

Situated Representation

Situated Representation

PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit Maastricht,
op gezag van de Rector Magnificus,
Prof. mr. G.P.M.F. Mols,
volgens het besluit van het College van Decanen,
in het openbaar te verdedigen
op donderdag 1 december 2005, om 12:00 uur

door

Michel van Dartel

Promotores: Prof. dr. E.O. Postma
Prof. dr. H.J. van den Herik

Leden van de beoordelingscommissie:
Prof. dr. A.J. van Zanten (voorzitter)
Prof. dr. J. van Heerden
Prof. dr. J-J.Ch. Meyer (Universiteit Utrecht)
Prof. dr. J.M.J. Murre
Prof. dr. T. Ziemke (University of Skövde)



SIKS Dissertation Series No. 2005-19

The research reported in this thesis has been carried out under the auspices of SIKS,
the Dutch research School for Information and Knowledge Systems.

ISBN 90 5278 491 4

Universiteit Maastricht

©2005 Michel van Dartel

Cover design: Daradar, www.daradar.com

Preface

During my M.Sc. research I studied the behaviour of a model of a rat learning to reach a goal position in a maze. The virtual rat exhibited several intelligent strategies to reach the goal position. Being a naïve cognitive psychology student I did what most researchers in psychology do; analyse the rat's artificial brain for clues on the origins of its intelligent behaviour. After months of performing analyses, the results obtained consisted of some mere correlations between the recorded activation and the rat's position in the maze. However, the analyses did not provide any answers on how and where the rat represented knowledge about the maze and the goal position. I slowly started to realise that the rat's brain only held part of the answer, while the environment held the remainder.

One of the first things I (re)discovered in my Ph.D. research is that the current notion of knowledge representation is conceptually underdeveloped; this holds in particular for artificial intelligent systems situated in the real world. While the established view on cognition (the computational theory of mind) regards knowledge representation as a purely system-internal affair, the more recent view on cognition (the theory of situated cognition) explains knowledge representation in terms of the dynamic interaction between a system and its environment, i.e., in terms of a system's *situatedness*. So far, in artificial intelligence, the notion of knowledge representation has not grown with this development, but is still deeply rooted in the established view. The thesis conceptually improves the notion of knowledge representation in artificial intelligence and accomodates adequately the situated view on cognition. The improvement is formulated in the *operationalisation of situated representation* and is supported by several experiments with simulated robot models of cognitive behaviour. The operationalisation is a first step towards the conceptual improvement of the notion of knowledge representation, and paves the way for the general acceptance of the theory of situated cognition in artificial intelligence and the cognitive sciences.

This scientific contribution was only possible with the help of many friends and colleagues. During my four years as a Ph.D. researcher I worked at the Institute of Knowledge and Agent Technology (IKAT) in the Department of Computer Science of the Universiteit Maastricht. During this period I was fortunate to benefit from the guidance by Eric Postma and Jaap van den Herik. I am very grateful to both. At IKAT, I enjoyed many inspiring conversations with Rens Kortmann, Joyca Lacroix,

Guido de Croon, Ben Torben-Nielsen, Sander Spek, Steven de Jong, Paul Vogt, Dirk Bollen and many other colleagues and students, and received support from the secretariat and system administration. Looking back, I would like to thank Sandra Smeets, Herco Fonteijn, and Bram Bakker, for directing my scientific career in its early stages.

I am also very thankful for the support, good times, and cooperation in my semi-professional work to all my friends, especially to Ramon Schelleman who designed the cover of this thesis. To my parents and brother I am grateful for the freedom, opportunity, and support they have given me. Finally, I would like to express my gratitude to Swantje Langeheine for her love and support.

Michel van Dartel
Maastricht, 2005

Contents

Preface	v
Contents	vii
1 The nature of representation	1
1.1 The computational theory of mind	3
1.2 The theory of situated cognition	4
1.3 Problem statement and research questions	5
1.4 Methodology	6
1.4.1 Conditions for models of situated cognition	6
1.4.2 Three agent models	7
1.4.3 The essence of our investigation	9
1.5 Outline of the thesis	9
2 The notion of representation in situated systems	11
2.1 Situated representation	11
2.2 Internal representation	14
2.3 External representation	15
2.4 A new operationalisation of representation	16
2.5 Chapter summary	17
3 Situated representation	19
3.1 Perceptual ambiguity	19
3.2 Coping with perceptual ambiguity	20
3.2.1 Natural agents	20
3.2.2 Artificial agents	21
3.3 The active categorical perception model	22
3.3.1 The environment	22
3.3.2 The agent	23
3.3.3 The categorisation task	28
3.3.4 Perceptual ambiguity	28
3.3.5 The evolutionary algorithm	29

3.4	Experiments and results	30
3.5	Analyses of behavioural strategies	32
3.5.1	Observation of behaviour	32
3.5.2	Sensory state-transition diagrams	34
3.6	Discussion	38
3.6.1	Summary of results	38
3.6.2	Discussion of results	39
3.7	Chapter conclusions	40
4	Internal representation	41
4.1	Symbol manipulation in situated agents	41
4.2	The Tower of London task	43
4.3	The situated Tower of London model	45
4.3.1	The test	45
4.3.2	The agent's sensors and actuators	45
4.3.3	The agent's neurocontroller	50
4.3.4	The evolutionary algorithm	51
4.3.5	The performance measures	52
4.4	Experiment	53
4.5	Results	53
4.6	Analyses of behaviour	55
4.6.1	Produced behaviour	55
4.6.2	Internally simulated behaviour	60
4.7	Evaluation by comparison	63
4.7.1	Comparison with human subjects	63
4.7.2	Comparison with other situated models of the ToL task	65
4.7.3	Comparison with other internal simulation models	66
4.8	Discussion	66
4.9	Chapter conclusions	67
5	External representation	69
5.1	The basics of macroscopic analysis	69
5.2	The foraging model	71
5.2.1	The environment	71
5.2.2	The agent	72
5.2.3	The evolutionary algorithm	73
5.3	Experiments and results	74
5.4	A macroscopic analysis	74
5.5	A microscopic analysis	78
5.6	Discussion	79
5.6.1	Applying macroscopic analysis	80
5.6.2	Relation to the framework of scientific research	80
5.7	Chapter conclusions	82

6	The operationalisation of situated representation	85
6.1	Advantages of the new operationalisation	87
6.1.1	External representation and internal representation	87
6.1.2	The representation debate	88
6.1.3	Situated accounts of cognition and awareness	90
6.1.4	The symbol grounding problem	91
6.1.5	The relation between the operationalisations	92
6.2	Implications of the new operationalisation	92
6.2.1	Artificial Intelligence	92
6.2.2	Cognitive neuroscience	93
6.2.3	Cognitive psychology	94
7	Conclusion	97
	References	99
	Appendices	115
A	Environmental complexity of Acp	117
A.1	Experiments and results	117
A.1.1	Variable horizontal velocities	117
A.1.2	Variable vertical velocities	119
A.2	Discussion and conclusion	119
B	Categorisation through internal simulation	121
B.1	ACP*	121
B.1.1	Output-input feedback mechanism	121
B.1.2	Occlusion mechanism	123
B.1.3	Normalisation mechanism	123
B.2	Experiments and results	123
B.3	Discussion and conclusions	126
	List of abbreviations and symbols	129
	Summary	131
	Samenvatting	135
	Curriculum vitae	139
	SIKS Dissertation Series	141

Chapter 1

The nature of representation

Representation is a central concept in Artificial Intelligence (AI) (Schank, 1990; Brachman, Levesque, and Reiter, 1991; Brooks, 1991; Davis, Shrobe, and Szolovits, 1993). In the field of AI, a representation is generally considered an object that codes for relevant aspects of the world (Winograd, 1990). More specifically, according to Newell (1980, p.156), something is a representation when “an entity X designates an entity Y relative to a process P, if, when P takes X as input, its behaviour depends on Y”. Assume that X is a sign pole, Y is a speed limit, and P is the process of slowing down to the speed indicated by the sign pole, then the sign pole *represents*, i.e., is a representation of, the speed limit. Would X be a rule in someone’s mind concerning maximum speed learned in a traffic school, then X would be a *mental representation* of the speed limit.

There are two leading contemporary theories on the nature of mind; the more traditional *representational theory of mind* and the relatively novel *theory of situated cognition*.

The representational theory of mind holds that all mental states and activities involve mental representation (Field, 1978; 1994; Fodor, 1987; Markman and Dietrich, 2000b). According to the representational theory of mind, knowledge, perception, memory, dreams, hopes, fears, et cetera, all involve mental representation. For instance, when looking at Malevich’s *Black circle*, shown in figure 1.1, a circle is perceived because the CIRCLE-representation in the mind is activated by the image projected on the retina.

In the theory of situated cognition, the mind is regarded as a direct result of interaction with the environment (Clancey, 1997; Clark, 1997; Pfeifer and Scheier, 1999). According to the theory of situated cognition the (involuntary) movements of our eyes over Malevich’s *Black circle* are a prerequisite for perceiving the black circle. In other words, the pattern of change in the projection on our retinas as our eyes move over the painting is crucial for perceiving it, and not some CIRCLE-representation in the mind.

The two contrasting views on the nature of mind pose a challenge for AI. The

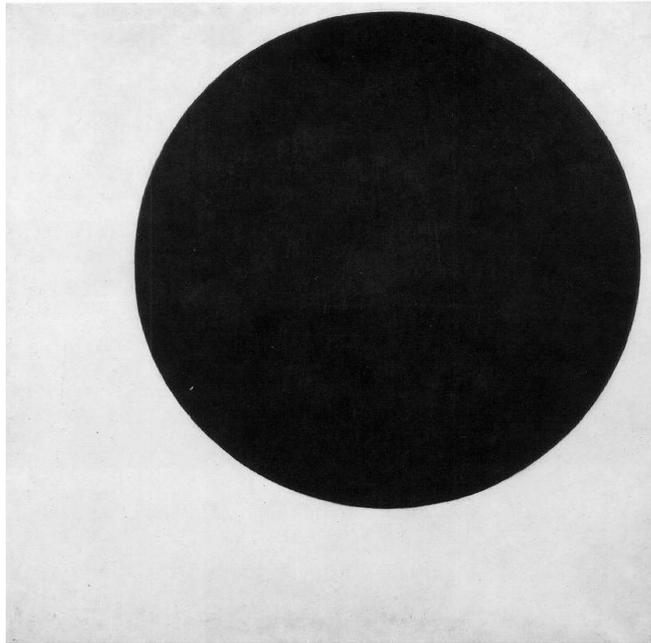


Figure 1.1: *Black circle* (1913, Oil on canvas) by Kasimir Malevich (1878-1935).

representational theory of mind traditionally has been embraced by AI because it allows a relatively easy formalisation of mental processes. In the theory of situated cognition it is more difficult to express mental processes in terms of common formalisms (such as logic). However, the latter theory offers a seamless incorporation of the interaction between an intelligent system and its environment. This thesis addresses the challenge posed by the incorporation of interaction in a theory of mind, by aiming at the formalisation of mental processes within the theory of situated cognition.

