (sensory) approaches somewhere in between. If the grounding considerations ... are valid, then this

6. “The body is another aspect of the outside world that can be theoetically set aside as we study cognitive processes.” (van Gelder, 1995, p. 233)

7. Embodiment presumes a body through which the real world can be experienced. The term is discussed in more detail later in this chapter.

8. Note that the role of embodiment differs for the classical computationalist and the enactive (or embodied) approach to cognition. For the classical computationalist approach, embodiment is merely a way of providing a link (i.e. a transducer) from the outside world to the symbols. For the enactive approach, embodiment is central as cognition is considered to emerge from physical agent-environment interaction.

> expectation is hopelessly modular and there is really only one viable route from sense to symbols: from the ground up. (Harnad, 1990, p. 345)

One of the more successful robot project based on the classical computationalist philosophy was the top-down Shakey project (Nilsson, 1984) from the late sixties at Stanford Research Institute, SRI. Shakey was based on the sense-model-plan-act (SMPA) or the functional decomposition framework (Brooks, 1991), where the role of sensing was exclusively to build an internal representation model of the real world. After building the internal representation, the planner could ignore the real world and instead use the representations to come up with an appropriate action. As sensor and motor activity were temporally and conceptually separate functions of the internal planner, they could subsequently be analysed and treated independent of each other.

In spite of the usage of the most powerful computers of that time and carefully designed simple environments, the sense-to-model step took immensely long time, because of the necessity of building accurate representations of the current situation required for successful behaviour. According to Brooks,

> [t]here was at least an implicit assumption ... that once the simpler case of operating in a static environment had been solved, then the more difficult case of an actively dynamic environment could be tackled. None of these early SMPA systems were ever extended this way. (Brooks, 1991, p. 570)

The central dilemma thus is how to correlate the internal representations with the physical changes occurring when an agent acts in the real world. This problem is referred to as the frame problem (McCarthy & Hayes, 1969). Questions were raised if the SMPA framework was really a realistic model e.g. when considering the time frames within which animate agents interact with their environments.

> There was a requirement that intelligence be reactive to dynamic aspects of the environment, that a mobile robot operate on time scales similar to those of animals and humans, and that intelligence be able to generate robust behavior in the face of uncertain sensors, an unpredicted environment, and a changing world. (Brooks, 1991, p. 570)

Following this inquiry, key realizations about the organization of intelligence emerged (some are mentioned in (Brooks, 1991)), which ultimately gave rise to the behaviour-based AI approach. Characteristic key aspects of this approach are situatedness, embodiment, intelligence, and emergence.

> **[Situatedness]** The robots are situated in the world - they do not deal with abstract descriptions, but with the here and now of the world directly influencing the behavior of the system. ... [The] world ... provides continuity ... [that] can be relied upon, so that the agent can use its perception of the world instead of an objective world model. ... *The world is its own best model.* (Brooks, 1991, original emphasis)

While abandoning the representationalist approach of "having representations of individual entities in the world", Brooks advocates spatial and functional representations (or deictic representations; cf. Agre, 1988; Chapman, 1990), "where the system has representations in terms of the relationship of the entities to the robot" (Brooks, 1991, p. 583). Beer also points out that empirical research in the context of situated agents, has shown that the computationally demanding and brittle task of planning can be significantly alleviated by the use of the immediate situation to guide behaviour (Beer, 1995). Arguments have been raised that deictic representations are not representations but rather registrations[9], since the world is not a pregiven reality, instead the agent plays a central role in creating its own world, i.e., it brings

9. A term which "[Smith (1996) refers to as] to do or be oriented towards the world *in such a way that it presents or arranges or constitutes itself as a world*." (Chemero, 1997, p. 46, original emphasis)

forth its world (cf. Varela *et al.*, 1991, Chemero, 1997). This view is elaborated further later in this section.

Another key aspect of the behaviour-based AI approach is embodiment:

> **[Embodiment]** The robots have bodies and experience the world directly - their actions are part of a dynamic with the world and have immediate feedback on their own sensations. ... [O]nly an embodied intelligent agent is fully validated as one that can deal with the real world. ... [O]nly through physical grounding can any internal symbolic or other system find a place to bottom out, and give 'meaning' to the processing going on within the system. ... *The world grounds regress.* (Brooks, 1991, original emphasis)

This bottom-up approach is also stressed by Wheeler, who emphasizes that cognitive science should proceed by studying embedded and situated controller architectures for simpler adaptive agent-environment interaction, which could provide guidelines on how more complex controller architectures should be approached (Wheeler, 1993).

Brooks also emphasizes that the overall behaviour does not need to have an internal locus of control:

> **[Emergence]** The intelligence of the system emerges from the system's interactions with the world and from sometimes indirect interactions between its components[10] - it is sometimes hard to point to one event or place within the system and say that is why some external action was manifested. ... Intelligence can only be determined by the total behavior of the system and how that behavior appears in relation to the environment. ... *Intelligence is in the eye of the observer.* (Brooks, 1991, original emphasis)

The most well-known alternative to the SMPA-framework which is based on the behaviour-based AI approach is the subsumption architecture (Brooks, 1986), in which overall behaviour of the agent is determined by (emerges from) the subsumption constraints on the out-

10. The components may for instance be layers in the subsumption architecture (Brooks, 1986; see below).

puts of multiple behaviour-producing layers. The mutually independent behaviour-producing layers work in parallel and are hierarchically organized, where higher levels (e.g. “Avoid objects”) have higher priority and can subsume lower levels (e.g. “Build maps”; see Figure 1).

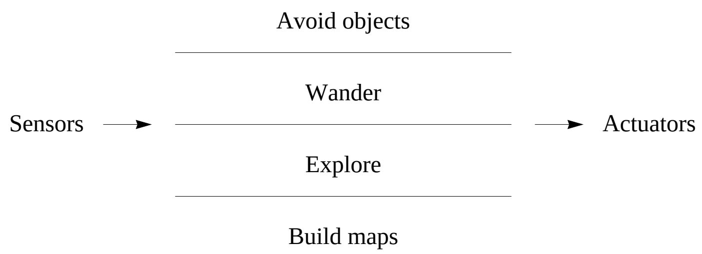

Figure 1: Behaviour-based decomposition in Brooks’ subsumption architecture.

This architecture captures the essence of the behaviour-based approach in that denies the use of central reasoning systems and (manipulable) representations, and focuses instead on decentralized processing.

Brooks further considers the agent and environment as coupled[11] (see Figure 2), which becomes apparent in the following passage:

> **[Intelligence]** [T]he source of intelligence is not limited to just the computational engine. It also comes from the situation in the world ... and the physical coupling of the robot with the world. (Brooks, 1991)

Brooks continues arguing, that

> ... the simple things to do with perception and mobility in a dynamic environment ... are a necessary basis for ‘higher-level’ intellect. Therefore, I proposed

11. “Two theoretically separable dynamical systems are said to be coupled when they are bound together in a mathematically describable way, such that, at any particular moment ... each system fixes the principles governing change in the other system.” (Wheeler, 1994, p. 10)

> looking at simpler animals as a bottom-up model for building intelligence. ... [W]hen ‘reasoning’ is stripped away as the prime component of a robot’s intellect, ... the dynamics of the interaction of the robot and its environment are primary determinants of the structure of its intelligence. ... *Intelligence is determined by the dynamics of interaction with the world.* (Brooks, 1991, original emphasis)

According to Wheeler, the dynamical coupling between agent and environment is *the* fundamental mechanism of situatedness. This dynamical coupling relation can be a subject of study if the level of abstraction is raised to a level, where the entire agent-environment coupling is regarded as one single system (see Figure 2), in which the patterns of interaction (i.e., behaviour) between agent-environment are properties of that system (Beer, 1995). Based on this observation, Beer developed a general theoretical framework for characterization of agent-environment interaction aiming at explanation and design of autonomous agents[12] (using dynamical systems theory, DST) by assessing the central problem for autonomous agents: to generate appropriate behaviour at the right time despite continuous internal (state) and external (environmental) changes (Beer, 1995).

12. An autonomous agent is in Beer’s definition “any embodied system designed to satisfy internal or external goals by its own actions while in continuous long-term interaction with the environment in which it is situated.” (Beer, 1995, p. 173)

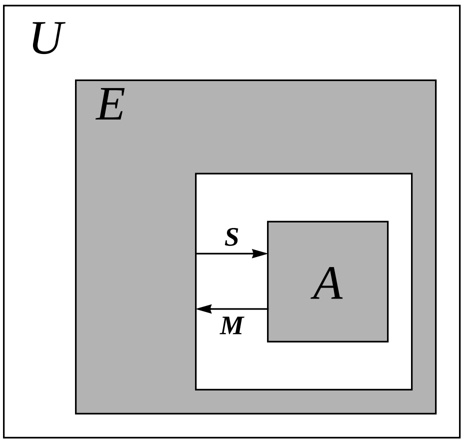

Figure 2: The coupled agent-environment system as one dynamical system[13].

In Figure 2, A denotes the agent, E the environment, $\boldsymbol{S}$ all effects that E has on A, and $\boldsymbol{M}$ all effects that A has on E, and U denotes the entire coupled system. The mutual influence between A and E is described in the following passage:

> Any action the agent takes affects its environment in some way through $\boldsymbol{M}$, which in turn affects the agent itself through the feedback it receives from its environment via $\boldsymbol{S}$. Likewise, the environment's effects on an agent through $\boldsymbol{S}$ are fed back through $\boldsymbol{M}$ to in turn affect the environment itself. (Beer, 1995, p. 182)

When considering the two coupled dynamical systems, A and E as one single dynamical system, U, the coupling relation between A and E implies that the dynamical property of U becomes a superset of the sum of the individual dynamical properties of A and E. U therefore exhibits a "richer range of dynamical behaviour" than could either system individually. Put it differently, in general, the dynamics of a coupled system differs radically from the same system uncoupled, therefore behaviour of a situated agent must be considered as a con-

13. "... [A] dynamical system [is] any system that we can effectively describe by means of an evolution equation; an equation which tells us how the state of the system evolves [over] time ..." (van Gelder, 1992, p. 15)

sequence of the agent-environment interaction, and not purely as a causal effect of its internal state alone, no matter how complex its internal dynamics (processes) may be (Wheeler, 1993).

> Behaviour is a feature of a system in which an environmentally-embedded agent and an agent-embedding environment evolve together through time, in a process of mutual and continuous feedback." (Wheeler, 1993, p. 15)

Once establishing the source of behaviour, the notion of adaptive fit was introduced to address the issue of what constitutes "appropriate" behaviour in an agent-environment interaction[14]. An adaptive fit exists between an agent and an environment as long as the interaction satisfies some given constraints C put on that interaction, or more informally, as long as the agent yields adequate performance on the task for which it was designed. The problem of how to design agents that exhibit certain behaviour is formulated in the synthesis problem, which, in Beer's framework is stated as:

> *The Synthesis Problem*. Given an environment dynamics E, find an agent dynamics A and sensory and motor maps ***S*** and ***M*** such that a given constraint C on the coupled agent-environment dynamics is satisfied. (Beer, 1995, p. 186, original emphasis)

The analysis problem is concerned with explaining the mechanisms underlying the behaviour, and in Beer's framework it is formulated as:

> *The Analysis Problem*. Given an environment dynamics E, an agent dynamics A and sensory and motor maps ***S*** and ***M***, explain how the observed behavior $\boldsymbol{M(x_A)}$ of the agent is generated. (Beer, 1995, p. 193, original emphasis)

Extending these ideas is the enactive approach, which views knowledge and embodiment as inseparable; knowledge is dependent on being in a world.

14. The notion is thus necessary for understanding the relationship between the agent and its environment.

> The central insight of this nonobjectivist orientation is the view that knowledge is the result of an ongoing interpretation[15] that emerges from our capacities of understanding. These capacities are rooted in the structure of our biological embodiment but are lived and experienced within a domain of consensual action and cultural history. They enable us to make sense of the world; or in more phenomenological language, they are the structures by which we exist in the manner of "having a world." (Varela *et al.*, 1991, p. 150)

As a natural consequence of the above argument, Varela argues against the representationalist approach by stating that

> ... scientific progress in understanding cognition will not be forthcoming unless we start from a different basis from the idea of a pregiven world that exists "out there" and is internally recovered in a representation. (Varela *et al.*, 1991, p. 150)

By considering cognition as being embodied action, Varela *et al.* take a position between idealism, or cognition as projection of the inner world, and realism, or cognition as recovery of pregiven outer world (the position of classical computationalists). Action is seen as inseparable from perception (as sensor and motor processes are inseparable), and the definition of action becomes perception *and* action, thus radically different in this context to the earlier definitions (e.g. in the SMPA framework). Based on these insights, the enactive approach to cognition was introduced with the overall concern of determining "common principles or lawful linkages between sensory and motor systems that explain how action can be carried perceptually guided in a perceiver-dependent world." (Varela *et al.*, 1991, p. 172) The key aspects of this approach are stipulated as

> ... (1) perception consists in perceptually guided action and (2) cognitive structures emerge from the recurrent sensorimotor patterns that enable action to be perceptually guided. (Varela *et al.*, 1991, p. 172)

15. The notion of interpretation in this context refers to "the enactment or bringing forth of meaning from a background of understanding." (Varela *et al.*, 1991, p. 149)

In the enactive approach, perceptually guided action concerns the study of how a perceiving agent can guide its actions in a perceiver-dependent world which is constantly changing, partly as a consequence of the agent's own activities.

The starting point for understanding the nature of perception thus becomes a matter of understanding the sensorimotor structures of the perceiving agent, since these structures resemble the way in which the agent is embodied.

2.2 Related work

The adaptive behaviour community, considering highly diverse problems (e.g. obstacle avoidance, legged locomotion) and approaches (e.g. top-down/bottom-up, computational/ dynamical), is unified by the common aim of understanding the nature and the underlying mechanisms of adaptive behaviour in natural and artificial agents. A particular subgroup of this community employ ANNs and combine them with evolutionary algorithms for evolving autonomous agent controllers for adaptive behaviour (e.g. to evolve complete architectures (Cliff, 1992), or ANN weights (Beer, 1996)). The motivation for employing evolutionary algorithms comes from the realization of three major problems associated with the design of complex control systems:

> It is not clear *how* a robot control system should be decomposed.
>
> Interactions between separate sub-systems are not limited to directly visible connecting links between them, but also include interactions mediated *via the environment*.
>
> As system complexity grows, the number of potential interactions between subparts of the system grows *exponentially.*
>
> (Harvey *et al.*, 1996, p. 1)

Evolutionary robotics minimizes the a priori assumptions by employing the use of evolutionary techniques for incremental evolution of increasingly complex robot systems, where the only benchmark is the overall behaviour (Harvey *et al.*, 1996). It thus alters the human designer's role to focus on *what* the robot's task is rather than on *how* the controllers should be designed for exhibiting behaviour coping with the task.

Although generally not explicitly concerned with cognition as such, adaptive behaviour research has implications to cognitive science, not just in providing models of "intelligent" and adaptive behaviour, but also to confirm or question the validity of different theoretical frameworks and subsequently also their philosophical underpinnings.