According to Beer (2003b), debates on cognition should be carried out in the presence of concrete examples. Therefore, the thesis' approach to the formalisation of mental processes within the theory of situated cognition is by implementing concrete examples of cognition. The concrete examples used are agent models (i.e., simulated robot models) of cognitive behaviour.

Section 1.1 briefly discusses the merits of the leading contemporary version of the representational theory of mind, i.e., the computational theory of mind. The theory of situated cognition is discussed in section 1.2. In section 1.3, the problem statement and the research questions of this thesis are formulated. Our approach towards answering these research questions is discussed in section 1.4. In section 1.5, the outline of the thesis is given.

1.1 The computational theory of mind

The computational theory of mind (Putnam, 1960; Fodor, 2000; Horst, 2003) is the leading contemporary version of the representational theory of mind. The theory holds that the brain is a kind of computer and that mental processes are computations on mental representations (Pitt, 2002). These mental representations are considered to be symbolic (Newell and Simon, 1976; 1976; Newell, 1980) or sub-symbolic (Smolensky, 1988) units of information. Newell and Simon (1985, p.37) claim that “symbols lie at the root of intelligence”, which is derived from their *physical symbol-system hypothesis*, which holds that:

“...intelligence is the work of symbol systems. Stated a little more formally, the hypothesis is that a physical symbol system [...] has the necessary and sufficient means for general intelligent action” (Simon, 1969, p.23).

The computational theory of mind makes the formal study of the mind feasible, because it allows formalisation of mental processes in computational models. Applications of computational modelling techniques that gained insight into mental processes include models of planning (Nilsson, 1984), REM sleep (Crick and Mitchison, 1983), speech (Sejnowski and Rosenberg, 1987), categorisation (Harnad, 1987), attention (Nosofsky, 1987; Shiffrin, 1988), recall and recognition (Ratcliff, Shiffrin, and Clark, 1990; Shiffrin, Ratcliff, and Clark, 1990), creativity (Boden, 1990), emotion (Morén and Balkenius, 2000), general intelligence (Laird, Newell, and Rosenbloom, 1987), and expert memory (Simon and Gilmarin, 1973).

For instance, by implementing a computational model of expert chess memory called MAPP (Memory-Aided Pattern Perceiver), Simon and Gilmarin (1973) provided insight into the mental processes concerning expert memory. MAPP formalised the hypothesis that expert chess players recall board positions by first recognising chunks on the board and then storing pointers to these chunks in a short-term memory. This hypothesis was inspired by the ‘chunking theory’ of expert memory, as formulated by Chase and Simon (1973a; 1973b). This chunking theory was based on empirical findings by de Groot (1946; 1965), who found that skill in chess depends on a player’s ability to store and relate useful patterns of chess positions, instead of the ability to store board positions in general. MAPP supported the chunking theory of expert memory by reproducing these empirical findings.

Although computational modelling of the mind has been a fruitful endeavor (Churchland and Sejnowski, 1992; Fodor, 2000), its fundamental assumption of (symbolic) representation has been questioned by, e.g., Brooks (1990), Thelen and Smith (1994), van Gelder (1995), Clancey (1997), and Dreyfus (2002a; 2002b) in the theory of situated cognition.

1.2 The theory of situated cognition

The theory of situated cognition opposes the representational theory of mind by regarding thinking as an activity, i.e., as interaction with the world, instead of as passive manipulation of symbols. To illustrate this, Clancey (1997, p.2) writes:

“In bicycling, every twist and turn of the steering wheel and every shift in posture are controlled not by manipulation of the physics equations learned in school, but by a *recoordination* of previous postures, ways of seeing, and motion sequences. Similarly, in reasoning, as we create names for things, shuffle around sentences in a paragraph, and interpret what statements mean, every step is controlled not by rote applying grammar descriptions and previously stored plans, but by adaptively reorganizing previous ways of seeing, talking, and moving.”

Models and theories that incorporate the interaction between an intelligent system and its environment show that intelligent behaviour and mental processes can be explained in terms of interaction (see, e.g., Wiener (1948), Gibson (1979), Maturana and Varela (1980), Bickhard and Richie (1983), Bickhard (2004), Agre (1995), Port and van Gelder (1995), Juarrero (1999), Beer (1990; 1995; 2000), O’Regan and Noë (2001b), Noë (2004)). Incorporating interaction in models of mental processes provided insight into selective attention (Slocum, Downey, and Beer, 2000; Beer, 2003b), categorisation (Beer, 1996; Nolfi and Marocco, 2001a; van Dartel *et al.*, 2005), motor development (Thelen, 1995), and various low-level behaviours (Beer, 1995; Pfeifer and Scheier, 1999; Nolfi and Floreano, 2000; Kortmann, 2003; van Dartel *et al.*, 2004b).

A particularly illuminating example of incorporating interaction in an explanation of mental processes is O’Regan and Noë’s (2001b) sensorimotor account of vision and visual consciousness, which explains the experience of seeing by the exploratory activity of the eyes. According to O’Regan and Noë, incorporating such exploration of the world can explain various perceptual phenomena, such as sensorimotor adaptation, visual filling in, visual stability despite eye movements, change blindness, colour perception, and sensory substitution. A particular instance of sensory substitution is tactile visual substitution, in which visual stimuli are transformed into tactile stimuli by a tactile visual substitution system (TVSS). A TVSS typically translates the luminance distribution of images captured by a camera into corresponding vibrations or electric signals on a skin area. If the TVSS’ camera moves in close correspondence to the eye and head movements, blind subjects have been reported to have the experience of seeing (see, e.g., Kay (1984), Easton (1992), Bach-y-Rita (1995)). This finding indicates that exploratory activity of the eyes, i.e., interaction with the world, is crucial for seeing, and not some (symbolic) representation in the subjects’ mind. In other words, interaction with the world seems to lie at the root of intelligence, and not symbols, as claimed by Newell and Simon (1976; 1985) in their physical symbol-system hypothesis.

1.3 Problem statement and research questions

In the representational theory of mind, representations are considered to be the symbolic (Newell and Simon, 1976; 1985; Newell, 1980) or sub-symbolic (Smolensky, 1988) units in the ‘brain’ of a system. They can be isolated from the system that uses them (Clancey, 1997). In the theory of situated cognition, representations cannot be readily isolated from the system that uses them, because their semantics cannot be analysed independently of the system itself and its ecological niche (see, e.g., Clancey (1997), Morrison (1998), Keijzer (2001)). Consequently, the problem that contemporary AI and cognitive science face leads us to the following problem statement.

What is the nature of representation in situated systems?

Within the theory of situated cognition, two contrasting views on the role of representation for cognition exist. Some advocates of the theory argue that the mind does not use representation at all, and the notion of representation should be discarded (Brooks, 1990; 1991; Thelen and Smith, 1994; van Gelder, 1995). Others argue that the notion of representation should be reconsidered to fit the theory (Bickhard, 1997; 2004; Clark, 1997; Cliff and Noble, 1997; Bechtel, 1998; Markman and Dietrich, 2000a, 2000b; Keijzer, 2001; 2002). As Haselager, de Groot, and van Rappard (2003, p.5) put it, the contrasting views result from the lack of “a proper operationalisation of the notion of representation”. The aim of this thesis is to provide cognitive science with a proper operationalisation of the notion of representation in situated systems. A proper operationalisation should enrich debates in the cognitive sciences on situated systems with a working definition of the notion of representation that clearly indicates when something is a representation.

Cummins (1989) suggests a distinction between (i) the *scientific* problem of representation, and (ii) the *philosophical* problem of representation. While the philosophical problem concerns the definition of mental representation (i.e., it answers the question (Cummins, 1996, p.1): “What is it for one thing to represent another?”), the scientific problem concerns the content, form, and implementation of representation. This thesis will focus on the scientific problem of representation, by studying the content, form, and implementation of representation in situated systems.

For our purposes, the problem statement stated above is translated into two research questions.

(i) *To what extent can we identify where the knowledge resides that is used by a situated system to perform a certain task?*

and

(ii) *How is this knowledge accessed and used by a situated system when performing a certain task?*

To answer these two research questions we employ the methodology of evolutionary robotics (Nolfi and Floreano, 2000).

1.4 Methodology

Below, we provide insight into our research approach. It is well known that the choice of a methodology may influence the results (Goodwin, 1995). Therefore, we consider it a matter of scientific soundness to state our point of departure clearly and to mention which conditions our approach contains. In subsection 1.4.1, the conditions that a concrete example of situated cognition, in our case an agent model, should satisfy are discussed. Subsection 1.4.2 introduces the three agent models that are employed as concrete examples in the following chapters. Subsection 1.4.3 discusses how these models are employed to answer the research questions stated in section 1.3.

1.4.1 Conditions for models of situated cognition

Models employed for debates on situated cognition should satisfy five conditions. A suitable model should address (i) *situatedness*, (ii) *embodiment*, and (iii) *cognition*, and should be characterised by (iv) *parsimony* and (v) *transparency*. Below these conditions are discussed in more detail.

(i) Situatedness

According to Clancey (1997, pp.1-2), “human thought and action is adapted to the environment, that is, *situated*, because what people *perceive*, how they *conceive of their activity*, and what they *physically do* develop together”. Therefore, a cognitive system, in a model that serves as a concrete example of cognition, should be able to act upon its environment and also be able to perceive the consequences of its own action, i.e., it should be *situated*.

(ii) Embodiment

A prerequisite for situatedness is that a system has a body to interact with the world, i.e., that it is *embodied*. Hence, embodiment and situatedness are strongly connected. In this thesis we follow Pfeifer and Scheier (1999) in that embodiment refers to the existence of a structural coupling of a system with the environment, rather than to having a physical body. This implies that embodiment can be simulated, as was demonstrated by Nolfi and Floreano (2000). It should be noted that a wide range of notions of embodiment exist (see, e.g., Ziemke (2003; 2004) and Chrisley and Ziemke (2002)).