One relevant question for an adaptive behaviour researcher aiming to contribute to cognitive science is therefore what sort of behaviours are the most appropriate for investigation given that they should be complex enough to raise cognitively interesting issues while remaining computationally tractable. Beer uses the term "minimally cognitive behaviour" to refer to behaviours having these properties. According to Beer, visually-guided behaviour is an excellent example since it includes phenomena such as "visual orientation, object perception, and discrimination, visual attention, perception of self-motion, object-oriented action, and visually-guided motion and manipulation" (Beer, 1996, p. 422).

There are numerous approaches to tackle the problems related to visually-guided behaviour. Despite the apparent differences of the considered AI approaches, they all, with the exception of SMPA, subscribe to the active perception hypothesis, which states that perception is an activity performed by an autonomous agent in the context of some adaptive behaviour, i.e., the indivisibility of perception and action (Wheeler, 1994).

The work presented here is based on Beer's work on visually-guided behaviour (Beer, 1996). The main motivation behind the latter work was to develop simple models of adaptive behaviour "for the purpose of elucidating the essential principles of a dynamical theory of adaptive behavior" (Beer, 1996, p. 422). The problem was approached by studying possible organizations of evolved internal dynamics of controllers employed in simple agents that exhibit visually-guided behaviour. The robot controller gained its dynamical properties through an artificial neural network (ANN) which provided a mapping from sensors to motors.

The experiment set consisted of three sets of experiments for visually-guided behaviour. All the experiments were designed with respect to the synthesis problem and the analysis step involved the mean fitness of the agents and qualitative assertion of their exhibited behaviours.

The first set of experiment is of immediate relevance for this project and is therefore considered in detail. It involved the use of agents with one degree of freedom motion capabilities (i.e., the agent could move along a line), whose task was to orient towards visual stimuli to catch a falling object (see Figure 3). Simple evolution strategies (ES) were used for evolving individuals, each of which were encoded as a vector of real numbers which represented the ANN connection weighs, unit biases, time constants (for the recurrent ANN), and gains[16] for the (sensor) input units. A population of 25 individuals was evolved over 50 generations.

16. The gains are values that correspond to sensor input difference sensitivity. "If the gains are too low, each [sensor] ray will show very little difference in response regardless of its length. If the gain is too high, each ray will essentially give a binary response at a very narrow range of distances..." (Beer, 1996, p. 423)

The agents employed five sensors uniformly distributed over an angle of $30°$ with the range of 220 [17] in an environment, whose dimensions were 400×275.

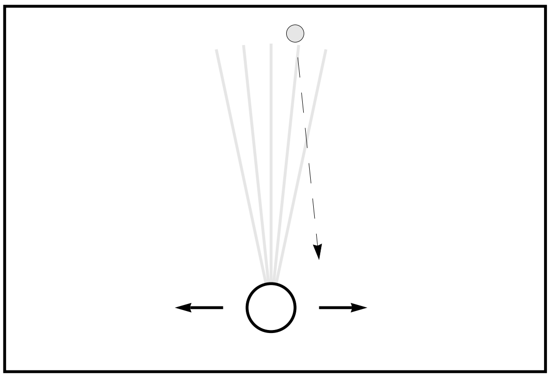

Figure 3: Beer's experimental setup for orientation experiments.

The task for the agent to handle was simplified by only allowing at most eight possible "trajectories" for falling objects. The falling object was a diameter of 26 and was dropped from the top of the environment with vertical speed in the range [0.5,5], horizontal speed in the range ± 6, and offset from the agent's centre in the range of ± 70. The fitness function to maximize was

$$200 - \frac{\sum_{i=1}^{NumTrials} d_i}{NumTrials} \qquad \text{(EQ 1)}$$

where $NumTrials$ was the total number of evaluation runs per individual (six to eight in the orientation experiments), and d_i was the horizontal distance between the agent and object centres at the *i*:th evaluation run when the object reached the same vertical position as the agent. Agents sensitive to objects with large horizontal and small vertical speeds used sensor

17. No unit of measurement was specified by Beer.

biases within the range -0.85 ± 0.18, whereas for agents sensitive to small horizontal and large vertical speeds, the sensor biases were within the range -2.06 ± 0.33.

Initially, the controllers employed feedforward ANNs, the deficiencies of which became apparent as the evolved agents only managed to orient towards objects following certain trajectories. The agents were unable to orient towards some objects following trajectories with large horizontal and small vertical speeds. As these falling objects passed out from the agent's sensor range (i.e., field of view), successful orientation was no longer possible using agents with feedforward ANNs, since these ANNs are purely reactive, requiring immediate sensory information. This deficiency was overcome by altering the ANN dynamics by employing a recurrent ANN architecture (Continuous Time Recurrent Neural Networks, CTRNNs). The internal dynamics of the recurrent ANNs allowed the agents to continue the object pursuit despite the object passing out of the agent's sensor range. The agents thus managed to handle a simple case of the object persistence problem.

The second set of experiments involved visual discrimination. The aim was to evolve agents for discriminating between three kinds of objects: circles, diamonds, and lines. The actual discrimination task was to catch the circular objects, while avoiding the non-circular objects. All evolved controllers exhibited dynamic object-pattern matching by active scanning through different foveation strategies. Foveation is a term for eye movement (for static-sensor agents: agent body movement) so that e.g. an object reaches the fovea, the eye-region which provides high acuity[18] vision in bright light (Bruce & Green, 1990).

18. Visual acuity is the angle between adjacent bars in the highest frequency (density) grating (i.e., parallel black and white bars of equal length) distinguishable from a plain field of the same average brightness as the grating. (Bruce & Green, 1990)

In the context of agents with very limited number of sensors, foveation is an effective method of increasing the sensors spatial resolution by altering position so that the sensor rays are brought to intermediate positions with respect to their positions before the movement. It is advantageous not only since it facilitates the discrimination task by increasing the resolution, it can also be utilized to guide the agent to move in a standard position with respect to the object.

All of these beneficial abilities were exploited by the evolved agents, which exhibited foveate-and-decide, and antifoveate-and-decide strategies. The former was initiated by foveation, which after successful discrimination turned to visual object orientation in the case of circular objects, and object avoidance in the case of non-circular objects. The antifoveate-and-decide strategy was initiated by antifoveation, where the object bearing was set to outermost (peripheral) sensors. After successful discrimination, the agent followed the same decide-strategy as in the foveate-and-decide strategy.

The last experiment set, concerned visual manipulation, where near-optimal handling evolved for the task of coordinating the movement of a one degree of freedom arm, with objects appearing in the visual field of the stationary agent. Because of its marginal relevance to current work, this experiment will not be considered any further.

There are several other approaches to visually-guided behaviour, some of which will be considered in the following exposition.

The most closely related work is done at COGS, University of Sussex (cf. Harvey *et al.*, 1996), where researchers focus on the question of how difficult it is to evolve internal dynamics of agents which exhibit desired behaviours. Evolutionary techniques are employed to incrementally evolve robot controllers with increasing complexity. The bottom-up approach taken allows complete agent network architectures, rather than just connection weights, to evolve using genetic algorithms (GAs) with increasing genotype lengths.

Recognizing the dynamical coupling in agent-environment interaction giving rise to the agent's behaviour, they argue, that the control system must itself be a dynamical system. The argument used as motivation for employing ANNs is that these are convenient forms of dynamical systems, some of which (CTRNNs) in principle can exhibit arbitrary degree of accuracy in replicating any other dynamical system with finite number of components (Harvey *et al.*, 1996).

Two sets of experiments conducted at COGS are of immediate interest for this project. The first set of experiments, a variation of the orientation experiments described above, concerned evolution of control systems for an agent whose task is to move to the center of a simulated circular arena. The controller used a specific neuron model for the nodes, with a fixed number of input and output nodes and arbitrary (evolvable) number of intermediate nodes and (evolvable) connections for arbitrary complexity. For introducing as little a priori assumptions as possible, the sensor acceptance angles and positions were also evolved. Two different successful controllers evolved, each of which positioned itself at the centre of the arena while spinning; either on the spot or in a minimum radius circle. The two controllers thus produced convergent behaviours, but utilizing divergent mechanisms for producing these behaviours. Although most analysed controllers lacked identifiable structure, occasionally controllers evolved which vaguely resembled a two-layer subsumption architecture[19] for visual guidance, in which spinning behaviour was subsumed by an approaching behaviour (Cliff, 1993).

A dynamical systems perspective was applied for analysis of the evolved networks, which revealed that the robot, despite presence of noise, is guaranteed to succeed at its task.

The second set of experiments concerned real-world discrimination and used the same kinds of specific networks as those employed in the first set. An incremental evolutionary methodology was employed, in which initially simple visual tasks and environments were used as a basis for proceeding to more difficult ones. The controllers of individuals succeeding in simpler tasks were thus chosen as basis for evolving controllers for more difficult tasks thus allowing to build competencies on the top of earlier ones in an integrative rather than modular way. Although the controllers evolved for orientation towards large objects performed very poorly in orientation towards a small object at the same position, once evolved correctly, the controllers were able to handle the generalized version of the static target orientation, namely orientation towards a moving object. The task difficulty was yet again increased to include discrimination between a triangular and a rectangular object. The evaluation function was altered to benefit closeness to the triangular object, and to penalize closeness to the rectangular object. This case, the evolution process took benefit of previously evolved small target orientation behaviour, and evolved triangle-discrimination very rapidly. The evolved agent now employed two visual sensors simultaneously for visual discrimination rather than one as in visual orientation.

19. The layers and the subsumption constraints corresponded to specific neurons. For details about Brooks' subsumption architecture, refer to section 2.1.

A variation of the orientation experiment was also conducted by Floreano & Mondana (1994), who employed GAs for evolving Simple Recurrent Networks (Elman, 1990) which served as controllers for a real robot. The task of the robot was to stay operative in its environment as long as possible given that its (simulated) batteries were linearly discharging with time. When moving to a specific (dark) region in the environment, the batteries were automatically and instantaneously recharged. The successfully evolved controllers exhibited behaviours for wandering around the environment while avoiding obstacles when batteries were charged. As the batteries reached a minimum threshold value, the robot automatically switched to a behaviour for orientation toward the recharging area. The threshold value was autonomously established by the controller and varied with distance to the recharging area.

Floreano & Mondana claimed the success on the orientation was due to exploiting the light gradient information. This claim is questionable when considering their subsequent observation that once the black paint was removed from the region, the robot, triggered by the threshold value, oriented towards the same area, and "started to explore its surroundings without leaving that zone until its battery was completely exhausted." (Floreano & Mondana, 1994). This task can thus be seen as a orientation experiment where temporal factors matter.

Whereas all the different approaches discussed above are bottom-up, there are also top-down approaches to visually-guided behaviour. While the bottom-up approaches make few a priori assumptions, the top-down approach is characterized by the deliberate a priori controller design, and is often employed when a cognitive model is to be validated.

A top-down approach is taken at AILab, University of Zürich (Scheier & Pfeifer, 1995b), where a categorization process is the outcome of a value map-guided interactions of a haptic and a visual system, where each multimodal system is carefully designed. In another AILab project (Scheier & Pfeifer, 1995a), a variant of Brooks' subsumption architecture is used, where the different layers correspond to simple processes (surprisingly, "obstacle avoidance" is considered as a simple process) rather than behavioural modules. Despite emergent classification, this approach cannot be considered as bottom-up, for several reasons. First, "obstacle avoidance" and "move along object" are considered as simple processes (despite clear rejection of behaviours at different layers) and are fixed a priori (hard-wired) in the controller. Second, classification behaviour is aided by a parameter, which encodes (by value addition) the object size when a "move along object"-process is active.

Whereas the both AILab projects view categorization as sensory-motor coordination, SMPA-framework approaches would have considered categorization as isolated perceptual subsystem. Briefly, the major problem of approaches utilizing this framework is concerned with recovery of sensor information into an accurate world model (representation). The recovery is done using the snapshot model, where sensory snapshots are taken of the environment, and through top-down inferential and search processes, relevant details of the environment are represented. The process of recovery from a single image is not just computationally demanding (cf. Brooks (1991)), but also in some cases impossible (e.g. consider partly hidden objects). Furthermore, the real-world time scales are just not in favour of the SMPA-framework, since the representation should in general be already inaccurate when completed, due to the dynamically changing real-world environment. The snapshot model that the SMPA-framework employs has also been rejected in a number of works (Wheeler, 1994). Since the approach to visually-guided behaviour using snapshot model and the SMPA-framework differs radically from this work, it will not be considered any further.

Chapter 3

Materials and methods

This chapter concerns the materials and methods used in this project and spans the following issues: design decisions and strategies, experimental conditions, tools for data collection and analysis. The last section discusses the repercussions that the design decisions, and the experimental conditions might have on the possible behavioural exhibition.

3.1 Design decisions and strategies

The design decisions concern controller design and choice of evolutionary algorithm. The overall design of the controller system was based on evolution of the internal dynamics of controllers by using evolutionary algorithms for finding appropriate ANN weight value settings, where the ANN is the central component in the controller architecture, responsible for the mapping of sensory input to motor output. Each agent (individual) thus used a controller with fixed architecture and fixed (evolved) weight settings, and therefore no on-line learning process was utilized in the agent.

The basic controller architecture was comparable with Beer's controller architecture but where the ANN architecture has been modified. The main motivation for the modification was to use a simpler ANN architecture while preserving the ANN dynamics, since ANN dynamics proved to be of importance for the experiments (Beer, 1996).

One ANN architecture which satisfies these requirements is the Simple Recurrent Network, SRN (Elman, 1990). The experiments were carried out using a fixed topology 5-3-2[20] SRN with a sigmoid unit activation function[21], where the five input units were dedicated for sensory input, and the two output units were dedicated for direct motor control of left and right motors. The sensor inputs and motor outputs were normalized according to the following scheme: the ANN inputs ranged from 0 (out of range) to 10 (close), and the ANN outputs were in the range of 0 (zero speed) to 1 (full motor speed forward). Refer to section 3.2 for further details about the sensors and motors.

The evolution algorithm that promised to be the best choice for evolving ANN weights were evolution strategies, ES (Rechenberg, 1973; Bäck & Schwefel, 1993). ES have proved to be superior to genetic algorithms (GAs) in global convergence performance in presence of epistatis; the parameter[22] interaction with respect to an individual's fitness function.

> Results of other research (Salomon, 1996b; Salomon, 1996c) strongly indicate that the independence of parameters is an essential prerequisite of the GA's high global convergence; the presence of epistatis drastically slows down convergence. The problem in the context of autonomous agents is that the parameters of control systems for real world applications are not independent. (Salomon, 1996a, p. 1)

Since ESs do not have these shortcomings but behave invariant in presence of epistatis (Salomon, 1996a), it should subsequently be more a appropriate choice. ES was also chosen

20. x-y-z stands for x input, y hidden, and z output units
21. S-shaped function defined as $\frac{1}{1+e^{-x}}$
22. The parameters refer to the arguments of the function to optimize (minimize or maximize), e.g. real-valued ANN weights.

by Beer for ANN weight evolution (Beer, 1996). The ES employed in the experiments followed Schwefel's (Bäck & Schwefel, 1993) empirically found suggestions for mutation, recombination, and selection.

The variant of ES employed in this project maintained a population of individuals (called parents and children), each consisting of a vector of n real valued numbers encoding for object variables (ANN weight values), which are variables coding for the function to minimize (or maximize, depending on the fitness function), and a strategy parameter consisting of one (1) standard deviation (no rotation angles were taken into consideration). The strategy parameter was adapted by mutations (this is also known as self-adaptation). The object variables were then mutated according to the new standard deviation. Mutations of individuals were carried out according to the following equations:

$$\sigma'_i = \sigma_i \cdot \exp(\tau' \cdot N(0,1) + \tau \cdot N_i(0,1)) \qquad \text{(EQ 2)}$$

$$\vec{x}' = \vec{x} + \vec{N}(\vec{0}, \vec{\sigma}') \qquad \text{(EQ 3)}$$

Where σ_i denotes the standard deviation for individual i, $N(0, 1)$ is a normally distributed one-dimensional random variable having expectation zero and standard deviation 1, and $N_i(0, 1)$ denotes that a new random variable of the above kind is generated for each new i. $\vec{N}(\vec{0}, \vec{\sigma})$ denotes the random vector distributed according to the generalized *n*-dimensional normal distribution having expectation $\vec{0}$ and standard deviations $\vec{\sigma}$. $\vec{x}$ denotes the vector of object variables, Schwefel suggests (Bäck & Schwefel, 1993, p. 4) that the factors τ and τ' should be set as:

$$\tau \propto \frac{1}{\sqrt{2\sqrt{n}}} \qquad \text{(EQ 4)}$$

$$\tau' \propto \frac{1}{\sqrt{2n}} \tag{EQ 5}$$

When producing new individuals (children) from the parent population, different mechanisms can be employed within ES (see Bäck, 1993 for an overview). In this project, a combination of three of the recombination (crossover) mechanisms were used for recombining object variables for creating the new individuals. The chosen mechanisms were similar to Salomon's (Salomon, 1996):

$$\text{without recombination: } x'_i = x_{S,i} \tag{EQ 6}$$

$$\text{(local) discrete recombination: } x'_i = x_{S,i} \text{ or } x'_i = x_{T,i} \tag{EQ 7}$$

$$\text{(local) intermediate recombination: } x'_i = x_{S,i} + \chi \cdot (x_{T,i} - x_{S,i}) \tag{EQ 8}$$

S and T denotes two randomly selected individuals from the parent population, and $\chi \in [0,1]$ is a uniform random variable.

The motivation for this was that, despite Schwefel's suggestions that intermediate recombination gave best results on empirical observations (Bäck & Schwefel, 1993, p. 4), they may not be as successful for the domain of optimizing ANN weight settings as for "conventional" parameter optimization. Specific to this domain is that the ANN works holistically, which requires that typically every weight is tuned to many (or all) other weights and the performance of the ANN is completely dependent on the tuning. Since different weight settings, as long as they are tuned correctly, can realize the same function approximation, ANN weights are not suited for interchanging between networks resulting in a network whose weights are a mixture of weights from two/more networks (cf. Meeden, 1996). In the majority of cases such an interchange would lead to worse function approximation than possible with any of the component networks. There are several proposed techniques for ANN weight recombination which take into consideration some of the properties of ANNs required for successful recombination (for some different GA techniques, see Mitchell, 1996). No such techniques were employed in this project, which might have had negative effects on the global convergence rate.

For the strategy parameter, a global intermediate recombination strategy was employed, as suggested by Schwefel in (Bäck & Schwefel, 1993, p. 4) where

$$\text{(global) intermediate recombination: } x'_i = x_{S_i}, i + \chi_i \cdot (x_{T_i}, i - x_{S_i}, i) \qquad \text{(EQ 9)}$$

which is similar to its local variant, besides that S_i, T_i, and χ_i are determined anew for each component of $\vec{x}$.

The deterministic selection strategy employed throughout all experiments followed Schwefel's recommendation of (μ, λ)-ES (Bäck & Schwefel, 1993) in which the μ best individuals are selected from λ children to form the new parent population. The population ratio was also based on Schwefel's $\mu / \lambda \approx 1/7$ using a population of 5 parents and 35 children in all experiments documented here.

The very simple fitness function to minimize, Φ,was similar to Beer's, although it considered the Euclidean distance between the robot and the object, since the robot had two degrees of freedom. The following fitness function was used in all experiments:

$$\Phi(a_i(t)) = \frac{\sum_{s=1}^{5} \sqrt{|rx_s - ox_s|^2 + |ry_s - oy_s|^2}}{5} \qquad \text{(EQ 10)}$$

$\Phi(a_i(t))$ is the fitness value for individual i at generation t, s is the evaluation trial, rx, ry are the horizontal and vertical positions of the robot respectively, ox, and oy the corresponding positions for the obstacle. Five evaluation trials were used for each individual, which means that the fitness values averaged over five runs determined the individual's fitness. The fitness value was only measured at the last time step of each run.

Chapter 4 includes details on how the fitness was determined for the individuals, and on additional constraints imposed on the fitness evaluation employed in some experiments, together with their motivation.

3.2 Experimental conditions

All the experiments were conducted in an extensively modified version of the Khepera Simulator (Michel, 1996), with the initial motivation to allow real-world validation on a Nomad 200 robot. The simulator was modified, so that the Khepera sensors would resemble Nomad 200 sonar sensors which required extension of the number of sensors from eight to 16, and the sensor range (corresponding to 660[23] points (cf. Figure 4)). The sensor beam width was slightly lower in the simulator than for the real Nomad; approximately $18°$ instead of the real Nomad's $25°$. The accuracy (scan point density) of the simulated sensors were left the same as for the simulated Khepera sensors. Without the sensor range extensions, the robot would end up looking for objects and eventually, just prior to bumping into them, detect their presence. The modifications also contributed to make the simulated robot's sensor configuration more alike that of the agent used by Beer.

23. 1 point corresponds to approximately 1 cm (9.7 mm) in the real world.

Despite the modifications, Beer's robots ended up being superior when it comes to sensor density (five to seven sensors each $30°$, compared to modified simulator's one each $22.5°$) and accuracy. The Khepera simulator takes into account noise on sensory readings and on motor actions, which was preserved in the modified simulator to achieve maximal fit with real world conditions[24]. A $\pm 5\,\%$ random noise was added to the sensor readings, $\pm 10\,\%$ to the (amplitude of the) motor speed, and $\pm 5\,\%$ to the direction (Michel, 1996).

The sensor inputs ranged from 0 (close) to 255 (far away), while the translational (speed) limitations of the motors ranged from -10 (full speed backward) to 10 (full speed forward).

The orientation experiments ran in a simulated 960×960 point environment surrounded by walls. In this environment, the diameter of the robot corresponded to 55 points. Maximum translational speed allowed (Euclidean) distance covering corresponding to approximately 4.5 points each time step.

In the experiments, a rhombus with the dimensions 25×25 points, was randomly positioned within an area of 500×500 points in the centre of the environment (activity area) to simplify the orientation task (cf. Figure 4).

24. Due to lack of time, no real world validation were carried out, despite availability of a real Nomad 200 robot and easy portability of available controller code to the Nomad 200.

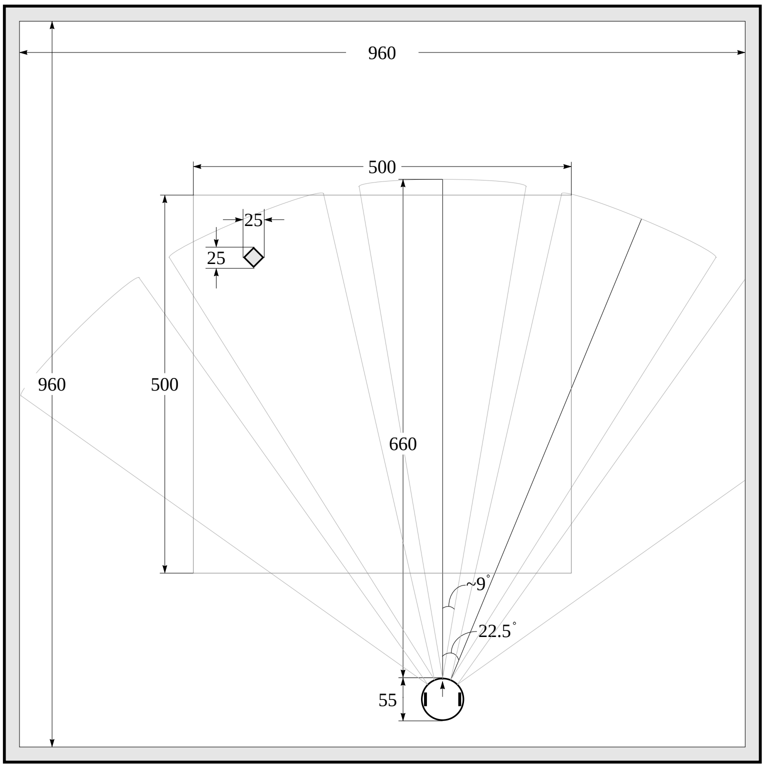

Figure 4: Dimensions of the environment, the agent, and the object.

The robot's position was determined anew at each run by randomly placing it at a distance of ten robot diameters from the object position in the environment. The heading of the robot was also determined anew (set to a random value) at each run. The evolution was thus not aided by having the obstacle within a certain maximal displacement (or offset) from the robot's heading[25].

In each run, the robot was allowed to move around until it collided with the object, or maximally 1500 time steps. This corresponds to maximal distance covering of approximately

25. In Beer's (1996) experiments for visually-guided orientation, the object's horizontal offset from the robot's center was at most 70 for the orientation experiments.

6450 points measured in Euclidean distance. Higher speeds resulted in higher area coverage at the expense of lower sensory coverage of the environment.

3.3 Data collection and analysis

Data was collected after a certain number of evolution runs when the individuals systematically succeeded in completing the task (having high fitness). The data collected included complete (normalized) sensory input data, data for all weights and unit biases, as well as all activation and output values. Robot position and robot heading (angle) was also saved for each step in the run. The collected data provided material for extensive analysis. External properties like qualitative behaviour of the robot were analysed using many runs and observations of trajectorial similarities for different individuals (or same individual with differing initial conditions). Internal properties like internal representations of the ANNs (hidden unit activations) were analysed via hierarchical cluster analysis, and additional visualisation was made using Gnuplot (Kelley & Williams, 1993).

3.4 Condition and design repercussions

This dissertation considers a bottom-up approach to visually-guided behaviour. Care has been taken to eliminate, to the largest extent possible, deliberate design decisions based on potentially misleading assumptions in order to avoid imposing unnecessary or too restrictive constraints.

The following list describes how the design decisions and strategies, and the constraints which these and other factors impose, impinge on the range of possible controller organizations:

1. The evolutionary algorithm is responsible for shaping the internal dynamics (ANN) of the controller. The algorithm does so by using heuristics (originally based on ideas of Darwinian evolution) to achieve faster than purely random convergence of controller dynamics by employing recombination and mutation techniques. The fitness function is used to guide the evolutionary algorithm since it determines which controllers are more appropriate. As this function can be based on the behavioural level (i.e., *what* behaviours are desirable) it transfers the task of internal controller design for desirable behaviour (i.e., *how* to realize these behaviours) to the evolutionary algorithm.

2. The limitations of the fixed ANN architecture constrain the possible complexity of the internal dynamics, thus the possible behaviours the robot can exhibit.

3. As the sensors and effectors (motors) of the robot have limited capacities, these constrain the possible behaviours the robot might exhibit, e.g. the robot can not start to orient to an object when, for instance, the object is out of range.

4. The number of (discretized) moves the robot was allowed to take was determined a priori and thus also imposed constraints on possible controller organization.

5. Finally, as the experiments were conducted in an environment surrounded by walls, which also restricted the possible organization of the controllers.

Chapter 4

Results

This chapter presents the results obtained from analysis of the orientation experiments described in Chapter 3. The results are based on two experiments, one for each type of controller that evolved, but are representative for the both types of controllers. Before moving on to the details of the results, a few operational definitions need to be stipulated.

4.1 Definitions

The following notation will be used in this chapter to refer to the different robot sensors (see Figure 5): leftmost sensor (LL), left sensor (L), centre sensor (C), right sensor (R), and rightmost sensor (RR).

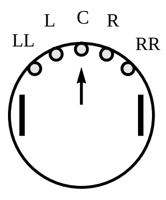

Figure 5: The robot sensor notation

In this chapter, the term “drop” refers to a sudden sensor reading decrease (e.g. 5 at time $t-1$, and 0 at time t) whereas the term “peak” is used to denote a sudden sensor reading increase.

4.2 The O1-type controller

The “minimal configuration” O1-type controllers were the most common controllers that evolved during the conducted experiments. The controllers evolved to exhibit good orientation performance within the first hundred generations. In the following, a representative O1-type controller, hereafter referred to as “O1” will be examined during an experimental run, E1. Figure 6 shows the averaged parent population fitness at each generation, and the trend, being the averaged parent population fitness over all generations for O1 during E1. The fitness remained unstable but close to optimal even when the number of generations exceeded one hundred, which is typical for the O1-type controllers. For details on how the fitness was measured, refer to Chapter 3.

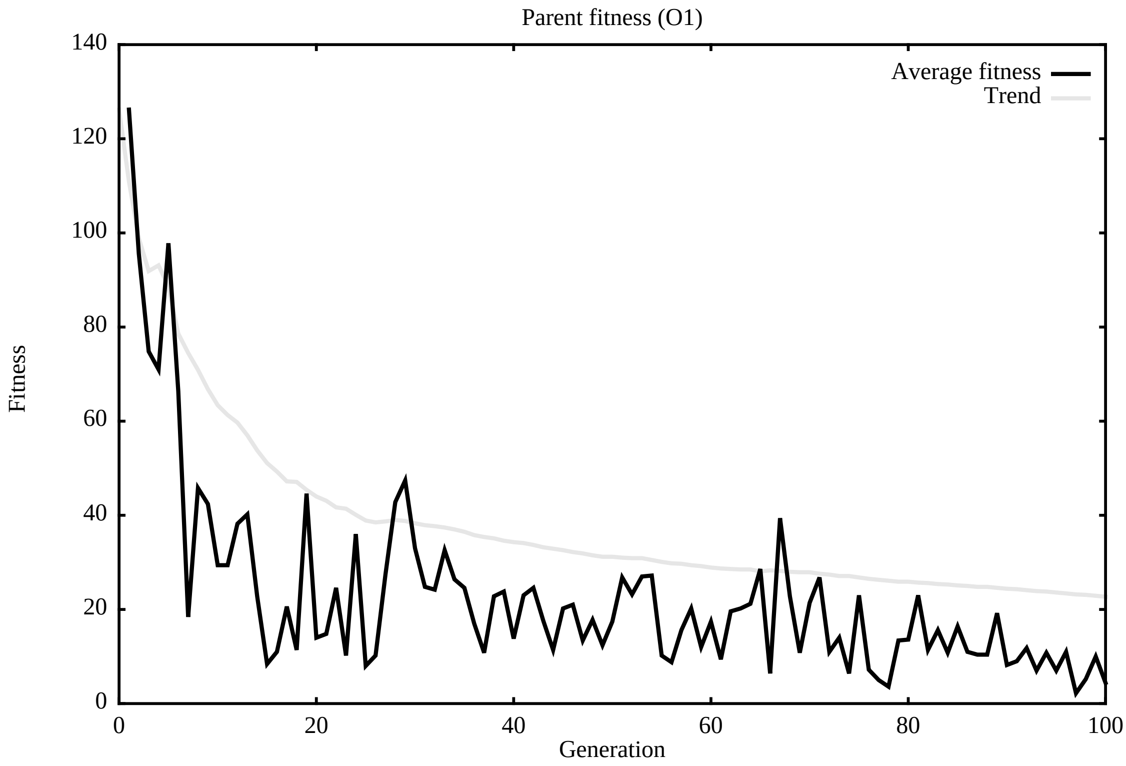

Figure 6: The O1 controller parent population fitness.