(iii) Cognition

Obviously, a system in a model of cognition should have *cognition*. According to Matlin (1994, p.2), cognition “involves the acquisition, storage, retrieval, and use of

knowledge”. As stated above, cognition is traditionally assumed to involve representation.

(iv) Parsimony

Beer (2003b) suggests that the best examples for debates on cognition are models of ‘minimally cognitive behaviour’ (see, e.g., Beer (1990; 1995; 1996; 2000; 2003b); Slocum *et al.* (2000), Goldenberg, Garcowski, and Beer (2004) for examples of minimal cognitive behaviour). The simpler the model is, given that it still satisfies the other conditions for models employed for debates on situated cognition, the easier it is to analyse. In other words, models should be selected according to the rule of *parsimony*.

(v) Transparency

Parameters and changes in a system and environment of a model that is employed for debates on situated cognition should be available for analysis, i.e., a model that is employed for debates on situated cognition should be *transparent*.

1.4.2 Three agent models

In the following chapters we employ three agent models as concrete examples of cognition. These are used to discuss the problem statement formulated in section 1.3. The three models are all agent models of minimally cognitive behaviour, i.e., “the simplest behaviour that raises issues of genuine cognitive interest” (Beer, 2003b, p.213). In all three models, we employ the methodology of evolutionary robotics (Nolfi and Floreano, 2000, p.18), because “the possibility of evolving robots that are free to select their way to solve a task by interacting with their environment may help us to understand how natural organisms produce adaptive behaviour”. This implies that agent behaviour in each model is learned by evolutionary optimisation (see, e.g., Cliff, Harvey, and Husbands (1993), Harvey *et al.* (1997), Nolfi and Parisi (1999), Nolfi and Floreano (2000); Nolfi and Marocco (2001a, 2001b); Nolfi (2002b)).

Below, all three models are briefly described. Moreover, we review how these models satisfy the conditions stated in subsection 1.4.1.

(1) An agent model of active categorical perception (Acp)

The development of an agent model of active categorical perception (ACP) was inspired by Beer’s (2003b) model of active categorical perception. As in Beer’s (2003b) model, in ACP, agents are optimised to catch small and avoid large falling objects. By avoiding large objects and catching small objects an agent can express its ability to categorise these two classes of objects.

An agent in ACP is embodied because it has a structural coupling with the environment through its sensors and actuators, and it is situated because it has the ability to observe the consequences of its own actions and respond to them. If an

agent in ACP is able to categorise more objects correctly than the agent would do when its actions were based on chance it can be attributed cognition, because to do so it has to acquire, store, retrieve, and use knowledge.

Compared to other models of categorical perception (see Harnad (1987) for examples) ACP is a very parsimonious one. It is even more parsimonious than the model by which it was inspired, because, among other simplifications, the environment and the agent's neurocontroller are less complex than the model presented by Beer (2003b). Finally, ACP is maximally transparent, because it is designed in such a way that any parameter, either environmental or internal to the agent, is accessible for analysis. A more detailed description of ACP follows in chapter 3.

(2) An agent model of the Tower of London task

The agent model of the Tower of London task is employed because solving Tower of London problems is assumed to involve symbol manipulation (see, e.g., Cooper (2002)). In the model of the Tower of London task, agents are optimised to find the best solutions to Tower of London problems.

As in the ACP model presented above, an agent in the model of the Tower of London task is embodied because it has a structural coupling with the environment through its sensors and actuators, and it is situated because it has the ability to observe the consequences of its own actions and respond to them. To solve Tower of London problems, agents need to plan their actions, a skill that is generally regarded as high-level cognition (Cooper, 2002). Hence, when an agent in the model is able to solve Tower of London problems cognition can be attributed.

Compared to other models of the Tower of London task (see, e.g., Dehaene and Changeux (1997)) the model presented here is very parsimonious. Finally, the model of the Tower of London task is transparent, because, as ACP, it is designed in such a way that any parameter, either environmental or internal to the agent, is accessible for analysis. A more detailed description of the model of the Tower of London task follows in chapter 4.

(3) An agent model of foraging

The agent model of foraging is based on the idea that “what we call *intelligence* in animals and men has its origins in the solution to finding and consuming food, and staying alive” (Mobus, 1999, p.592). In the model of foraging, agents are optimised to collect as much food as possible.

As in the two agent models presented above, an agent in the model of foraging is embodied because it has a structural coupling with the environment through its sensors and actuators, and it is situated because it has the ability to observe the consequences of its own actions and respond to them. If an agent in the model is able to collect more food items than it would do when based on chance, cognition can be attributed, because to do so the agent has to acquire, store, retrieve, and use knowledge.

Compared to other models of foraging (see, e.g., Langton (1995)), the agent model of foraging presented here is a very parsimonious one. Finally, the model of foraging is transparent, because, as in both models presented above, it is designed in such a way that any parameter, either environmental or internal to the agent, is accessible for analysis. A more detailed description of the model of foraging follows in chapter 5.

1.4.3 The essence of our investigation

The agent models of cognition presented in subsection 1.4.2 are employed to answer the research questions presented in section 1.3.

The model of active categorical perception (ACP) is employed to investigate *where* the knowledge resides that is used to perform active categorisation. This issue is approached by testing to what extent reactive agents and non-reactive agents can cope with perceptual ambiguity and by analysing their behaviour under different task conditions and neural structures. Analysis of perceptual-state transitions in the agents is conducted to investigate the strategies that are employed by the agents to cope with perceptual ambiguity, i.e., *how* the represented knowledge to perform active categorisation is successfully used.

The Tower of London model is employed to investigate if a model of minimally cognitive behaviour can perform a task that requires symbol manipulation. If such symbol manipulation can be identified, it will reveal *where* the knowledge resides that is used to plan ahead in the Tower of London task, and *how* the represented knowledge is put to use by the agents in the model.

The agent model of foraging is employed to investigate whether a macroscopic analysis can reveal a universal property of adaptive behaviour. If such a universal property in the foraging behaviour of agents can be recognised, it answers the question of *where* the knowledge resides that is used to perform successful foraging behaviour. Furthermore, if such a universal property can be revealed, it localises the represented knowledge in the underlying statistics of the behaviour. An additional microscopic analysis of the evolved neural structure of an optimised agent is performed to investigate *how* represented knowledge is used by the agent.

1.5 Outline of the thesis

This chapter introduced the problem of the nature of representation in situated systems (in sections 1.1 and 1.2) and the two research questions derived from the problem stated in section 1.3. It also introduced the methodology adopted (in section 1.4), and the models employed (in subsection 1.4.2) to answer the two research questions.

The outline of the remainder of this thesis is as follows. Chapter 2 elaborates on the notion of representation in situated systems and discusses the models introduced in subsection 1.4.2 in the light of situated representation. In chapter 3, the ACP

model is described in more detail and is employed to answer both research questions of section 1.3 for this model by investigating to what extent reactive agents and non-reactive agents can cope with perceptual ambiguity. In Chapter 4, the model of the Tower of London task is described and employed to answer both research questions of section 1.3 for this model by investigating whether situated agents can solve Tower of London problems. Chapter 5 describes the model of foraging in more detail and employs the model to answer both research questions of section 1.3 for this model by investigating whether macroscopic analysis can reveal a universal property of adaptive behaviour.

A general discussion in chapter 6 formulates a new operationalisation of representation that applies to situated systems by combining the experimental results, and discusses the advantages and implications of this new operationalisation. Finally, chapter 7 provides a general answer to the two research questions and concludes on the problem statement by giving our contribution to the formalisation of mental processes within the theory of situated cognition.

Chapter 2

The notion of representation in situated systems

The term *situated representation* (Rosenschein and Kaelbling, 1996; Morrison, 1998) denotes representation that is associated with situated systems. Chapter 1 indicated that cognitive science is in need of an operationalisation of (the notion of) situated representation. To achieve such an operationalisation of situated representation, we first need to elaborate on the notion of situated representation. This chapter will discuss representation in situated systems (section 2.1) and the two types of representation that they use: internal representation (section 2.2), and external representation (section 2.3). This discussion will guide our empirical study of situated representation in artificial systems in chapter 3, chapter 4, and chapter 5. Section 2.4 discusses the three agent models, that were introduced in chapter 1 (see subsection 1.4.2), in the light of situated representation.

2.1 Situated representation

As stated in chapter 1, approaches based on the computational theory of mind explain cognition in terms of formal descriptions, such as rules (Newell and Simon, 1972) and scripts (Schank and Abelson, 1975). Rules and scripts represent (sequences of) actions that are typical for frequently experienced events. When the rules and scripts are applied instantly in response to the appropriate preconditions, the system acts in *open loop*, i.e., the system is controlled directly by an input signal, without the benefit of feedback (see figure 2.1). The theory of situated cognition considers open-loop systems insufficient to model cognition. In the words of Clancey (1995):

“...all processes of behaving, including speech, problem solving, and physical skills, are generated on the spot, not by mechanical application of

scripts or rules previously stored in the brain. Knowledge can be represented, but it cannot be exhaustively inventoried by statements of belief or scripts for behaving. Knowledge is a capacity to behave adaptively within an environment; it cannot be reduced to (replaced by) representations of behaviour or the environment.”

In other words, representations underlying cognition are part of *closed-loop* systems. A closed-loop system receives feedback by perceiving the results of its own output. The advantage of feedback is that a system can monitor whether it is making incremental progress towards its goal, rather than having to retain perfect knowledge to plan how to reach the goal before execution, as an open-loop symbol system has to do. The problem that open-loop systems face is that updating an internal symbolic

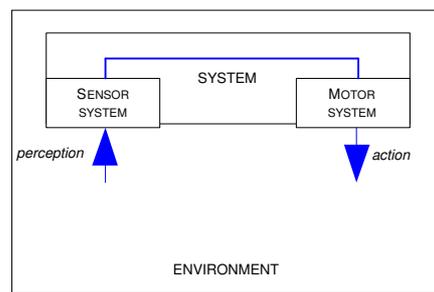


Figure 2.1: An open-loop system.

model is too costly in terms of computational resources to function in a continuously changing world (Norman, 1993). This was already illustrated more than twenty years ago by the robot Shakey (Nilsson, 1984), which was designed to navigate through an environment by means of a script-based planning system. Although Shakey was able to perform simple tasks, such as following a route, it could not cope with dynamic changes in the environment.