Another characteristic of the O1-type controllers is that they make use of only one of the five possible sensors for handling the orientation task. In all observed cases one of the outermost sensors (LL or RR) was used to direct the orientation towards an object. This sensor will hereafter be referred to as the "dedicated sensor". Subsequently, there was a general tendency for a turn in the same direction as the dedicated sensor during virtually the entire time. The rare occasions when the opposite motor activation was the highest, it resulted in forward motion because of the small activation differences for the two motors. As a consequence, either forward motion or turn in direction of the dedicated sensor was employed. O1 employed the leftmost sensor as dedicated sensor. Accordingly, there was a general tendency to turn left. Figure 7 shows the trajectory and, for each time step, the angle of heading of O1 during experiment E1. Refer to Appendix A for the corresponding sensor inputs and motor outputs.

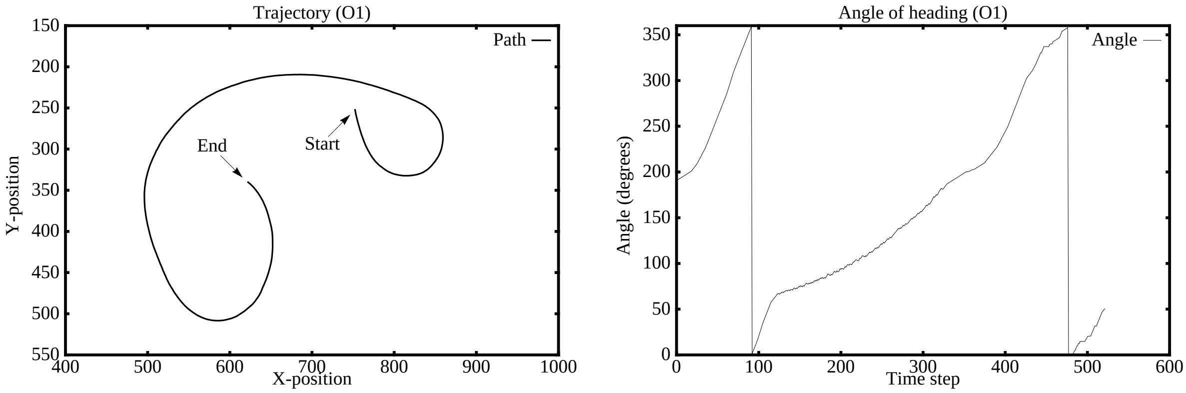

Figure 7: The trajectory and heading for O1 during E1.

Figure 8 shows the ANN employed by O1. The ANN has been split up in two parts, the left part showing the input (bottom), hidden (middle), and output (top) units, and the weights between them, while the right part shows the context (bottom), hidden (middle) and output (top) units, with corresponding inter-unit and unit bias weights. Solid lines denote excitatory connections, while dotted lines indicate inhibitory connections. The thickness of the connections refer to the associated magnitude (thick for strong connection, thin for weak connection). For exact weight and unit bias values for O1, refer to Appendix A.

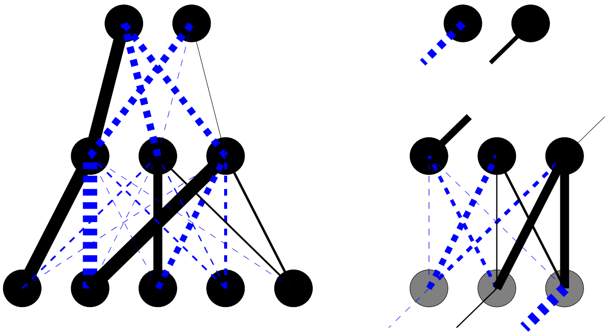

Figure 8: The O1 controller ANN.

The ANN employs recurrent connections which makes it difficult to analyse. The analysis proceeded by correlating the hidden unit activations with the behaviour of O1 in the experimental run, E1.

A close examination of the hidden unit activation reveals a partitioning according to turn sharpness and direction. Turn sharpness refers to the difference in motor speeds, where bigger difference leads to lower radius turn. Figure 9 shows the hidden unit activation partitioning according to turn sharpness and turn direction. The figure indicates the difference between the two motor outputs, where the marked area corresponds to right turn, whereas the unmarked corresponds to left turn.

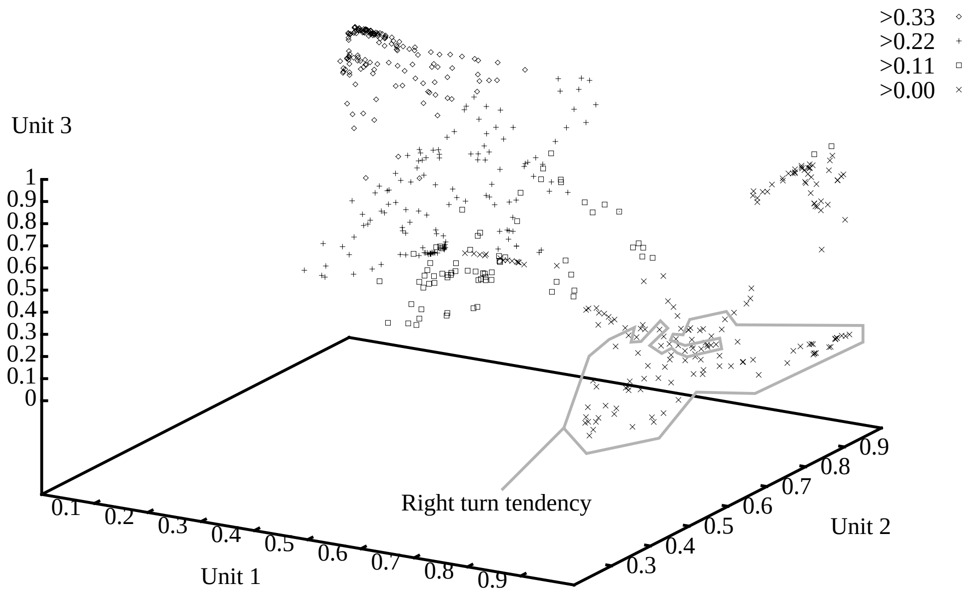

Figure 9: Hidden activation partitioning: turn sharpness and direction (O1).

Note in Figure 9 that the right turn tendency results in forward motion as the turn sharpness is low. In Figure 8, a high activation in the leftmost sensor inhibits the general left turn tendencies and amounts in forward motion as a result of the strong connection between the unit corresponding to the leftmost (dedicated) sensor and the unit corresponding to left motor output.

Hierarchical cluster analysis on the hidden unit activations of O1 during E1 also revealed that there are clusters corresponding to different behaviours employed during the run. The behavioural clustering can be seen in Figure 10, where the different markings refer to different behaviour clusters. Note that the figure only includes the seven biggest clusters (in the markings, the roman numeral after the initial letter indicates the cluster depth level in the dendogram[26]).

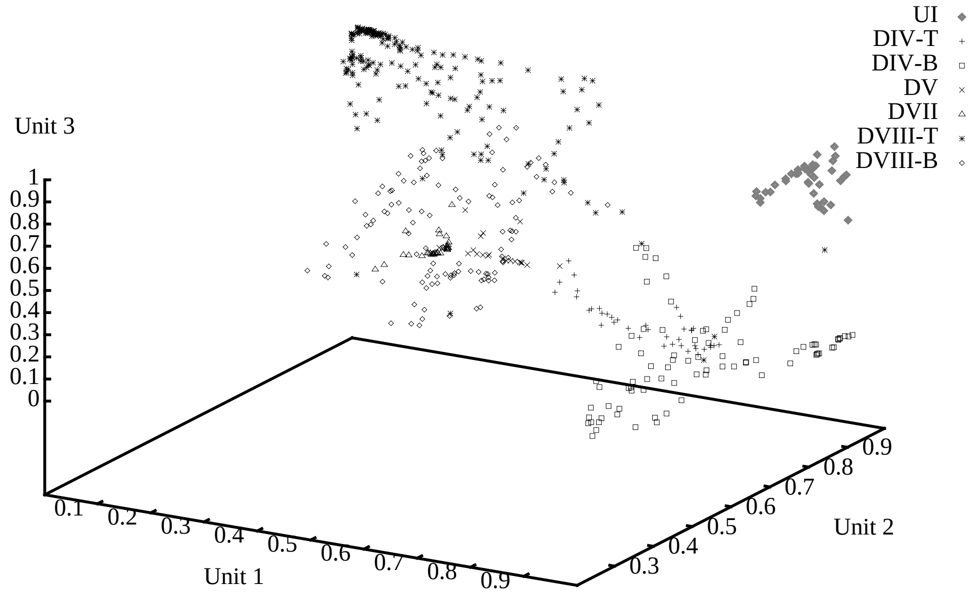

Figure 10: Behavioural clustering of hidden unit activations for O1 during E1.

While the whole behaviour of O1 during run E1 was organized in eleven clusters (sub-behaviours), the overall behaviour can be summarized according to the following key behaviour categories:

1. Behaviour for normal functioning and obstacle avoidance, where the former refers to moving around in the environment and amounts in a left turn tendency. The turn sharpness depends on the values of the leftmost (dedicated) and the left sensor, implying an intertwining relationship between the normal functioning and obstacle avoidance behaviour. The relation between sensor influence and turn sharpness is proportional

26. As the dendograms are very large, they do not lend themselves for easy visualization, which is why the clusters are shown in this form.

for the left sensor, and inversely proportional for the leftmost sensor. Turn sharpness is increased (i.e. sharper left turn) when the leftmost sensor reading *decreases* or the sensor reading for left sensor *increases*. The turn sharpness is decreased (i.e. less sharp turn to the left) when the sensor reading for the leftmost sensor *increases* or the sensor reading for the left sensor *decreases*. This means that orienting can be ignored despite high sensor readings in the dedicated sensor, which is desirable when, for example a wall rather than an object that gives high sensor readings for the dedicated sensor.

This behaviour is also exhibited when the dedicated sensor reading is zero (absent).

The hidden unit subspaces corresponding to this behaviour are indicated as DIV-T, DV, DVII, DVIII-B and DVIII-T in Figure 10.

2. The behaviour for orienting towards an object is achieved by combining forward motion with behaviour for normal functioning. The similarities of the both motor output values[27] depend on the difference between the left and leftmost (dedicated) sensors, as noted in the previous point. Subsequently, the turn sharpness decreases when the difference between the left and the leftmost (dedicated) sensor value increases. Forward motion is carried out when leftmost (dedicated) sensor reading is present, while the left sensor reading is absent (zero).

   The hidden unit subspaces corresponding to this behaviour are indicated as DIV-B, UI in Figure 10.

Qualitatively speaking, the robot moves forward when the sensor reading for the leftmost (dedicated) sensor is high (or peaks) compared to the left sensor, and turns to the right when

27. Similar output values for both motors lead to forward motion.

the leftmost (dedicated) sensor is absent or suppressed by comparably high left sensor reading.

A behaviour not analysable using hierarchical cluster analysis, was observed during experiment E1, when the dedicated sensor value (fixed on an object) drops to zero without O1 regaining high sensor reading. The reasons for this (loss of fixation) are the constraints imposed by the evolved ANN outputs on the corresponding turns when it comes to turn sharpness. Since the sensor inputs effectively constrain the maximal magnitudes of the output values, they therefore also constrain the maximum sharpness of the turns. This implies that O1 can only move according to certain trajectories, which determines whether a detected object can be oriented towards or not. (An example of the latter behaviour can be found in Appendix A (time steps 326 to 444) during which the robot loses the object to which it oriented (at time step 325) despite continued turn, i.e. normal functioning behaviour.)

4.3 The O2 controller

Contrary to the O1-type controllers, O2-type controllers were only observed to evolve very rarely. In the following analysis, a typical O2-type controller, hereafter referred to as "O2" will be examined during a sample experimental run, E2. Within the first hundred generations, O2 evolved to exhibit near optimal orientation performance, superior to the performance of the O1-type controllers. This becomes apparent when comparing the average fitness values for the parent population of the both controllers O1 (see Figure 6), and O2 (see Figure 13). The average fitness for the parent population at each generation, as well as the trend, being the averaged parent population fitness over the entire generation are shown in the Figure 13.

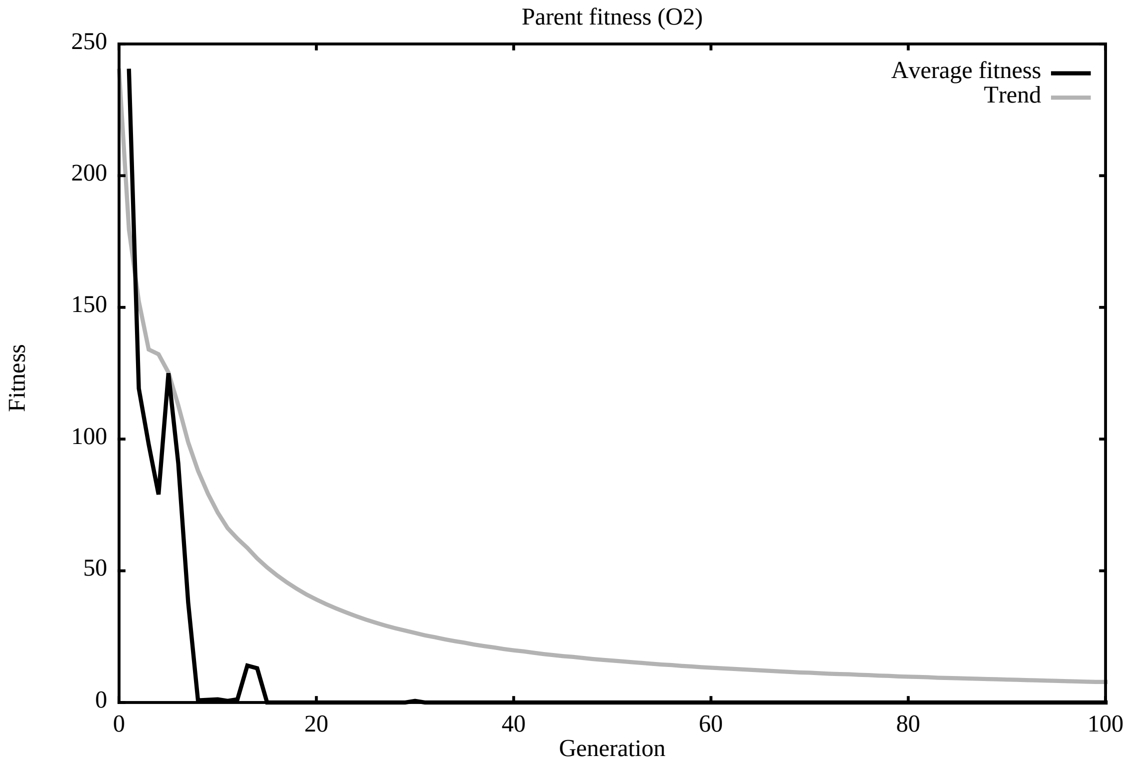

Figure 11: The O2 controller parent population fitness.

As the parent population size is small[28], the population consists of almost identical individuals, it would suggest that the O2 controller is either more stable when it comes to smaller architectural (ANN weight) differences, or that it is more successful in the orienting task possibly due to the constraints imposed on the O1-type controllers by the trajectorial limitations (see section 4.2 for details on the latter). The agents' success in the orienting task was, apart from the fitness, not explicitly measured in the experiments.

28. Five individuals out of 35 in the entire population (see Chapter 3 for further details).

Another characteristic of the O2-type controllers is that, contrary to O1-type controllers, they utilized four of out of five sensors for coping with the orientation task. The dedicated sensors were the leftmost (LL), left (L), centre (C), and the right (R) sensors. Accordingly, the possible motor outputs ranged from forward motion to turn in either direction. Figure 14 shows the trajectory and the heading of an agent using O2 during experiment E2. The sensor inputs and the corresponding motor outputs for E2 can be found in Appendix A.

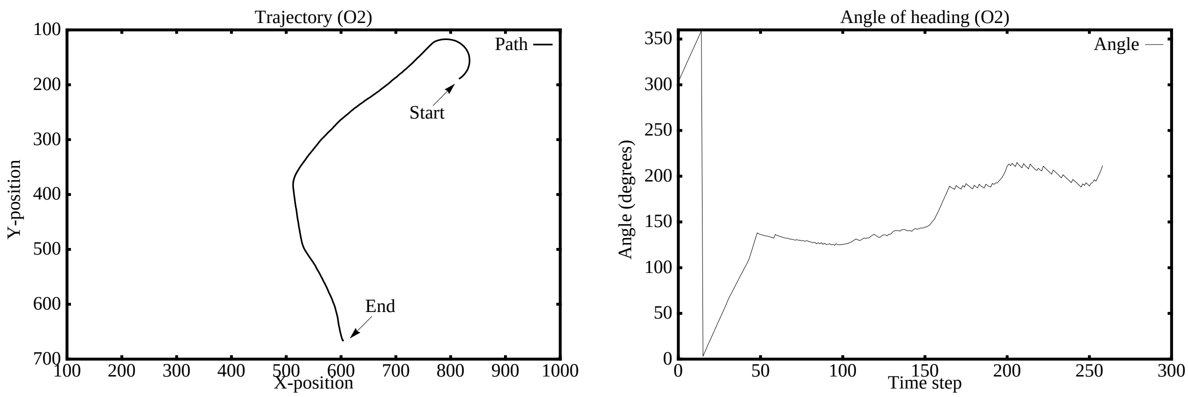

Figure 12: Trajectory and heading for O2 during E2.

The ANN employed by the O2 controller is depicted in Figure 15. (refer to section 4.2 for details on ANN notation). For exact weight and unit bias values for O2, refer to Appendix B.

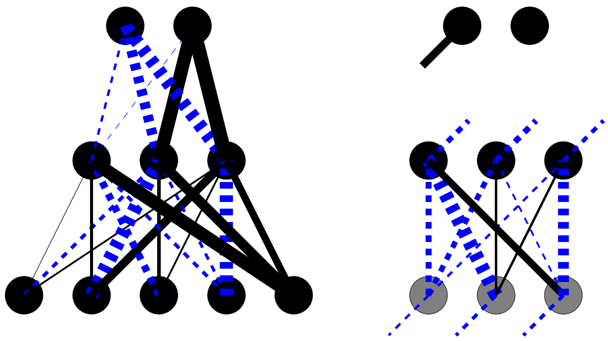

Figure 13: The O2 controller ANN.

Figures for normalized sensor inputs and motor outputs for the O2 controller can be found in Appendix B. An examination of the motor outputs reveals that there is throughout E2 a general tendency for left turn and forward motion (however right turn tendencies occasionally lead to right turns, and are not almost entirely suppressed as in O1, where the right turn tendencies lead to forward motion).

Inspection of the activations at the hidden layer reveals that the partitioning according to turn sharpness, easily identifiable in the hidden unit activation space of O1-type controllers (cf., Figure 9), is much more complex in the O2 controller. The graph is a plot of the hidden unit activation with respect to turn sharpness and turn direction. The figures (values) indicate the difference in activation between the left and right motor outputs (higher difference results in sharper turn). The marked area corresponds to right turn, whereas the marked area corresponds to left turn.

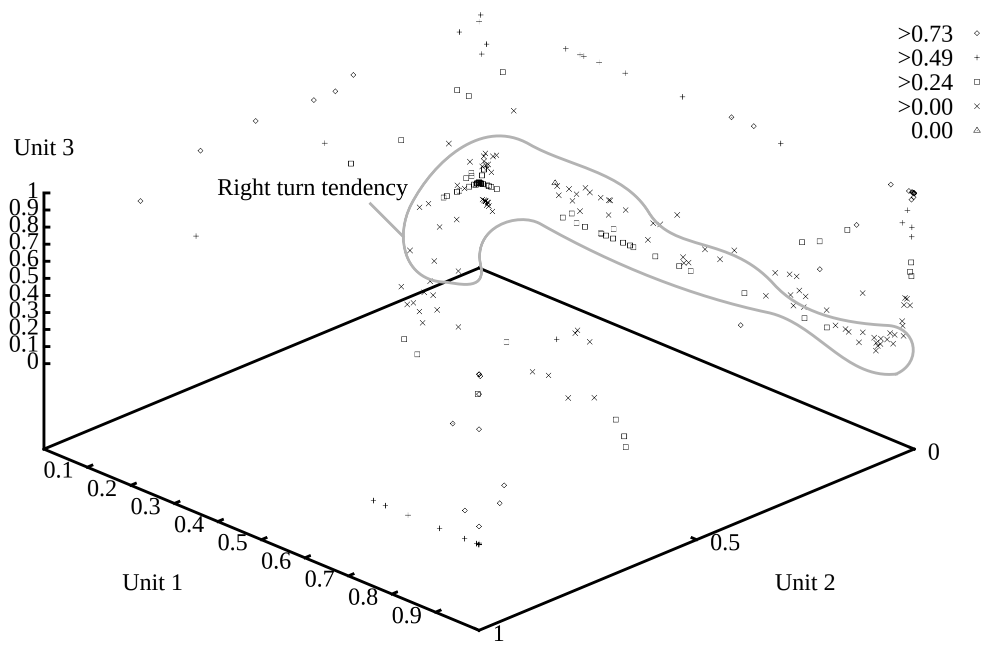

Figure 14: Hidden activation partitioning: turn sharpness and direction (O2).

Hierarchical cluster analysis on the hidden unit activations of the ANN in O2 during E2 revealed a behaviour clustering, which is depicted in Figure 16, where the different markings refer to different behaviours. The figure includes the whole hidden unit activation space.

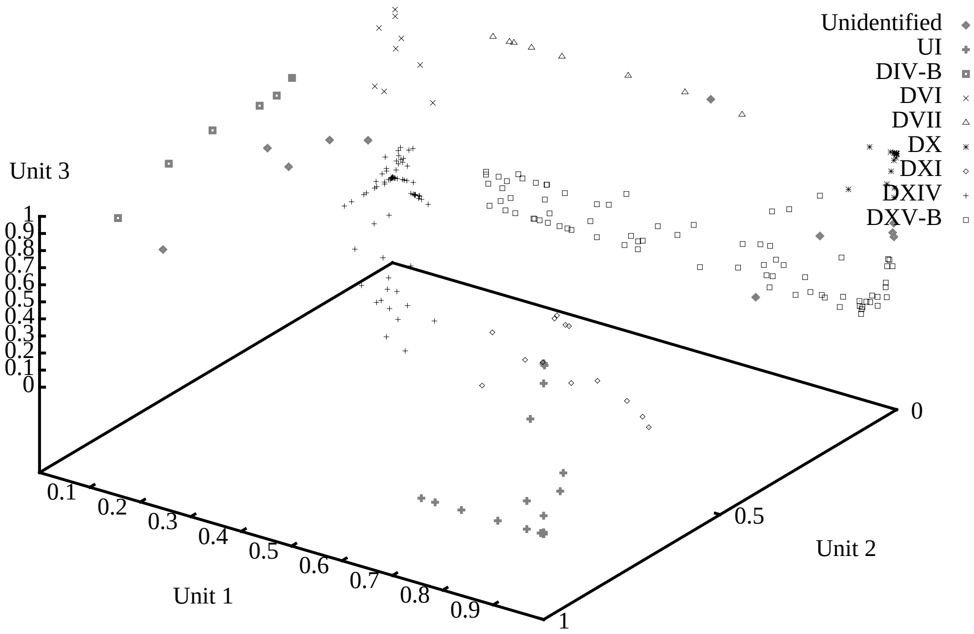

Figure 15: Behavioural clustering of hidden unit activations for O2 during E2.

The overall behaviour of O2 during E2 can be summarized according to the following behaviour categories:

1. Behaviour for normal functioning for moving around in the environment when no sensor readings are active (above zero), in which case there is a small right turn tendency[29].

   The hidden unit subspaces corresponding to this behaviour are denoted as DXI, and DXIV in Figure 17.

29. The difference between the motor activations are approximately 0.1.

2. Behaviour for searching is employed when a sensor value peak, present at time $t-1$ and employed by some behaviour module for exhibiting orientation behaviour, drops at time t, in which case a turn is initiated. The agent makes a right turn when the sensor reading for the left (L) sensor *drops*, and a left turn when the sensor reading for the right (R) sensor *drops*.

   Corresponding hidden unit subspaces are denoted as DIV-B, and DXIV in Figure 17.

3. Behaviour for obstacle avoidance which results in a left turn when all sensor readings are high. The behaviour is ended when the rightmost (RR) sensor reading drops to zero (vanishes).

   The hidden unit subspaces corresponding to this behaviour are indicated as UI, and DX in Figure 17.

4. Behaviour for orienting towards an object is achieved by combining behaviours for normal functioning (no objects within sensor range) with behaviours for orienting (when the object is within dedicated sensor range). The agent makes a right turn when sensor reading for either centre (C) or right (R) sensor *peaks*, and a sharp left turn when the sensor reading from left (L) sensor *peaks*.

   The corresponding hidden unit subspaces corresponding are marked as DVI, DVII, DXIV[30], and DXV-B in Figure 17.

30. Only for right turn when peak on centre (C) sensor.

The behaviour module DXV-B further exhibits behaviour corresponding to motion in three directions (left, right, and forward), where the orientation behaviour is triggered by sensor reading on leftmost (LL) sensor while absent (zero) sensor reading on centre sensor (C). The behaviour ends when the leftmost (LL) sensor drops to zero.

Chapter 5

Discussion and conclusions

The first observation to be made is that orientation behaviour appears within the first one hundred generations. As so often with evolutionary algorithms, the observations made during the conducted experiments show that minimal configuration adequate for good performance (cf. Figure 6) is evolved most frequently.

5.1 The O1-type controllers

The "minimal configuration controllers" (the O1-type controllers) employed only one of the five sensors for coping with the orientation task, although an additional sensor was utilized for achieving simple obstacle avoidance behaviour. In all observed cases, this dedicated sensor was positioned at an outermost position (the leftmost or rightmost sensor).

Subsequently, the overall turn tendency evolved to be towards the same direction as the dedicated sensor (e.g. left turn tendency when leftmost sensor was dedicated). An observation of the hidden unit activation for an ANN employed in O1-type controllers, revealed a partitioning corresponding to turn sharpness and direction, confirmed the overall turn tendency on one motor, and indicated that high activation for the opposite motor only resulted in forward motion (i.e. no turn).

Hierarchical cluster analysis of the same hidden unit activations also revealed clusters corresponding to different behaviours (or rather sub-behaviours). These behaviours can be roughly organized in three categories:

1. Behaviour for normal functioning, which is employed when the agent moves around in its environment "looking" for the object towards which to orient (i.e. when the object is outside the sensor range). This behaviour is very simple and amounts only in a motion which is a light turn towards the same direction as the dedicated sensor (biased towards the "strong" motor). This slow turning behaviour allows the agent to cover (from both area and sensor point of view) a big portion of its environment by exploiting the fact that the environment is rather limited and surrounded by walls.

2. Behaviour for simple obstacle avoidance, e.g. avoiding walls. This behaviour amounts in turn in the same direction as the dedicated sensor with respect to the agent centre. It is activated when the sensor reading next to the dedicated sensor[31] increases to comparable readings to that of the dedicated sensor. This very limited form of obstacle avoidance behaviour thus *subsumes* the object tracking behaviour. This behaviour can also be seen as a very simple form of *discrimination*, exploiting the fact that the object is never positioned so as to give high sensor reading on the outermost and next to outermost sensors at the same time step[32].

31. i.e. the left (L) sensor when the leftmost (LL) sensor is dedicated

32. This behaviour thus also exploits the condition that the object is never positioned close to a wall, in which case a discrimination would not have been possible. The latter was observed during the initial orientation experiments, where the objects could be placed at an arbitrary position within the environment, with the result of very poor (close to random) performance.

3. Behaviour for object tracking, which is employed once the agent has moved to a position where the object produces high sensor readings on the dedicated sensor. This behaviour amounts in forward motion as long as the dedicated sensor reading remains high (provided that the obstacle avoidance behaviour is not invoked).

The orienting behaviour is exhibited by combining the behaviour for normal functioning with the behaviour for object tracking. The object tracking behaviour is responsible for keeping the object “in sight”, while the normal functioning behaviour is evoked when the object is “lost from sight”. The object lost from sight as a result of forward motion implies that the object is positioned (and should therefore be sought for) at the same side as the dedicated sensor. As the normal functioning behaviour amounts in turns in the same direction as the dedicated sensor, it inherently starts seeking for the object in order to regain object sensor reading on the dedicated sensor. The normal functioning thus *inherently* copes with the *object persistence problem*[33]. This result also suggests that simple reactive (feedforward) ANNs would have succeeded in the orienting task by employing a similar strategy. The orientation behaviour can be described in terms of a subsumption architecture (see Chapter 2), where the different behaviour modules and the relationships (control and organization structures) between these modules are evolved rather than designed os in Brooks’ original subsumption architecture[34]. Figure 18 depicts the O1-type controller interpreted as a subsumption architecture, where the obstacle avoidance behaviour can subsume the object tracking behaviour, which in turn can subsume the behaviour for normal functioning.

33. The observations thus suggest that the object persistence problem is limited to the environment in which the experiments were conducted, since the behaviours succeed in finding the object after it has disappeared from view by exhibiting a small radius turn, thus exploiting the (area) limitations of this restricted environment.

34. Similarly, Ziemke (1996) presented experiments in which behaviour ‘modules’ and their organization emerged (as structured internal state spaces) from on-line learning in recurrent connectionist agent controllers.

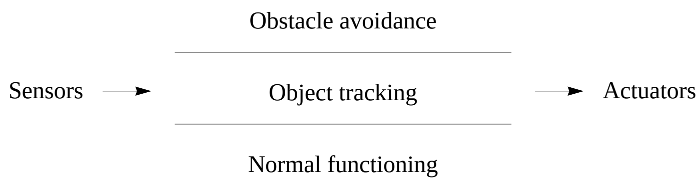

Figure 16: The O1-type controller as a subsumption architecture.