In contrast to symbolic representation, representation in situated systems does not involve perfect knowledge of the outside world, because situated systems are closed-loop systems. Situated representation encompasses sensory systems and motor systems, and is related to the way these systems are coordinated to function in a dynamic environment. As Clark (2001, p.88) states it:

“The internal representation of worldly events and structures may be less like a passive data structure or description and more like a direct recipe for action.”

Situated systems do not represent the knowledge in the sense-think-act cycle of ‘classic robotics’, in which a symbolic model of the external world has to be retained. They represent the knowledge that coordinates sensory systems and motor systems in

the closed loop of the ‘new robotics’ (see, e.g., Connell (1989) and Brooks (1991)), in which, “[f]ollowing the core ontology of interactivism, representations aren’t things, but rather, processes” (Morrison, 1998, p.206).

In psychology and related fields, the closed loop of situated systems is referred to as *the sensorimotor loop* (see, e.g., Thelen and Smith (1994)). The sensorimotor loop is illustrated in figure 2.2. In situated systems, the coordination between sens-

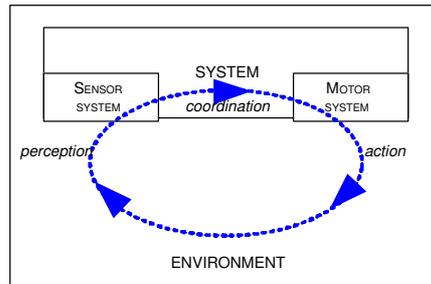


Figure 2.2: The sensorimotor loop.

ory systems and motor systems, i.e., *sensorimotor coordination*, underlies cognition (Braitenberg, 1986; Thelen and Smith, 1994; Pfeifer and Scheier, 1999; O’Regan and Noë, 2001b). Natural brains discover sensorimotor contingencies, i.e., the structure of changes in sensory input as a result of motor output (O’Regan and Noë, 2001b), which form the foundation for cognitive processes. For instance, according to O’Regan and Noë (2001b), the experience of seeing is caused by the mastery of the structural changes that occur on the retina in relation to the way we move our eyes (see section 1.2). By internalising this feedback relation between sensory systems and motor systems, a situated system (given the capable neural architecture) can think about objects or events in their absence (Hesslow, 2002). Figure 2.3 depicts the *internalised* sensorimotor loop.

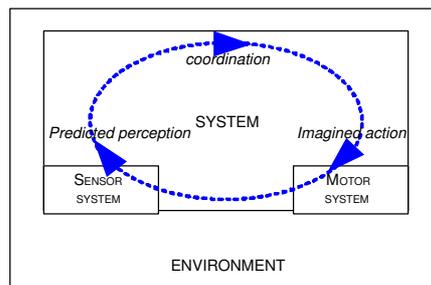


Figure 2.3: The internalised sensorimotor loop.

The relation between sensorimotor coordination and cognitive processes is largely unexplored, with a few notable exceptions (see, e.g., Thelen and Smith (1994), Port and van Gelder (1995), Pfeifer and Scheier (1999), Nolfi and Floreano (2000), O'Regan and Noë (2001b), Noë (2004)). The main reason is the dominance of the established view, which explains cognition in terms of symbol manipulation. Although symbol manipulation models may be fruitful to describe high-level processes in isolation, they do not reveal how cognitive processes are rooted in the interaction with the environment. Approaches based on the computational theory of mind ignore how the interaction with the environment contributes to the development and evolution of cognition (Norman, 1993; Noë and Thompson, 2002), they ignore the situatedness of cognitive systems.

The nature of representation depends on whether a situated system exhibits *reactive* behaviour or *non-reactive* behaviour. Reactive behaviour is defined as the motor patterns that are caused only by the immediate perceptual state of the animal (Balkenius, 1995, p.79). In contrast, non-reactive behaviour is defined as the motor patterns that are caused by something different from or more complex than the immediate perceptual state. Cognition is generally associated with the latter type of behaviour (Cruse, 2003). However, it has been shown that reactively behaving situated systems can perform simple cognitive tasks (Braitenberg, 1986; Nolfi and Floreano, 2000). The experiments to be reported in chapter 3 support these findings. Reactive agents that are able to perform cognitive tasks indicate that cognition can be produced on the basis of the immediate perceptual state of the system only, and therefore without representing information over time. An operationalisation of situated representation is required to account for these cases in which cognitive behaviour is produced both in the absence and presence of *internal* representation. Hence, we study both non-reactive and reactive systems in the experiments reported in the next chapters. When situated systems produce cognitive behaviour in the absence of internal representation, they are argued to represent *externally*. The next two sections will discuss the notions of internal representation and external representation, respectively.

2.2 Internal representation

As mentioned in section 2.1, situated systems can represent internally and represent externally (see, e.g., O'Regan (1992), Nolfi (2002b) for human examples and robot examples, respectively). The computational theory of mind puts emphasis on internal representation. The human mind does indeed have an impressive capacity to represent information internally. As a case in point, Luria (1968) demonstrated an exceptional case that suggests that the human mind may have an unlimited capacity to store information internally. He demonstrated this by giving his subject, a news reporter named S., very long strings of numbers, words, and nonsense syllables. Even years later, S. was able to recall these strings, generally without mistake.

However, S. was a mnemonist that utilised his synesthesia to achieve an excep-

tional performance on recall tasks. In general, humans need highly specific external cues to recall internally represented information (see, e.g., Tolman (1948)). This point has been forcefully stated by the behaviourist school of psychology (see, e.g., Skinner (1938)). The actual existence of internal representations that can be manipulated was experimentally shown in a classic study by Cooper and Shepard (1973). They measured how long it took students to determine whether a rotated letter was mirrored. Examples of the stimuli used by Cooper and Shepard are depicted in figure 2.4. They found that the time it took to determine whether a letter was



Figure 2.4: Examples of stimuli used by Cooper and Shepard (1973).

mirrored was directly proportional to the time it took to rotate them. For instance, it took students longer to decide whether the stimulus letter labelled ‘5’ in figure 2.4 was mirrored than it took them to answer the same question for the stimulus letter labelled ‘3’, because the rotation angle of the former was larger than that of the latter. According to Cooper and Shepard (1973) this indicates that, to perform this task, humans perform a mental rotation of the physical object. Hence, the object has to be represented internally to do so.

In general, humans do not store and manipulate internally all information that they want to recall at later times. In contrast to Luria’s (1968) subject S., humans are often confronted with the limits of their mental capacities (Norman, 1993). To cope with these limitations, humans use all kinds of external memory aids, such as grocery lists, electronic organisers, and calendars, that do not merely act as inputs and stimuli to internally stored information (Chambers and Reisberg, 1985; Zhang, 1997). In psychology and related fields, such external memory aids are called *external representations*.

2.3 External representation

External representation is defined by Zhang (1997, pp.179-180) as

“the knowledge and structure in the environment, as physical symbols, objects, or dimensions (e.g., written symbols, beads of abacuses, dimensions of a graph, etc.), and as external rules, constraints, or relations embedded in physical configurations (e.g., spatial relations of written digits, visual and spatial layouts of diagrams, physical constraints in abacuses, etc.)”.

Although external, the representations involved largely depend on perceptual systems to be picked up, analysed, and processed (Zhang, 1997). The latter is what makes external representations different from presentations; a perceptual system has to be able to pick up, analyse, and process them. In terms of the example given at the beginning of chapter 1 this implies that the sign pole only represents the speed limit when it is in view of the driver. Once the sign pole is out of view, it is debatable whether it still represents the speed limit. According to Zhang and Norman (1994), internal representation partly depends on knowledge in the world while external representation partly depends on knowledge in the mind. This may imply that knowledge is represented somewhere on the hypothetical continuum between external and internal representation depicted in figure 2.5. The symbol-system hypothesis of clas-

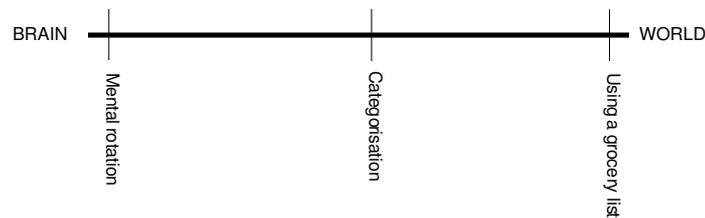


Figure 2.5: Representational continuum

sic AI does not relate internal representation to external representation. Symbolists argue that all knowledge has to be represented internally to mediate actions, which has resulted in a general ignorance of the importance of external representation for cognition (Zhang, 1997). An operationalisation of situated representation should account for representation at any location on the continuum between internal and external representation. In our investigations of situated representation in the next chapters we explore both forms of representation.

2.4 A new operationalisation of representation

In our quest towards understanding the nature of representation in situated systems, the interaction with the environment plays a pivotal role. Therefore, our method is a bottom-up approach addressing minimal cognitive behaviour (Beer, 1996; 2003b) in situated agents by adhering to the five conditions stated in chapter 1 (i.e., situatedness, embodiment, cognition, parsimony, and transparency), the reactive and non-reactive behaviour of situated systems (discussed in section 2.1), and the two types of situated representation (discussed in section 2.2 and section 2.3).

In an attempt to clarify the notion of representation in situated systems, we start by investigating simple agent-environment systems in which reactive agents and non-reactive agents are optimised to perform minimal cognitive behaviour.

In chapter 3 we investigate the notion of situated representation by focussing on an agent that has to decide whether to catch or avoid objects using incomplete or ambiguous sensory information. Implicitly, objects have to be categorised in order to make the right decision. Categorisation is a minimal cognitive behaviour because, it can be modelled in relatively simple models, but is often argued to be at the basis of cognition as a whole (Harnad, 1987; Beer, 2003b). Evolutionary optimised agents are analysed to clarify how the cognitive process of categorisation is rooted in the interaction with the environment, i.e., in their sensorimotor coordination. By evolving the performance of both reactive agents and non-reactive agents on the categorisation task, we expect the results to demonstrate the difference between internal representation and external representation in situated systems.

In chapter 4 we investigate the notion of internal representation by employing an agent-environment system in which an agent has to perform the Tower of London task. Since the Tower of London task is a task that is typically associated with internal symbol manipulation, we expect the acquired insight from this agent model to clarify the nature of internal representation in situated cognition.

In chapter 5 we investigate the notion of external representation by employing another agent-environment system in which agent foraging behaviour is optimised. The environment is composed of random distributions of food elements. Successful agents are tuned to the statistics underlying the environmental distribution of elements. We analyse the behaviour of optimised agents by averaging over many instances of agent-environment interactions. We expect the acquired insights from this agent model to emphasise the role of the environment and clarify the nature of the external representation in situated cognition.