The success of this object orientation strategy is constrained by the range of possible motor outputs corresponding to range of turn sharpness that is possible to exhibit by the agent.

5.2 The O2-type controllers

The second type of controller, the O2-type controller, was observed to evolve very rarely. Agents utilizing this controller type exhibited near optimal behaviour, superior to that of the agents employing O1-type controllers. The O2-type controllers used four of the five possible sensors[35] for successful handling of the orientation task.

Accordingly, the possible motor outputs ranged from forward motion to turn in either direction, i.e. no turn suppressed as in the O1-type controllers. Examination of the hidden unit activations for a sample run revealed a far less clear partitioning according to turn sharpness than in the case of the O1-type controllers.

Hierarchical cluster analysis of the hidden unit activations unveiled clusters corresponding to different (sub-)behaviours, which can coarsely be organized in the following categories:

1. Behaviour for normal functioning is exhibited when no objects are visible for the agent, and results in a small radius turn to the right.

35. The dedicated sensors were the leftmost (LL), left (L), centre (C), and the right (R) sensors.

2. Behaviour for searching for an object "lost from sight". This behaviour is exhibited when an object is lost from the dedicated sensor range, where the dedicated sensor was used in orienting behaviour in the preceding time steps. The behaviour amounts in a turn, where the direction depends on the dedicated sensor employed for the orientation behaviour before the sensor reading drop. The agent carries out a right turn when the sensor readings for the left (L) sensor drops, or a left turn when the sensor readings for the right (R) sensor drops. This suggests that the controller exhibits *active search*, which is different from the *passive search* (normal functioning) employed by the O1-type controllers. This is probably due to the controller taking advantage of the dynamics of the recurrent ANN in order to produce output according to the dedicated sensor that gave high readings most recently.

3. Behaviour for obstacle avoidance resulting in left turn when all sensory readings are high, and is subsumed by normal functioning when the sensor reading for the right-most sensor (RR) drops to zero, i.e. when the robot has turned away from the obstacle. Similarly to the O1-type controller obstacle avoidance behaviour, this behaviour too exploits the fact that an object cannot be positioned so as to give high sensor reading on two sensors simultaneously (at the same time step), although the behaviour utilizes all sensors rather than just two, as in the case of the O1-type controller. This behaviour can thus be seen as simple form of *discrimination* behaviour.

4. Behaviour for object tracking is employed when the agent has positioned itself so as the object produces high sensor readings on a dedicated sensor. This behaviour amounts in right turn, when the object gives high sensor reading for the centre (C) or right (R) sensors, and left turn for high sensor reading in the left (L) sensor.

Similarly as for the O1-type controllers, the O2-type controllers can be interpreted as a subsumption architecture with evolved behaviour modules and control organization and structures. Figure 19 shows the O2-type controller interpreted as a subsumption architecture (cf. section 2.1), where the obstacle avoidance behaviour can subsume the object tracking behaviour, which can subsume the active search behaviour, which in turn can subsume the normal functioning behaviour.

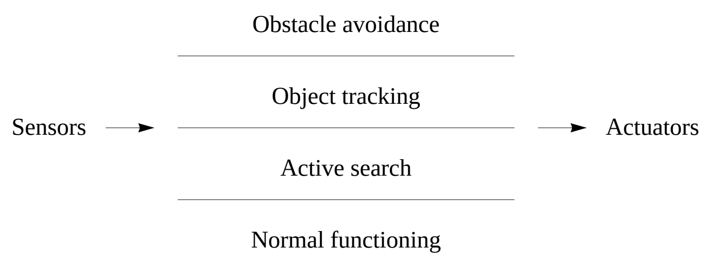

Figure 17: The O2-type controller as a subsumption architecture.

The orienting behaviour is exhibited by employing the behaviours for object tracking with active search. This behaviour can be regarded as using a *foveation* strategy, where the agent keeps the object in front of it by successive left and right turn (cf. Figure 14).

There also exists a behaviour cluster, which is responsible for exhibiting a whole foveation strategy for orienting towards an object giving high sensory reading in the leftmost (LL) sensor. This behaviour is subsumed by the regular behaviour tracking (e.g. at high sensor readings in the centre (C) sensor).

5.3 Future work

The field where evolutionary robotics and computational neuroethology are used in the investigation of underlying mechanisms of adaptive behaviour such as visually-guided behaviour is virtually unexplored and leaves plenty of work to be done. Some possible directions for future work are:

Orientation and discrimination in presence of multiple objects. Initial work was done by Beer in visually-guided discrimination. He reports that the agents evolve to successfully handle simple discrimination tasks (with one object present at a time) but the speed with which successful individuals evolve (i.e. the convergence ratio) is very slow compared to the orientation task. The discrimination task could be extended to include orientation and discrimination in presence of multiple objects. Slow convergence ratio was confirmed by preliminary results from initial experiments in the latter task, where the emergent behaviours employed strategies of O1-type controllers for orienting towards an object, where in the presence of two objects, the agents alternately oriented to the both objects, resulting in an oscillating behaviour between the two objects.

Recombination with ANN property preservation. Recombination techniques that take ANN properties into consideration would probably increase the convergence ratio.

Incremental evolution and integration of behaviours. Agent controllers, evolved for one task, could be used as starting points for evolution of more complicated task, or for integration of behaviours. This approach is taken by researchers at COGS, University of Sussex, where a framework called SAGA (Species Adaptation Genetic Algorithms (Harvey, 1992)) has been developed for open-ended incremental evolution, and takes into consideration that the task specifications might not be known from the outset or might change during the incremental evolution.

Dynamic analysis. As recurrent ANNs are dynamical systems, these could beneficially be analysed using dynamical analysis methods, such as dynamic systems theory (DST).

Real world validation. In order to prove that the agent can exhibit successful behaviours in real world conditions, it needs to be validated in the real world using a real robot.

5.4 Conclusions

As discussed in the first chapter, it is all too easy to ascribe cognitive phenomena to an agent to assume that each phenomena has corresponding agent internal cognitive and psychological structures. As the complexity of the agent's behaviours reflect the complexity of its environment, a way of avoiding this pitfall is by studying concrete models of situated and embodied agents, to illuminate the essential principles underlying the exhibition of behaviours associated with cognitive phenomena. As the understanding of deliberate reasoning poses tractability problems when considering the (e.g. structural) complexity of the brain and body and the descriptively defined (vague) terms referring to corresponding ascribed cognitive phenomena, adaptive behaviour research engages in studying the mechanisms underlying adaptive behaviour, which allow animals to remain viable in their environment, and ignore deliberate reasoning as long as it does not restrain the adaptiveness of behaviour.

This work complies to this standpoint of adaptive behaviour research by taking a bottom-up approach to artificially evolve internal dynamics (connections of recurrent ANNs) of controllers employed in simulated embodied and situated agents which are engaged in visually-guided orientation, a form of adaptive behaviour raising issues about cognitive phenomena such as visual orientation, object perception and discrimination, and visually-guided motion. Situatedness and embodiment allow the grounded internal dynamics of the visually-guided agents to emerge from their interaction with the environment as a result of the artificial evolution process (cf. the enactive approach, Chapter 2). The latter was guided by a very simple fitness function, where the distance between the agent and the object (to which the agent should orient) determined the measure on the agent's success in the task. This approach allowed for evolution of controllers with emergent behaviours, and structure and control organizations. The emergent behaviours included simple discrimination, obstacle avoidance, object tracking, passive and active search, the latter exploiting the dynamics of the recurrent ANNs. The emergent structural and control organizations introduced hierarchical (priority) constraints on the different behaviours, allowing behaviours (with high priority) to suppress other behaviours (with lower priority).

In conclusion, it can be stated that by employing evolutionary mechanisms for evolving ANN dynamics, where the agent's closeness to an object is rewarded, dynamics can evolve which allows the agent to exhibit behaviours which can be associated with simple forms of *passive search*, *active search*, *discrimination*, *foveation*, and *object persistence.*

Acknowledgements

I would like to express my profound gratitude to some persons without whom this work would not have been possible to accomplish - my supervisor Tom Ziemke, for commenting my occasionally very confusing drafts and for giving me valuable advices throughout the entire project, and my brother László Biró, to whom I dedicate this work, for motivating me to carry on and complete this work in time.

I would like to thank the members of the Connectionist Research Group, Randall Beer, and Susanna Grahammar for helping me in one way or another at the initial stages of the project.

I am also grateful to have received the Swedish-Austrian scholarship of the University of Skövde, which allowed me to attend the NTCS-97 conference in Vienna.

Thank you, Anne-Marie Daniels for having shown great tolerance and understanding, during the entire project.

Bibliography

* Agre, P.E. (1988). The Dynamic Structure of Everyday Life. AI Technical Report 1085, MIT.

Bäck, T. & Schwefel, H. (1993). An overview of evolutionary algorithms for parameter optimization. *Evolutionary Computation*, **1** (1): 1-23.

Beer, R.D. (1990). *Intelligence as Adaptive Behavior: An Experiment in Computational Neuroethology.* San Diego: Academic Press.

Beer, R.D. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, **72**: 173-215.

Beer, R.D. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. In Maes P., Mataric M., Meyer J.-A, Pollack J., Wilson S., eds., *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior (SAB96),* 421-429. MIT Press.

* Boden, M.A. (1994). Autonomy and artificiality. In Boden, M.A., ed. *The Philosophy of Artificial Life*. Oxford University Press. Also available as Cognitive Science Research Paper 307, University of Sussex.

Brooks, R.A. (1986). A Robust Layered Control System for a Mobile Robot. In *IEEE Journal of Robotics and Automation*, **2** (1): 14-23. Also available as AI Memo 864, MIT.

Brooks, R.A. (1991). Intelligence Without Reason. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI-91)*, 569-595. San Mateo: Morgan Kaufmann.

Bruce, C., Green, P. (1990). *Visual perception: physiology, psychology and ecology.* Exeter: Lawrence Erlbaum Associates.

* Chapman, D. (1990). Vision, Instruction and Action. AI Technical Report 1204, MIT.

Chemero, A. (1997). Two types of anti-representationalism: A taxonomy. In *Proceedings from the International Conference New Trends in Cognitive Science (NTCS’97)*, ASoCS Technical Report 97-01, 44-48. Vienna: Austrian Society of Cognitive Science.

Cliff. D.,T. (1994). Neuroethology, Computational. Cognitive Science Research Paper 338, University of Sussex.

Cliff, D., Husbands, P., Harvey, I. (1992). Analysis of Evolved Sensory-Motor Controllers. Cognitive Science Research Paper 264, University of Sussex.

Cliff, D.T., Husbands, P., Harvey, I. (1993). Evolving visually guided robots. Cognitive Science Research Paper 220, University of Sussex.

Dorffner, G. (1997). Radical connectionism - a neural bottom-up approach to AI. In Dorffner, G., ed., *Neural Networks and a New Artificial Intelligence*, 93-132. London: International Thomson Computer Press.

Eysenck, M.W., Keane, M.T. (1995). *Cognitive Psychology: A Student's Handbook*. Hove: Psychology Press.

Floreano, D., Mondana, F. (1994). Active Perception, Navigation, Homing, and Grasping: An autonomous Perspective. In Gaussier P., Nicoud J.-D., eds., *From Perception to Action Conference*. Los Alamitos: IEE Computer Society Press.

Franklin, S., Graesser, A. (1996). Is it an Agent, or just a Program?: A Taxonomy for Autonomous Agents. In Mueller, J.P., Wooldridge, M., Jennings, N. (Eds), *Intelligent agents III: Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages.* Berlin: Springer-Verlag.

Ginsburg, H.P., Opper, S. (1988). *Piaget's Theory of Intellectual Development*, 3rd edn. Englewood Cliffs: Prentice Hall.

Harnad, S. (1990). The symbol grounding problem. *Physica D*, **42**: 335-346.

Harvey, I. (1992). Species adaptation genetic algorithm: The basis for a continuing SAGA. In Varela, F.J., Bourgine, P., eds., *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, 346-354. MIT Press.

Harvey, I., Husbands, P., Cliff, D. (1994). Seeing the Light: Artificial Evolution, Real Vision. In Cliff, D., Husbands, P., Meyer, J.-A., Wilson S., eds., *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior (SAB94)*. MIT Press. Also available as Cognitive Science Research Paper 317, University of Sussex.

Harvey, I., Husbands, P., Cliff, D., Thompson, A., Jakobi, N. (1996). Evolutionary Robotics: the Sussex Approach. In Pfeifer, R., Brooks, R.A., eds., *Robotics and Autonomous Systems, Special Issues on "Practice and Future of Autonomous Agents"*.

Kelley, C., Williams, T. (1993). Gnuplot. UNIX version 3.5.

* McCarthy, J., Hayes, P.J. (1969). Some philosophical problems from the standpoint of artificial intelligence. In Meltzer, B., Michie, D., Swann, M., eds., *Machine Intelligence*, **4**: 463-502. Edinburgh: Edinburgh University Press.

Meeden, L.A. (1996). An Incremental Approach to Developing Intelligent Neural Network Controllers for Robots. In *IEEE Transactions on Systems, Man, and Cybernetics*, Volume 26, **3**: 474-485.

Michel, O. (1996). *Khepera Simulator version 2.0 User Manual*. University of Nice-Sophia Antipolis, Valbonne, France (`http://wwwi3s.unice.fr/~om/`)

Mitchell, M. (1996). *An Introduction to Genetic Algorithms*. MIT Press.

Newell, A. (1980). Physical Symbol Systems. *Cognitive Science*, **4**: 135-183.

* Nilsson, N.J., ed., (1984). Shakey the Robot. Technical Note 323, SRI A.I. Center.

Palmer, S.E. (1978). Fundamental Aspects of Cognitive Representation. In Rosch, E., Lloyd, B.B., eds., *Cognition and Categorization*, 259-274. Hillsdale: Lawrence Erlbaum Associates.

* Rechenberg, I. (1973). *Evolutionsstrategie*. Stuttgart: Frommann-Holzboog. In *Evolutionsstrategie '94* (1994). Stuttgart: Frommann-Holzboog.

Salomon, R. (1996a). Increasing adaptivity through evolution strategies. In Maes P., Mataric M., Meyer J.-A., Pollack J., Wilson S., eds., *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior (SAB96)*, 411-420. MIT Press.

* Salomon, R. (1996b). Performance Degradation of Genetic Algorithms under Coordinate Rotation. To appear in *Fifth Annual Conference on Evolutionary Programming (EP96)*.

* Salomon, R. (1996c). Implicit Independence Assumptions; a Notorious Problem for Genetic Algorithms. To appear in *Proceedings of the International Symposia on Soft Computing and Intelligent Industrial Automation (SOCO96)*.

Scheier, C., Pfeifer, R. (1995a). Classification as Sensory-Motor Coordination: A Case Study on Autonomous Agents. AI Lab Technical Report 95.05, Computer Science Department, University of Zürich.

Scheier, C., Pfeifer, R. (1995b). Categorization in a real-world agent using haptic exploration and active perception. AI Lab Technical Report 95.15, Computer Science Department, University of Zürich.

Schöneburg, E. (1994). *Genetische Algorithmen und Evolutionsstrategien: Eine Einführung in Theorie und Praxis der simulierten Evolution*. Addison-Wesley (Deutschland).

Sharkey, N.E. (1997). Artificial Neural Networks for Coordination and Control: The Portability of Experiential Representations. To appear in *Robotics and Autonomous Systems*, special issue *"Robot Learning: The New Wave"*.