2.5 Chapter summary

The symbol-system hypothesis (see section 1.1) suffices as an operationalisation of representation in non-situated systems in classic AI. However, symbolic representation cannot account for representation in a situated context. We discussed situated systems and the two types of representation that situated systems use: internal representation and external representation. We also indicated how this discussion guides our empirical investigations of representation in situated systems (chapter 3), internal representation (chapter 4), and external representation (chapter 5).

Chapter 3

Situated representation

In this chapter¹ we investigate a simple agent-environment system in which situated agents are optimised to perform minimal cognitive behaviour (cf., Beer (1996; 2003b)). By investigating the effective sensorimotor mapping of situated agents, we attempt to clarify the notion of situated representation. The perceptual ambiguity in the agent-environment system forces successful situated agents to represent adequately the information they need to perform the task, which provides a unique opportunity to study situated representations.

3.1 Perceptual ambiguity

Perceptual ambiguity occurs when identical sensory states require different responses. Such ambiguous sensory states are a problem for agents that have to select an appropriate action. Beer (1996) stated that only non-reactive agents, i.e., agents that base their actions on the current as well as past sensory states (see subsection 2.1), can cope with perceptually ambiguous tasks. Contrasting Beer's (1996) statement, Nolfi (2002b) showed that reactive agents, i.e., agents reacting directly and exclusively to the current sensory state (see subsection 2.1), are able to cope with perceptual ambiguity as well. Nolfi's finding agrees with the observation that numerous animals behave reactively (Tinbergen, 1951; Lorenz, 1973; Balkenius, 1995) and are still able to cope with perceptual ambiguity, presumably omnipresent in their natural environments. The discrepancy between Beer's statement and Nolfi's finding indicates a difference in opinion concerning the nature of representation in situated systems.

¹This chapter is reprinted by permission of Sage Publications Ltd from van Dartel M.F., Sprinkhuizen-Kuyper I.G., Postma E.O., and van den Herik H.J. (2005) Reactive agents and perceptual ambiguity. *Adaptive Behavior*, Copyright (©International Society of Adaptive Behavior, 2005). The experiments and results reported here were presented at the 15th Belgium-Netherlands Conference on Artificial Intelligence in Nijmegen, The Netherlands (see van Dartel *et al.* (2003)). The author would like to thank his co-authors and the publishers for their kind permission to reuse relevant parts of the articles in this thesis.

This chapter focuses on the discrepancy between Beer’s statement and Nolfi’s finding by addressing two research questions: (1) to what extent can reactive and non-reactive agents cope with perceptual ambiguity, and (2) which behavioural strategies are employed when doing so? To answer these questions, the model of active categorical perception¹ inspired by Beer’s (2003b) model briefly discussed in chapter 1, is used.

The outline of this chapter is as follows. In section 3.2, we discuss the discrepancy between Beer’s statement and Nolfi’s finding in detail. In section 3.3, the active categorical perception model is presented. The model emphasises the active nature of perception in the agents studied. Through simulation studies, we evaluate the extent to which reactive and non-reactive agents are able to cope with different levels of perceptual ambiguity, and investigate which behavioural strategies are employed when doing so. Section 3.4 describes the experiments conducted and reports on the results. In section 3.5, the behaviour of optimised agents is analysed. Finally, the results are discussed and chapter conclusions are given in section 3.6.

3.2 Coping with perceptual ambiguity

This section provides the context to better understand the discrepancy between Beer’s statement and Nolfi’s finding. In subsection 3.2.1 and subsection 3.2.2, we discuss how natural agents and artificial agents cope with perceptual ambiguity, respectively.

3.2.1 Natural agents

Natural agents are experts in coping with perceptual ambiguity. For instance, they cope with perceptual ambiguity when confronted with camouflaged preys (Brower, 1958) or opponents (Hartcup, 1979). In everyday life, perceptual ambiguity occurs in recognition tasks such as face recognition (see, e.g., Postma, van Dartel, and Kortmann (2001), Lacroix *et al.* (2004)). As was stated in section 2.3, many researchers have shown that natural agents can adequately cope with perceptual ambiguity by use of external representation, i.e., by the use of the environment as an external memory (Chambers and Reisberg, 1985; O’Regan, 1992; Hutchins, 1995; O’Regan and Noë, 2001b; Gray and Fu, 2004). External memory is often referred to as the use of objects or characteristics in the environment to alleviate the internal-memory load (see, e.g., Kirsh and Maglio (1992), Kirsh (1995), Hutchins (1995)). Recently, O’Regan and Noë (2001b) argued that the complete environment of a natural agent may serve as an external memory. To illustrate this, O’Regan and Noë discuss the blind spot as an example.

The blind spot of an eye is the part of the retina where no photosensitive cells are present (because there the optic nerve leaves the eye). The blind spot results in a ‘gap’ of 5 to 7 degrees in the human visual field, creating perceptual ambiguity

¹Available for download at <http://www.cs.unimaas.nl/mf.vandartel/ACP.htm>

by partially or completely occluding objects from sight. In general, humans are not aware of the fact that objects are often partially or completely occluded from sight by their eyes' blind spot. It is argued that, to cope with the perceptual ambiguity caused by the blind spot, some mechanism in the brain 'fills in' the occluded part of the visual field (cf., Ramachandran and Gregory (1991)). However, O'Regan and Noë (2001b) argue that the perceptual ambiguity caused by the blind spot is coped with by simply moving the eyes when the information is needed. In this way, the environment is used as if it were an external memory (O'Regan, 1992; O'Regan and Noë, 2001b). It is of interest to researchers in artificial intelligence and cognitive science whether artificial agents can use similar strategies.

3.2.2 Artificial agents

Beer (1996) performed experiments in which artificial agents were optimised to catch downward falling circles by moving horizontally at the bottom of an environment. The circles fell down at a variable vertical and horizontal velocity. Agents were able to sense the falling circles by the use of upward directed 'sensor rays'. Beer (1996) observed that optimised reactive agents were outperformed by non-reactive agents optimised on the task. An analysis of the behaviour of the reactive agents revealed that they were unable to cope with objects moving beyond their sensor range. When objects move beyond an agent's sensor range, a sensory state is ambiguous to an agent. Beer (1996) argued that reactive agents lack the required internal dynamics to cope with perceptual ambiguity, but that non-reactive agents can cope with this ambiguity because their internal dynamics allow them to organise "their behavior according to sensory stimuli that are no longer present" (*ibid*, p.424). Probably for this reason, Beer (1996; 2003b) did not investigate the performance of reactive agents on the active categorical perception task in his subsequent experiments.

In contrast to Beer, Nolfi (2002b) reported that reactive agents can cope with perceptual ambiguity. He studied a model in which reactive agents were positioned in a circular environment. The environment consisted of 40 cells positioned on the circle, 20 on each semicircle. The task for the agents was to end up in the left semicircle when starting from an arbitrary initial position. Cells on both semicircles were assigned a random number in the range $[0, 19]$, and could be sensed by the agent. Depending on the number, the agent had to decide whether to move one cell clockwise or counterclockwise. Nolfi's task is perceptually ambiguous (i.e., a so-called perceptual aliasing problem (Nolfi and Parisi, 1999; Nolfi, 2002b)), because each number occurred in both semicircles. Hence, the number did not reveal in which semicircle the agent was located. However, Nolfi (2002b) found that some of the reactive agents evolved successful behavioural strategies by the use of 'attractor states' in the left semicircle. In the task, an attractor state corresponds to a pair of cell values that evoke opposite actions such that the agent is captured between the cells in the left semicircle. The same cell values appear in the right semicircle and act as a 'repellor state' by evoking movements towards the left semicircle. In this way, agents always moved back and forth between the two cells of the pair while in

the left semicircle, and away from both cells of the pair while in the right semicircle.

Comparing Beer’s (1996) statement with Nolfi’s (2002b) finding raises two questions: (1) to what extent can reactive agents deal with perceptual ambiguity, and (2) which behavioural strategies are employed by reactive agents when doing so? To address these questions, we introduce the active categorical perception model.

3.3 The active categorical perception model

The active categorical perception model (ACP), introduced in subsection 1.4.2, is based on the extension of Beer’s (1996; 2003b) model that was briefly described in section 3.2. ACP differs from Beer’s model in several details. The most important differences are in the sensory mechanism (an additional type of sensor and binary sensory states). These differences are the result of our wish to control the level of perceptual ambiguity in ACP. As in Beer’s (1996; 2003b) model, in ACP agents are optimised to catch and avoid falling objects. Categorical perception is active in the model; it is established via an effectual interaction between agent and environment (Bajcsy, 1988; Aloimonos, Weiss, and Bandopadhyay, 1988). Active perception requires that the model is situated (Clancey, 1997), i.e., that the agent is able to interact with its environment. Below, the model is outlined in terms of environment (section 3.3.1), agent (section 3.3.2), categorisation task (section 3.3.3), perceptual ambiguity (section 3.3.4), and evolutionary algorithm (section 3.3.5).

3.3.1 The environment

A two-dimensional grid G_t of size $x_{max} \times y_{max}$ defines the environment in which the agent acts at time t . For all experiments we set $x_{max} = 20$ and $y_{max} = 10$. The objects and agents are allowed to move through the left and right boundaries of the environment, defined by $x = 0$ and $x = x_{max} - 1$, and to re-appear at the opposite side of the environment (i.e., the environment is defined as a cylinder). In the model, time passes in discrete steps. Each object is placed somewhere in the row $y = y_{max} - 1$ at $t = 0$. By convention we mention the leftmost x -coordinate of the object as the starting position of the object. The law of gravity in the model causes objects to fall. Objects hit the floor defined by $y = 0$ at $t = y_{max} - 1$. Note that the place of an object and the time are related by $t = y_{max} - 1 - y$. An object, located at position (x, y) at time t , is represented by a sequence of ones: $G_t((x + j) \bmod x_{max}, y) = 1$ for $j \in \{0, 1, \dots, j_{max} - 1\}$, with j_{max} the width of the object. All other cells in grid G_t have a value of 0. Two classes of objects are defined: small objects ($j_{max} = 2$) and large objects ($j_{max} = 4$). For the initial horizontal position of an object, x is selected from $x \in \{0, 1, \dots, x_{max} - 1\}$. At time $t = 0$ an object is placed at position (x, y_{max-1}) in the top row of grid G_0 . Then, all grids G_t for $t > 0$ are defined as

$$G_{t+1}(x, y) = G_t((x - 2d) \bmod x_{max}, y + 1) \quad (3.1)$$

for $(x = 0, 1, \dots, x_{max} - 1)$ and $(y = 0, 1, \dots, y_{max} - 2)$, with $d \in \{-1, +1\}$ a direction parameter that is fixed between $t = 0$ and $t = t_{max}$. The object defined by the sequence of ones in G_t moves to the left for $d = -1$ and to the right for $d = +1$. Figure 3.1 illustrates the movement of an agent and object in the environment (grid G_t) over four consecutive simulation time steps (denoted by $t = 4$ to $t = 7$). The large object (represented by 4 black grid cells) falls to the left. The four circles in the bottom row of each grid represent the sensors of the agent (described in subsection 3.3.2); they are activated (gray circles) by the presence of an object in the same column. In the figure, the agent moves four grid cells to the left in each time step. The movement of the agent depends on the activation of the sensors and on the structure of the agent.

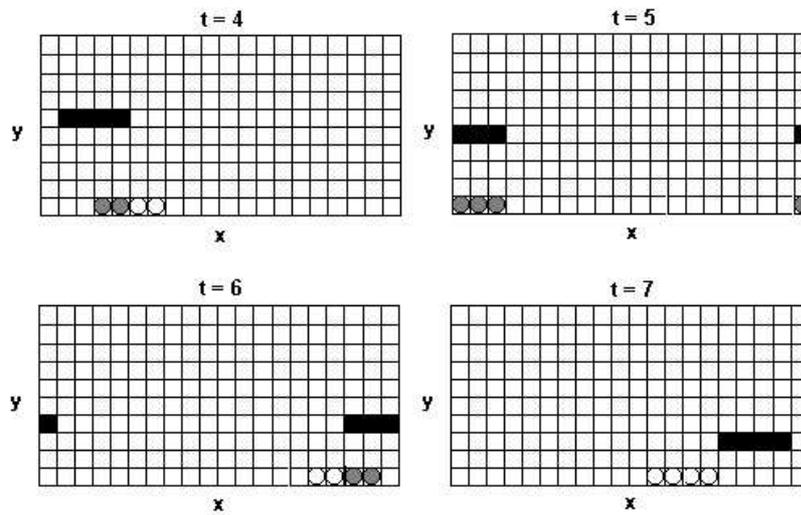


Figure 3.1: Movement of an agent and object over four consecutive simulation time steps (from $t = 4$ to $t = 7$). Black grid cells represent the object and circles represent the active (gray) and inactive (white) sensors of the agent.

3.3.2 The agent

The agent consists of a neurocontroller that receives environmental input through an array of sensors. A motor system moves the agent according to the output of the neurocontroller, resulting in the agent either catching or avoiding the object.

Neurocontroller

The ability of reactive agents to cope with perceptual ambiguity is studied by using two feed forward neurocontrollers; (i) a perceptron and (ii) a multilayer perceptron. The analogous ability of non-reactive agents is studied by using a recurrent neurocontroller (Elman, 1990). The three different neural networks correspond to the three most common variants of simple neural networks. Two of these variants produce reactive behaviour, the perceptron and the multilayer perceptron, while the third, the recurrent neural network, can produce non-reactive behaviour. The topologies of the three neural networks are illustrated in figures 3.2, 3.3, and 3.4. In all three figures rectangles represent the different layers of the neural network, and circles represent the nodes within each layer. A straight arrow indicates full connectivity between layers; a bend arrow indicates a one-to-one copy of activation between layers. Below, we briefly describe the neural networks.

Figure 3.2 illustrates the topology of the perceptron (P). The P consists of an input layer of four nodes, one for each sensor. The input nodes are all connected to a single output node providing the motor output (i.e., the action) for the agent. Figure 3.3 illustrates the topology of the multilayer perceptron (MLP). The MLP has a similar structure, but is extended with a hidden layer of four nodes that is fully connected to both the output and input layer. The P-controlled and MLP-controlled agents cannot store any information internally, i.e., they are reactive. Figure 3.4 illustrates the topology of the recurrent neural network (RNN). In the RNN four context nodes are added to the structure of the MLP. The context nodes are fully connected to the hidden layer. These nodes contain exact copies of the activation of the hidden nodes at time $t - 1$, which is used as additional input to the hidden layer at time t . The activation of all context nodes is initialised to zero at $t = 0$. Recurrent connections enable an agent to retain past sensory information over time. This allows RNN-controlled agents to behave non-reactively. In all three networks a bias node with a constant output of 1 is connected to the hidden nodes, if present, and the output node. The activation function of the input nodes is linear, i.e., activation of input nodes is equal to the binary input received by the sensors. Context nodes and output nodes also have linear activation functions. However, the activation of the hidden nodes is determined by the hyperbolic tangent (*tanh*) of the *net* input. The real-valued connection weights of the neural networks are initialised according to a uniform distribution with $[-1, 1]$, and are optimised by the evolutionary algorithm described in subsection 3.3.5.

Array of sensors

The sensor array of the agent consists of two types of sensors: functional sensors and blind sensors. Each agent has s functional sensors ($s = 4$, for all experiments outlined in section 3.4), with each functional sensor connected to one input node of the neurocontroller. A variable number of blind sensors b ($b \in \{0, 1, 2, 3\}$), i.e., sensors that are not connected to the input layer, are inserted into the sensor array; they

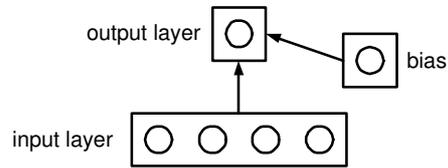


Figure 3.2: Topology of the perceptron (P).

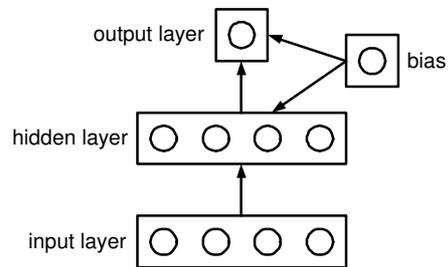


Figure 3.3: Topology of the multilayer perceptron (MLP).

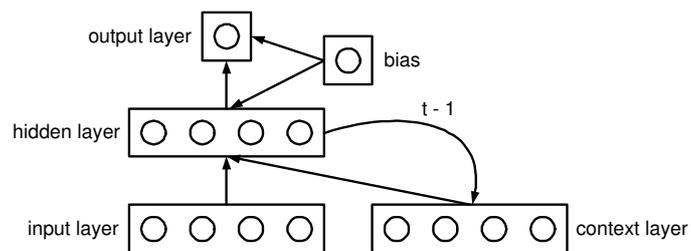


Figure 3.4: Topology of the recurrent neural network (RNN).

determine the level of perceptual ambiguity (see subsection 3.3.4). Blind sensors are always positioned in the middle of the sensor array. The sensors occupy neighbouring grid cells and are constrained to the bottom row of the grid G_t , $G_t(x, 0)$ for all t . The activation of the sensor at position x at time t is represented by $I(x, t)$ and is defined as:

$$I(x, t) = \sum_{y=0}^{y_{max}-1} G_t(x, y) = G_t(x, y_{max} - 1 - t) \quad (3.2)$$

An agent with four functional sensors and two blind sensors is illustrated in figure 3.5. The two leftmost sensors are activated ($I(x, t) = I(x + 1, t) = 1$) by the object. We note that the agent's sensory state does not contain any information regarding its own position or its distance to the object.

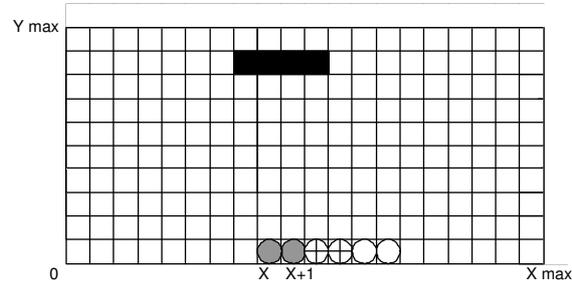


Figure 3.5: An agent with six sensors, of which two are blind sensors (crossed circles), sensing a large object (black cells) with two of its functional sensors (gray circles).

The motor system

In order to catch and avoid objects, agents can move to the left, move to the right, or stand still in the bottom row of G_t . The real-valued output of the neurocontroller rounded to the nearest integer value, expressed as a , defines the size and direction of the next step taken by the agent, i.e., the number of grid cells moved to the left (negative output) or right (positive output). If $a = 0$, the agent does not move. For a sensor positioned at x , the new position after movement is defined by:

$$(x + a) \bmod x_{max} \quad (3.3)$$

A movement of the agent leads to a new position of the sensor array and, consequently, to a new sensory state determined by the new position of the agent and equation 3.2.

Catch or avoidance behaviour

The behaviour of an agent is evaluated when the object reaches the bottom row of grid G_t at $t = y_{max} - 1$. Because the environment is defined as a cylinder, the bottom row forms a circle with x_{max} different positions (i.e., $x \in \{0, 1, 2, \dots, x_{max} - 1\}$). The relative distance between the centre of the object and the centre of the agent in ACP is defined as the smallest angle α between both centre positions, where $\alpha \in \{0, u, 2u, \dots, (x_{max}/2)u\}$, with u representing the (angular) unit in degrees, i.e., $u = \frac{360}{x_{max}}$. An object is caught at time $t = y_{max} - 1$ by an agent (with any number of blind sensors) if and only if $\alpha \leq 4.5u$, and avoided otherwise.

Setting the criteria for behavioural evaluation in this way ensures that objects that activate sensors at $t = y_{max} - 1$ are evaluated as being caught for agents with any number of blind sensors b . This is illustrated in figure 3.6, which depicts the relative distances between agent and object for two border cases of avoiding and catching, respectively. In subfigure 3.6(a), an agent with three blind sensors ($b = 3$)



Figure 3.6: Illustration of a small difference in separation between agent and object that causes a change in behaviour. Left: $\alpha = 5.5u$ results in avoiding the object at $t = y_{max} - 1$. Right: $\alpha = 4.5u$ results in catching the object at $t = t_{max} - 1$.

has a relative distance α of $5.5u$ from the object, which results in avoiding the object if it occurs at $t = y_{max} - 1$. A separation of a single unit (grid cell) less, as depicted in subfigure 3.6(b), yields catch behaviour if $t = y_{max} - 1$. Besides the cases in which sensors of the agent are activated at $t = y_{max} - 1$, there are cases in which objects that do not activate sensors at $t = y_{max} - 1$ are evaluated as being caught. This can only occur when agents have less than three blind sensors ($b < 3$), because the evaluation criterium ($\alpha \leq 4.5u$) for catching is constant for each number of blind sensors, while the horizontal range of an agent's sensor array decreases when the number of blind sensors decreases. An example is illustrated in figure 3.7.