Sharkey, N.E., Heemskerk, J. (1997). The neural mind and the robot. To appear in Browne, A.J., ed., *Perspectives in Neural Computing*. Bristol: IOP Press.

* Simon, H.A. (1969). *The Sciences of the Artificial*. MIT Press.

* Smith, B.C. (1996). *On the Origin of Objects.* MIT Press.

Thornton, C. J. (1992). *Techniques in Computational Learning: An introduction*. London: Chapman & Hall.

van Gelder, T. (1992). What might cognition be if not computation? In *Journal of Philosophy*, **92**: 345-381. Also available as Technical Report 75, Cognitive Sciences, Indiana University.

van Gelder, T. (1995). Modeling, Connectionist and Otherwise. In Niklasson, L.F., Bodén, M.B., eds., *Current Trends in Connectionism: Proceedings of the Swedish Conference on Connectionism, 1995 (SCC95).* Hillsdale: Lawrence Erlbaum Associates.

van Gelder, T. (1997). The Dynamical Hypothesis in Cognitive Science. *To appear in Behavioral and Brain Sciences*.

Varela, F. J., Thompson, E., Rosch, E. (1991). *The embodied mind: cognitive science and human experience*. Cambridge: MIT Press.

* Warren, W.H. (1995) Self-motion: Visual perception and visual control. In Epstein, W., Rogers S., eds., *Handbook of Perception and Cognition, volume 5: Perception of Space and Motion*. New York: Academic Press.

Wheeler, M. (1993). Explaining the Behaviour of Springs, Pendulums, and Cognizers. Cognitive Science Research Paper 284, University of Sussex.

Wheeler, M. (1994). Active Perception in Meaningful Worlds. Cognitive Science Research Paper 327, University of Sussex.

Wilson, S.W. (1985). Knowledge growth in an artificial animal. In *Proceedings of the First International Conference on Genetic Algorithms.* Pittsburgh: Lawrence Erlbaum.

* Wood, W.B. (1988). *The Nematode Caenorhabditis Elegans.* Cold Spring Harbor Laboratory, Cold Spring Harbor.

Ziemke, T. (1996). Towards Adaptive Behaviour System Integration using Connectionist Infinite State Automata. In Maes P., Mataric M., Meyer J.-A, Pollack J., Wilson S., eds., *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior (SAB96)*. MIT Press.

Ziemke, T. (1997). The 'Environmental Puppeteer' Revisited: A Connectionist Perspective on 'Autonomy'. In *Proceedings of the Sixth European Workshop on Learning Robots (EWLR-6)*.

Appendix A

Sensor inputs and corresponding motor outputs for controller O1 during experiment E1.

ANN connection weight and unit bias values for O1.

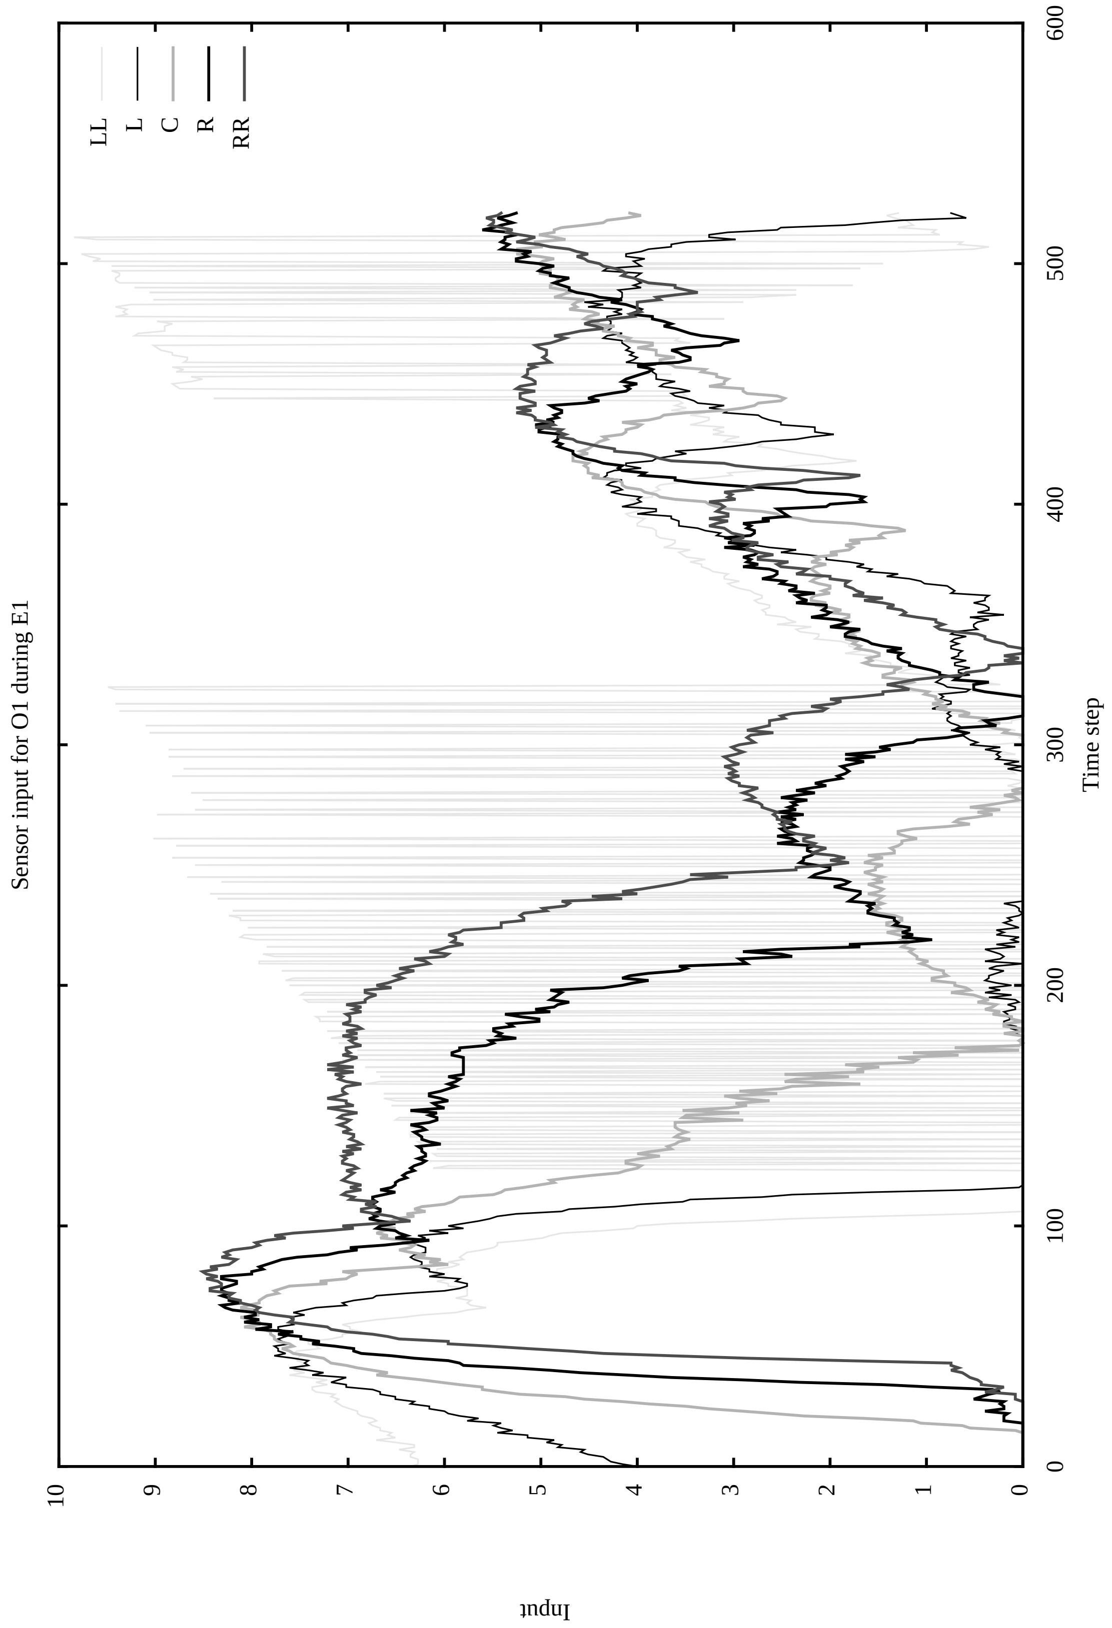
Sensor input for O1 during E1
LL
L
C
R
RR
Time step
Input
0
100
200
300
400
500
600
0
1
2
3
4
5
6
7
8
9
10

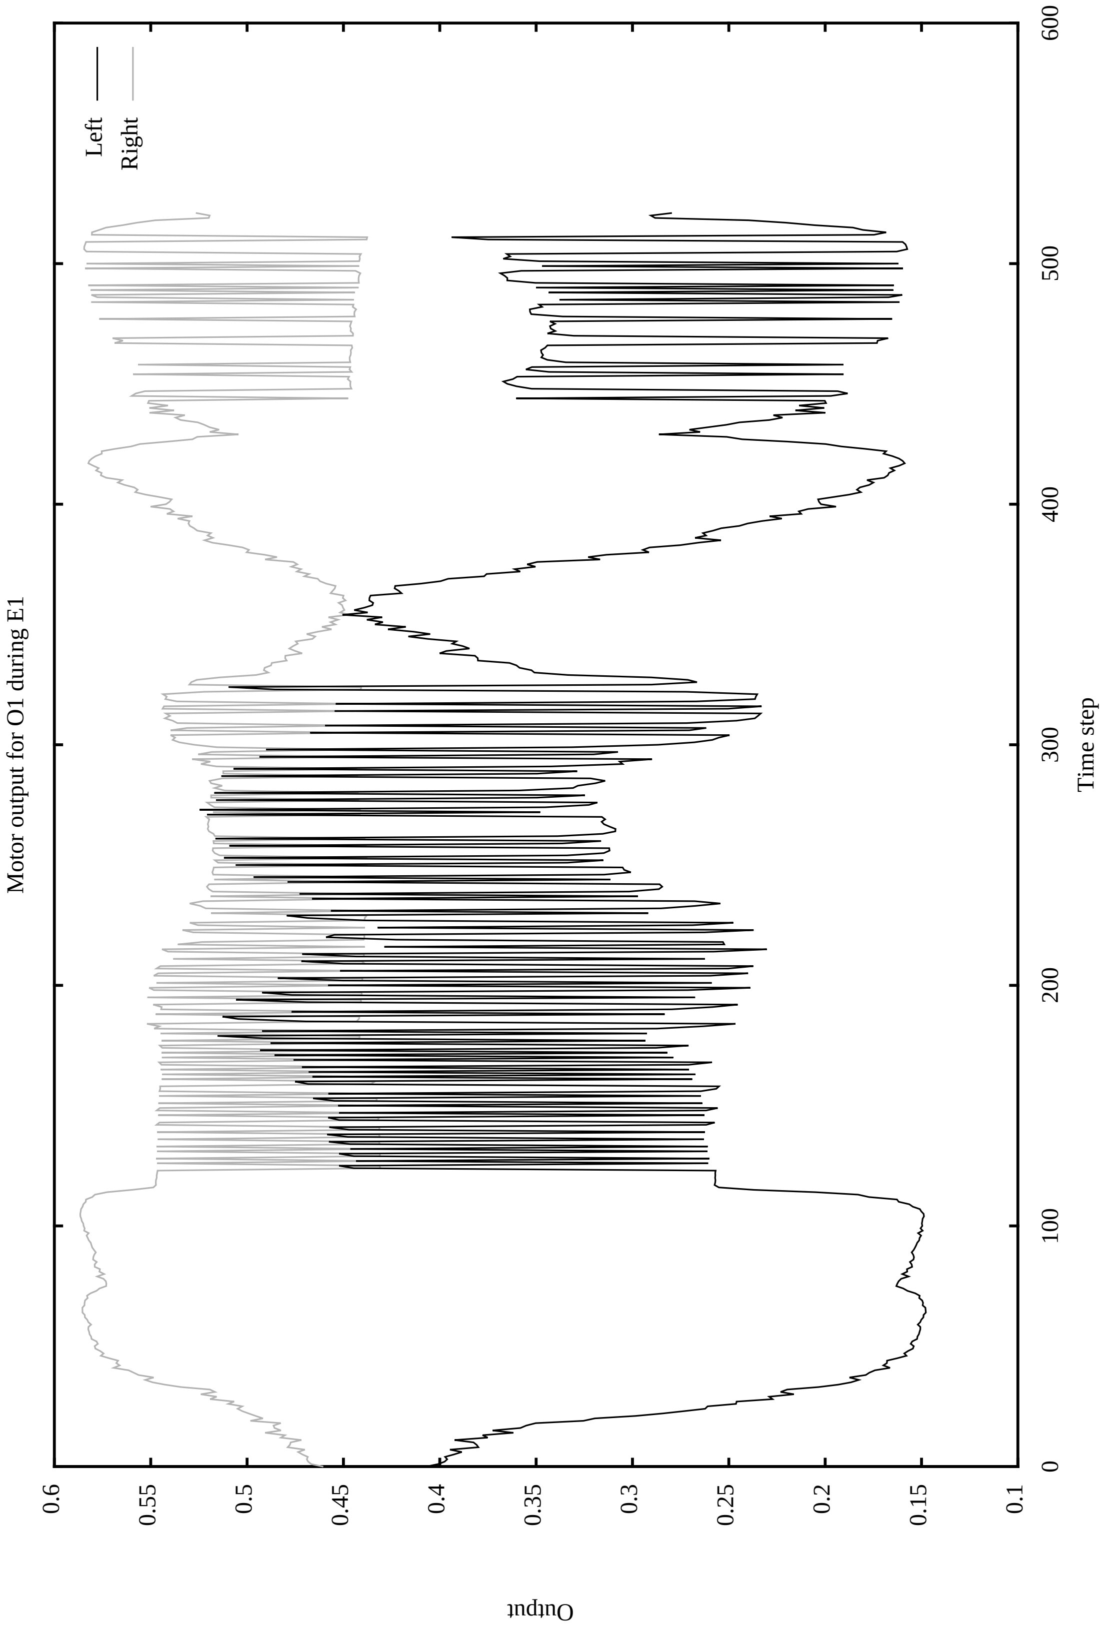
Motor output for O1 during E1
Left
Right
Time step
Output
0
100
200
300
400
500
600
0.6
0.55
0.5
0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1

ANN connection weight and bias values for O1.

```
Input[0]->Hidden[0] =  1.044734
Input[0]->Hidden[1] = -0.169445
Input[0]->Hidden[2] = -0.083699
Input[1]->Hidden[0] = -1.135065
Input[1]->Hidden[1] = -0.044608
Input[1]->Hidden[2] =  1.104705
Input[2]->Hidden[0] = -0.046301
Input[2]->Hidden[1] =  0.732977
Input[2]->Hidden[2] = -0.499084
Input[3]->Hidden[0] = -0.173365
Input[3]->Hidden[1] = -0.154885
Input[3]->Hidden[2] = -0.250949
Input[4]->Hidden[0] = -0.047572
Input[4]->Hidden[1] =  0.174874
Input[4]->Hidden[2] =  0.228863

Input[5]->Hidden[0] =  0.609141 (BIAS)
Input[5]->Hidden[1] =  0.006819 (BIAS)
Input[5]->Hidden[2] =  0.037894 (BIAS)


Context[0]->Hidden[0] = -0.084283
Context[0]->Hidden[1] = -0.522051
Context[0]->Hidden[2] = -0.319004
Context[1]->Hidden[0] = -0.328946
Context[1]->Hidden[1] =  0.154793
Context[1]->Hidden[2] =  0.746461
Context[2]->Hidden[0] = -0.040507
Context[2]->Hidden[1] =  0.243953
Context[2]->Hidden[2] =  0.717935

Context[3]->Hidden[0] = -0.042603 (BIAS)
Context[3]->Hidden[1] =  0.138986 (BIAS)
Context[3]->Hidden[2] = -0.924395 (BIAS)


Hidden[0]->Output[0] =  0.988394
Hidden[0]->Output[1] = -0.586872
Hidden[1]->Output[0] = -0.596363
Hidden[1]->Output[1] = -0.055335
Hidden[2]->Output[0] = -0.554677
Hidden[2]->Output[1] =  0.057118

Hidden[3]->Output[0] = -0.624696 (BIAS)
Hidden[3]->Output[1] =  0.350978 (BIAS)
```

Appendix B

Sensor inputs and corresponding motor outputs for controller O2 during experiment E2.