Since the centre of an object is used to determine α , the relative distance is independent from the size of an object. This independence of object class (small or large) results in an equal chance to avoid an object from both classes. Similarly, it results in an equal chance to catch an object from both classes. So, if an agent behaves randomly, i.e., when its behaviour is not influenced by sensory information, it will catch the same number of small objects as it will catch large objects, and avoid the same number of large objects as it will avoid small objects. Since the

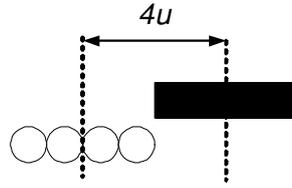


Figure 3.7: Illustration of an example case in which an object is caught by the agent, i.e., $\alpha \leq 4.5u$, if $t = y_{max} - 1$, although no sensors of the agent are active.

same value is subtracted from the fitness of an agent for catching an object from the incorrect class than is added to the fitness of an agent for catching an object from the correct type, and both types of objects are encountered an equal number of times, the fitness of such an agent will be equal to zero and its performance will be equal to 0.5 (see subsection 3.3.5 for details on fitness and performance). Maintaining the same behavioural evaluation criteria for agents with different values of b enables comparison between the different conditions (as done in section 3.4).

3.3.3 The categorisation task

The agent's task is to categorise the two classes of objects (small and large). agents are optimised to avoid large objects and to catch small objects. By correctly doing so an agent exhibits its ability to categorise.

3.3.4 Perceptual ambiguity

A sensory state is ambiguous to the agent when the state can result from the presence of an object from either object class (small or large). Hence, categorisation is not possible on the basis of a single ambiguous sensory state. In figure 3.1 only the sensory state at time $t = 5$ is not ambiguous to the agent, since this state can only result from the presence of a large object.

The number of blind sensors inserted in the middle of the sensor array determines the level of perceptual ambiguity. When the number of blind sensors is increased, the proportion of ambiguous sensory states increases too, i.e., it becomes more difficult for the agent to categorise objects. Figure 3.5 is an illustration of how the insertion of two blind sensors (crossed circles) makes a sensory state ambiguous: the two activated sensors (gray circles) may be caused either by a large or by a small object.

Table 3.1 shows the percentages of unique and ambiguous sensory states for the different numbers of blind sensors ($b \in \{0, 1, 2, 3\}$). Only nine or fewer different sensory states (depending on the number of blind sensors, as shown in table 3.1) can occur in the agent's sensor array during the task. This enables determination of the perceptual ambiguity in the model and, as will be shown, makes analysis of the sensorimotor mapping of optimised agents feasible.

number of blind sensors	number of possible sensory states	number of unique sensory states	number of ambiguous sensory states
$b = 0$	9	4 (44.4%)	5 (55.6%)
$b = 1$	9	4 (44.4%)	5 (55.6%)
$b = 2$	8	3 (37.5%)	5 (62.5%)
$b = 3$	7	0 (0%)	7 (100%)

Table 3.1: Number of possible, unique, and ambiguous sensory states for b blind sensors ($b \in \{0, 1, 2, 3\}$).

3.3.5 The evolutionary algorithm

The behaviour of the reactive and non-reactive agents depends on the weight values of their constituent neurocontrollers. In ACP we employ an evolutionary algorithm to optimise the behaviour of agents, since de Croon, van Dartel, and Postma (2005b) show that an evolutionary algorithm outperforms reinforcement learning techniques on a task similar to ACP. The evolutionary algorithm determines the weight values for all connections in the agent’s neurocontroller. After randomly initialising the weights of the neurocontrollers, the complete generation is tested on the active categorical perception task described above. An agent’s fitness F , i.e., the success of a tested agent, is calculated as:

$$F = (CC + CA) - (FC + FA) \quad (3.4)$$

with CC the sum of correctly caught objects, CA the sum of correctly avoided objects, FC the sum of caught objects that should have been avoided, and FA the sum of avoided objects that should have been caught. agents are tested on 80 trials, all possible starting positions times the number of object types times the number of directions in which objects can fall ($x_{max} \cdot 2 \cdot 2$). The performance of an agent is expressed by its success rate ($\in [0, 1]$), which is calculated by $(F + 80)/(2 \cdot 80)$.

During the experiments described in section 3.4, the following settings for the evolutionary algorithm were maintained; the number of generations was set to 20,000 and each generation was set to consist of 100 agents.

For the evolutionary algorithm to operate, an agent’s neurocontroller is represented in a genome by placing its connective weights at arbitrary positions in a vector. The algorithm uses the standard evolutionary techniques of reproduction, crossover, and mutation (Goldberg, 1986). Parents are selected by taking the five best-ranked agents and adding thirty winning agents from tournaments (Miller and Goldberg, 1995) of size three held among randomly selected agents. Sixty-four new agents are created by one-point crossover on the basis of random selection among the parents; two hundred mutations are performed over the complete new generation. The exact location of a one-point crossover in the genome is determined by random selection of a point in the genome. A mutation is performed by selecting a random weight in the genome of a randomly selected agent and adding a random value in the range

$[-0.2, 0.2]$ to the weight. Elitism is introduced to the algorithm by adding the best agent of the last generation to the new generation (Goldberg, 1986).

By constraining the input weights of all neurocontrollers we reduce the evolutionary search space. For all agents the value of the i -th input weight to the left of the centre is defined to be equal to minus the value of the i -th input weight to the right of the centre ($i \in \{1, 2\}$).

3.4 Experiments and results

Experiments were performed with each combination of neurocontroller (P, MLP, or RNN) and number of blind sensors ($b \in \{0, 1, 2, 3\}$). So, 12 different experiments were conducted. Each experiment was executed 15 times, over which the success rates of the best-performing agents were averaged. Comparing the results of the 12 experiments provides insight into the ability to cope with perceptual ambiguity.

Table 3.2 shows the average success rate (\overline{sr}) and standard deviation (sd) for all 12 combinations of neurocontroller and number of blind sensors (b). Figure 3.8 illustrates the average success rates in a bar chart.

All agents appear capable of performing categorical perception above the level of chance (i.e., $\overline{sr} > 0.50$), even when all sensory states are ambiguous to the agent, i.e., $b = 3$. This finding is most surprising for reactive (i.e., P-controlled and MLP-controlled) agents, since they cannot integrate sensory information over time as non-reactive (i.e., RNN-controlled) agents can.

For P-controlled agents, those with one blind sensor ($b = 1$) outperform agents without blind sensors ($b = 0$). It should be noted that for both cases, $b = 0$ and $b = 1$, the percentage of ambiguous sensory states is identical (i.e., 44.4%; see table 3.1). For larger numbers of blind sensors ($b = 2$ and $b = 3$) a decrease in performance of P-controlled agents is observed.

For MLP-controlled agents, a similar pattern of relative performance over the different values of b is observed. Still, MLP-controlled agents outperform the P-controlled agents for each value of b .

For RNN-controlled agents, a slightly different pattern of results is observed. RNN-controlled agents with $b = 1$ and $b = 2$ perform better than those with $b = 0$. For $b = 0$, the difference in performance between RNN-controlled and MLP-controlled agents is small. RNN-controlled agents outperform P-controlled agents whatever the value of b . The RNN-controlled agents outperform the MLP-controlled agents adequately, albeit with a considerable smaller margin.

Performance on the active categorical perception task is not linearly related to the perceptual ambiguity in the task. This is, for instance, shown by the differences in performance between agents with $b = 0$ and $b = 1$ (for all three neurocontrollers) while the number of ambiguous sensory states is equal for both values of b . This relation between performance and ambiguity can be attributed to two factors. First, it can be attributed to the change in sensor configuration as a result of altering the number of blind sensors. Second, it can be attributed to the difference in the

	$b=0$		$b=1$		$b=2$		$b=3$	
	\overline{sr}	sd	\overline{sr}	sd	\overline{sr}	sd	\overline{sr}	sd
P	0.7358	0.0011	0.7775	0.0013	0.7258	0.0008	0.6867	0.0002
MLP	0.8217	0.0018	0.8333	0.0010	0.7708	0.0020	0.7283	0.0016
RNN	0.8492	0.0031	0.9050	0.0037	0.9117	0.0037	0.8392	0.0030

Table 3.2: Average success rate (\overline{sr}) and standard deviation (sd) for all neurocontrollers and sensor configurations.

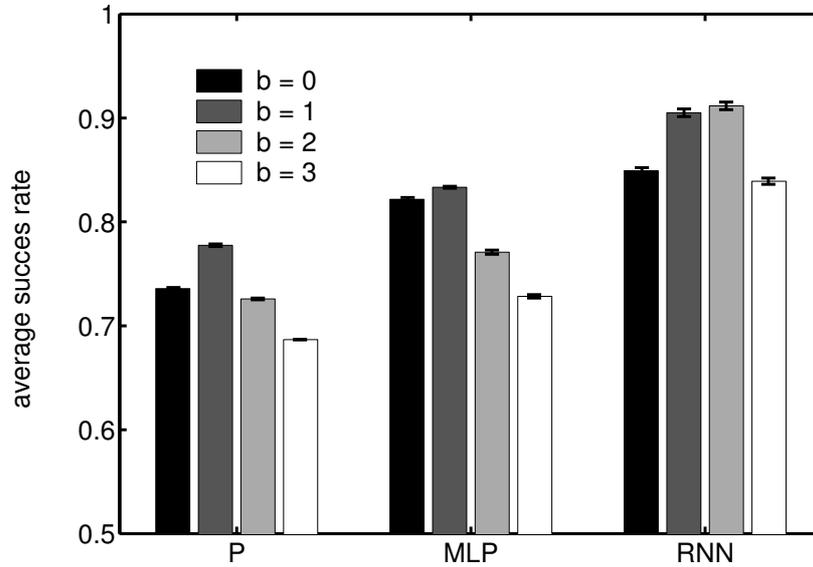


Figure 3.8: Average success rate of the best-performing agents on the active categorical perception task. Each bar represents the averaged performance of agents with a different combination of number of blind sensors (b) and neurocontroller (P, MLP, RNN).

spatial extent or scope of the sensor array. Presumably, both factors influence the performances in our study.