ANN connection weight and unit bias values for O2.

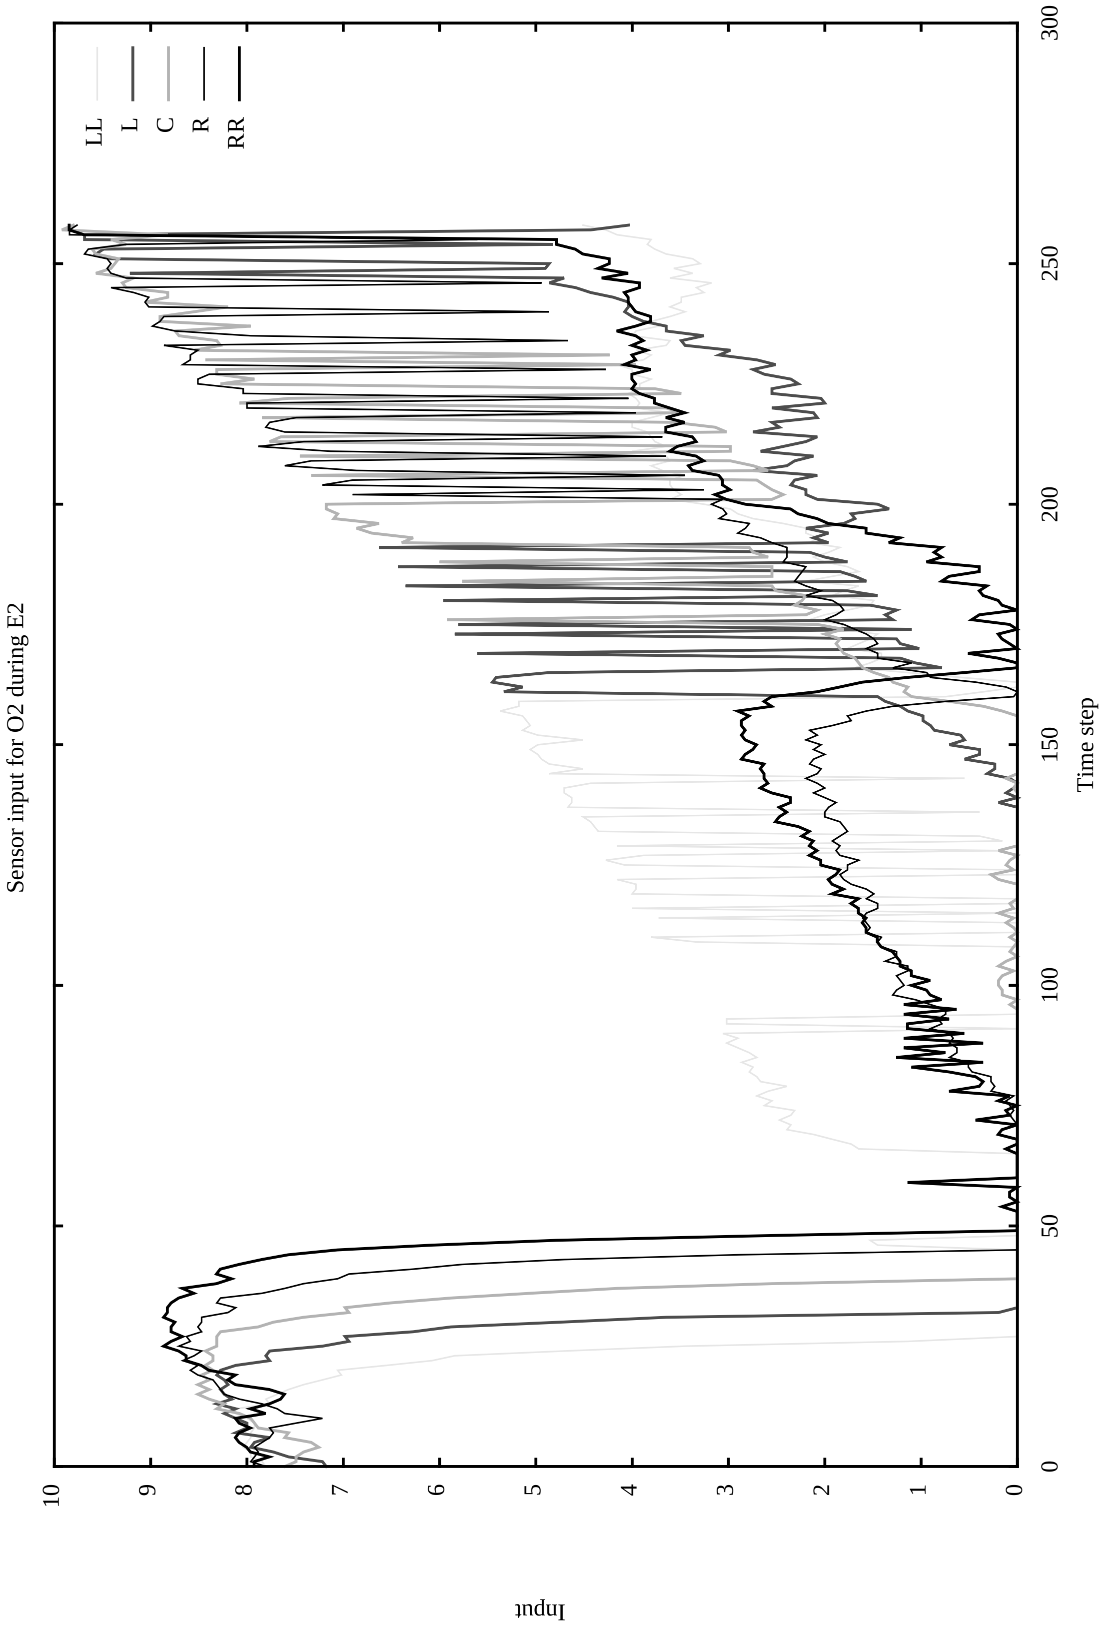
Sensor input for O2 during E2
Time step
Input
LL
L
C
R
RR
0
50
100
150
200
250
300
0
1
2
3
4
5
6
7
8
9
10

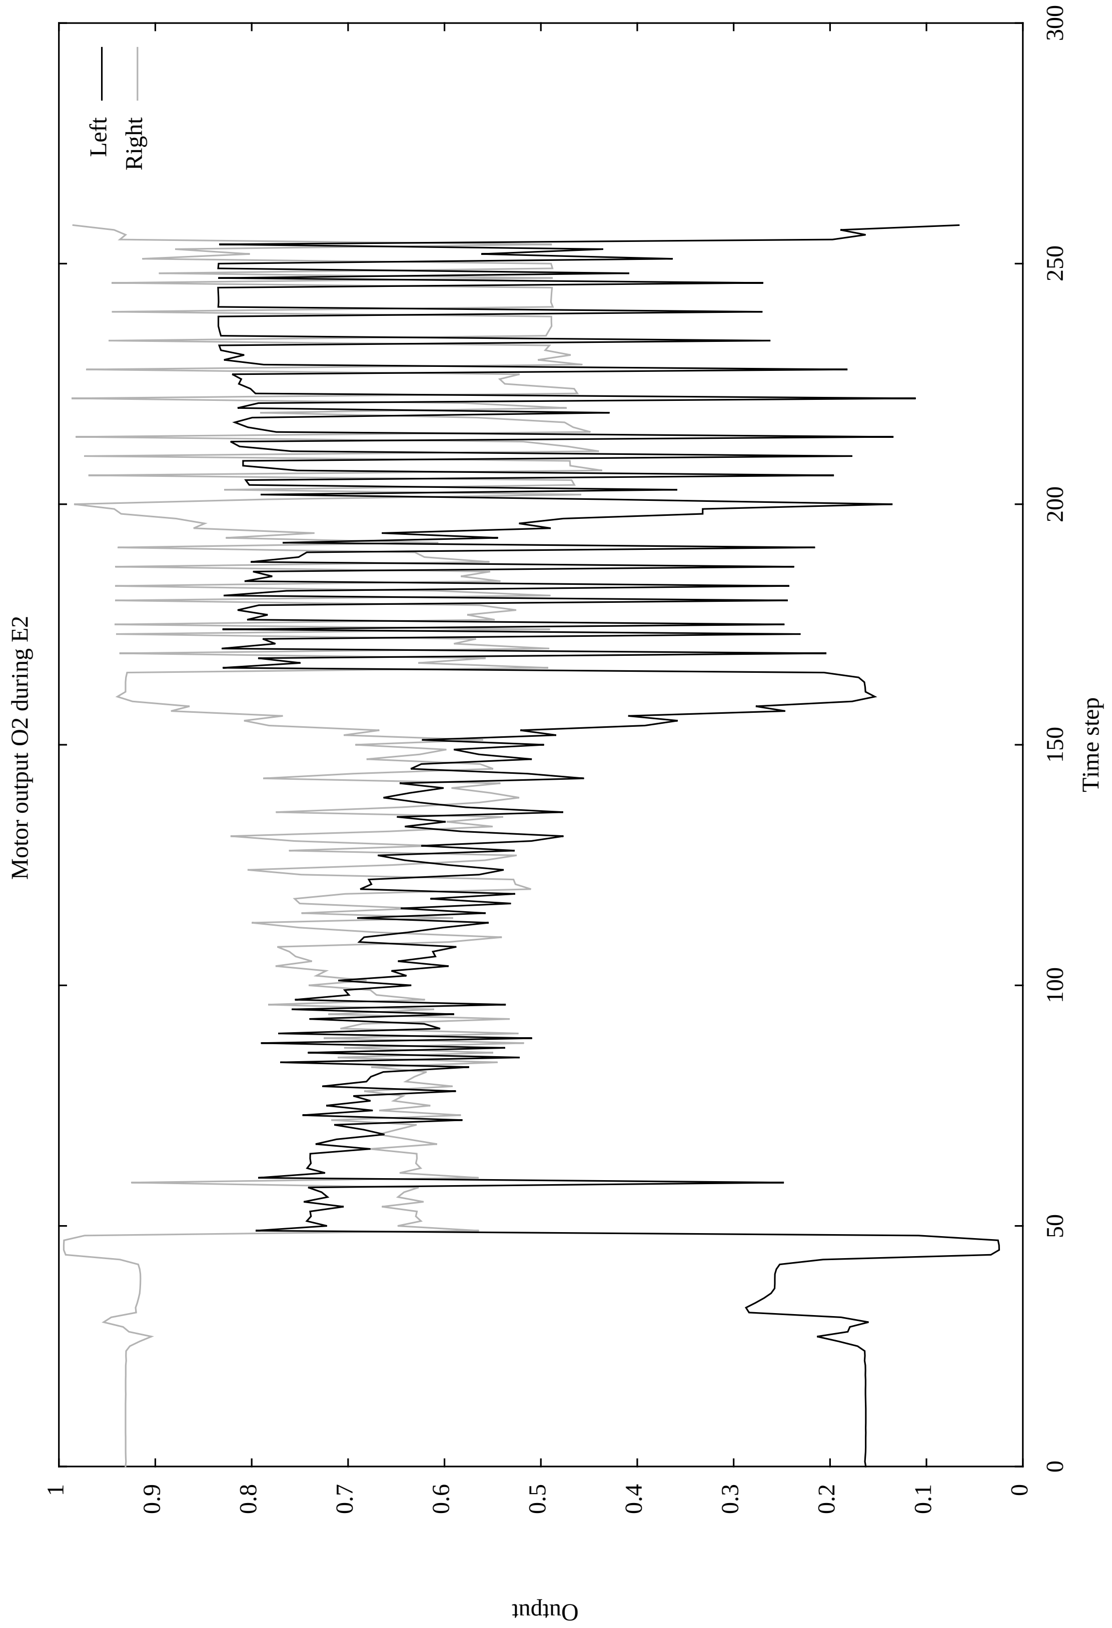
Motor output O2 during E2
Left
Right
Time step
Output
0
50
100
150
200
250
300
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0

ANN connection weight and bias values for O2.

```
Input[0]->Hidden[0] =  0.253060
Input[0]->Hidden[1] = -0.808560
Input[0]->Hidden[2] =  0.474743
Input[1]->Hidden[0] =  0.739581
Input[1]->Hidden[1] = -2.773240
Input[1]->Hidden[2] =  1.706559
Input[2]->Hidden[0] = -1.516205
Input[2]->Hidden[1] =  0.835460
Input[2]->Hidden[2] =  0.439668
Input[3]->Hidden[0] = -0.757605
Input[3]->Hidden[1] = -0.686142
Input[3]->Hidden[2] = -2.644522
Input[4]->Hidden[0] =  2.662894
Input[4]->Hidden[1] =  2.114125
Input[4]->Hidden[2] =  1.547377

Input[5]->Hidden[0] = -1.215724 (BIAS)
Input[5]->Hidden[1] = -1.253749 (BIAS)
Input[5]->Hidden[2] = -0.986022 (BIAS)


Context[0]->Hidden[0] = -1.285684
Context[0]->Hidden[1] = -1.127804
Context[0]->Hidden[2] = -0.625236
Context[1]->Hidden[0] = -2.820007
Context[1]->Hidden[1] =  0.583184
Context[1]->Hidden[2] =  0.590916
Context[2]->Hidden[0] =  1.467944
Context[2]->Hidden[1] = -0.420192
Context[2]->Hidden[2] = -2.237860

Context[3]->Hidden[0] = -0.581209 (BIAS)
Context[3]->Hidden[1] = -0.837435 (BIAS)
Context[3]->Hidden[2] = -0.759361 (BIAS)


Hidden[0]->Output[0] = -0.637832
Hidden[0]->Output[1] = -0.258609
Hidden[1]->Output[0] = -2.044924
Hidden[1]->Output[1] =  2.689275
Hidden[2]->Output[0] = -2.623067
Hidden[2]->Output[1] =  2.909031

Hidden[3]->Output[0] =  1.623818 (BIAS)
Hidden[3]->Output[1] = -0.049797 (BIAS)
```