In summary, the results reported above answer the first question stated in section 3.1, i.e., “to what extent can reactive and non-reactive agents cope with perceptual ambiguity?”, as follows. In the context of categorical perception, both reactive and non-reactive agents can cope with perceptual ambiguity to the extent that all sensory states are ambiguous to the agent (see $b = 3$ in table 3.2). When all sensory states are ambiguous to the agent, the average success rate amounts to 0.6867 for P-controlled agents, 0.7283 for MLP-controlled agents, and 0.8392 for RNN-controlled agents. Yet, although both types of agent can cope with perceptual ambiguity, non-reactive agents (i.e., RNN-controlled agents) outperform reactive agents (i.e., P and MLP-controlled agents) at all levels of perceptual ambiguity.

3.5 Analyses of behavioural strategies

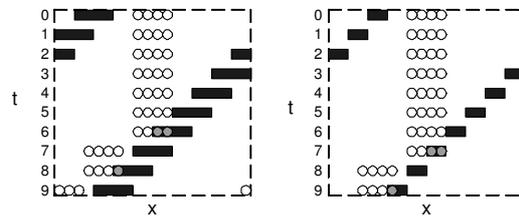
In section 3.4 we established to what extent reactive and non-reactive agents can cope with perceptual ambiguity. In this section we investigate the different behavioural strategies employed by reactive and non-reactive agents. In subsection 3.5.1, the behaviour and the sensory states of optimised agents are recorded and observed for both types of agent. In subsection 3.5.2, sensory state-transition diagrams are constructed for reactive agents only, for non-reactive agents this would lead to diagrams with infeasible complexity.

3.5.1 Observation of behaviour

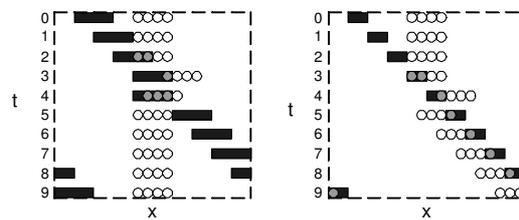
The behaviour and sensory states of optimised agents were recorded over all experiments. Figure 3.9 illustrates typical behaviour of optimised agents without blind sensors. Each subfigure shows one avoid (left panel) and one catch (right panel) sequence. For each of these sequences, the figure illustrates the relative horizontal position (x) between an object’s elements (crosses) and an agent’s sensors (circles) from $t = 0$ to $t = 9$.

Figure 3.9(a) shows the behaviour of a reactive P-controlled agent avoiding (left) and catching (right) an object. The examples show that the agent is capable of successful categorisation (i.e., avoiding large objects and catching small objects) by reacting only to its sensory states. Part of the behavioural strategy employed by the P-controlled agent is visible in figure 3.9(a). The agent remains passive until sensory state 0011 occurs. Subsequently, it exhibits behaviour that results in avoidance behaviour for large objects (left panel of figure 3.9(a)), or results in catching behaviour for small objects (right panel of figure 3.9(a)).

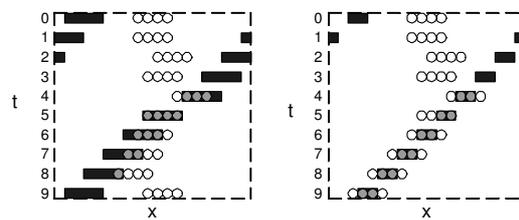
Figure 3.9(b) shows the reactive behaviour of an MLP-controlled agent avoiding (left) and catching (right) an object. In this case the reactive agent enters and remains in a ‘behavioural attractor’ (Thelen, 1995; Nolfi and Parisi, 1999; Nolfi and Marocco, 2001b) from $t = 5$ onward, in both the avoid and catch sequence.



(a) P-controlled reactive agent without blind sensors



(b) MLP-controlled reactive agent without blind sensors



(c) RNN-controlled non-reactive agent without blind sensors

Figure 3.9: Examples of avoid (left panels) and catch (right panels) behaviour during a single trial of optimised agents with three different types of neurocontrollers (P, MLP, and RNN) and no blind sensors (i.e., $b = 0$). The horizontal position (x) of the object (black rectangle) and the agent's functional sensors (circles) is depicted over time (t). White circles represent inactive sensors and gray circles represent sensors that are activated by the presence of an object.

A behavioural attractor is a single state or a repeating sequence of states, (i.e., a fixed point or n -cycle, respectively) to which an agent is attracted. A behavioural attractor arises because it yields evolutionary (reproductive) success associated with the sensory state(s).

Figures 3.9(a) and 3.9(b) illustrate how reactive agents cope with perceptual ambiguity despite their incapability to store sensory information. Although reactive agents are forced to react exclusively to current sensory states, their local actions do result in different sensory states and ‘physical’ states, i.e., a new sensory state and a changed relative position between agent and object. Ultimately, the sensorimotor loop results in the relative distance between agent and object that is required for successful behaviour.

Figure 3.9(c) shows the non-reactive behaviour of an RNN-controlled agent avoiding (left) and catching (right) an object. The non-reactive RNN-controlled agent displays different behaviour than both reactive agents. As expected, in this case, the action depends on previous sensory states. Similar sensory states (e.g., state 0000 at $t = 0$ to $t = 3$ in figure 3.9(c)) do not always yield the same action. Apparently, the RNN-controlled agent exploits its ability to retain sensory information (cf., Beer (1996)).

Figure 3.9(c) illustrates how the non-reactive agent incorporates behavioural strategies (e.g., such as remaining in sensory state 0110 (see the right panel of figure 3.9(c))). However, in this case, the categorisation problem is largely solved internally, and the agent’s behavioural strategy may change depending on the context (see, e.g., Bakker and van der Voort van der Kleij, 2000).

In conclusion, reactive agents have to exploit the interaction with the environment to cope with perceptual ambiguity, while non-reactive agents can cope with perceptual ambiguity internally. The latter agents solve the perceptual ambiguity by integrating sensory information over time. Hence, the second research question is only answered partially. It remains to be established how the reactive agents exploit their interaction with the environment, and how they arrive at adequate behavioural strategies as visualised in figures 3.9(a) and 3.9(b).

3.5.2 Sensory state-transition diagrams

Sensory state-action (SSA) mappings of optimised agents are extracted from the action and sensory states recorded during the experiments. SSA mappings show all possible sensory patterns and the corresponding action a , i.e., the length and direction of a step, taken by an agent. Examples of SSA mappings are shown in figures 3.10 and 3.12 for an optimised P-controlled and MLP-controlled agent, respectively (with $b = 0$). In these mappings, actions correspond to the value a (defined in subsection 3.3.2), which determines the size and direction of the step taken by an agent (i.e., the number of grid cells moved to the left (negative output) or right (positive output)). Sensory states are represented by ellipses with black or white circles, expressing the binary value of activation for each individual sensor (i.e., black = ‘1’ and white = ‘0’). From table 3.1 we know that only nine sensory states can occur (when $b = 0$).

We define sensory state-transition (SST) diagrams to analyse the behavioural strategies employed by reactive agents. SST diagrams are constructed on the basis of SSA mappings for each object class and direction of movement. The SST diagrams contain all trajectories of sensory states, instead of single instances such as given in figure 3.9. As argued below, SST diagrams make explicit how reactive agents make use of the interaction with their environment. As a result, they reveal how the environment acts as an external memory. In figures 3.11 and 3.13, the SST diagrams constructed on the basis of the SSA mappings in figures 3.10 and 3.12 are shown. Sensory states are represented by ellipses with black or white circles, as in the SSA mappings. The arrows in the diagrams represent transitions from one sensory state to another. Since the depicted agents are reactive, there can only be one successor state from each sensory state. However, this is not the case for the sensory state in which none of the sensors is active, i.e., the sensory state 0000. Although the agent behaves reactively, in the 0000 sensory state the object can be at many different positions relative to the agent. The arrow from sensory state 0000 to a target sensory state is labelled with a number indicating the maximal number of time steps before the transition to the target sensory state occurs, i.e., indicating the number of time steps the agent maximally remains in sensory state 0000. For instance, if the arrow from sensory state 0000 to a target sensory state is labelled ‘< 8’, then the transition from sensory state 0000 to the sensory state towards which the arrow points occurs within 8 time steps.

For small objects, only six of the nine sensory states of the SSA mappings depicted in figures 3.10 and 3.12 can occur. Hence, the diagrams in figures 3.11(a), 3.11(b), 3.13(a), and 3.13(b), contain only six sensory states. For large objects, eight of the sensory states can occur, which is why the corresponding diagrams (figures 3.11(c), 3.11(d), 3.13(c), and 3.13(d)) contain eight sensory states.

The SST diagrams of the optimised P-controlled agent, depicted in figure 3.11, reveal that the agent is optimised to avoid large objects, since all four diagrams show many arrows pointing towards sensory state 0000. The diagrams reveal behavioural attractors (n -cycles) in the P-controlled agent. Once the agent enters one of the three states 0000, 1000, or 1100 (for figures 3.11(a) and 3.11(c)), or 0000, 0001, or 0011 (for figures 3.11(b) and 3.11(d)), its future behaviour is restricted to these states. Note that the value of n in the n -cycles cannot be determined from the SST diagrams, since the order of the sensory states may vary within the cycle.

The SST diagrams of optimised MLP-controlled agents show a somewhat different behaviour. Besides n -cycles, their diagrams reveal the existence of fixed point attractors that lead to catch behaviour. The fixed point attractors correspond to sensory states 0001 and 1000 for small objects that fall to the right and small objects that fall to the left (figures 3.13(a) and 3.13(b)), respectively. The same attractor states also occur in the case of large objects that should be avoided (see figures 3.13(c) and 3.13(d)). However, in the large object cases, the agent only enters these states when the state is the initial state at $t = 0$ (see figure 3.13(c) and 3.13(d)).

The behavioural attractors in figure 3.13 clearly demonstrate how a reactive agent can make a categorical decision on the basis of information stored in the

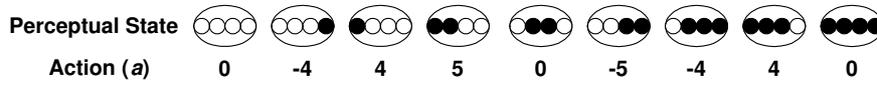
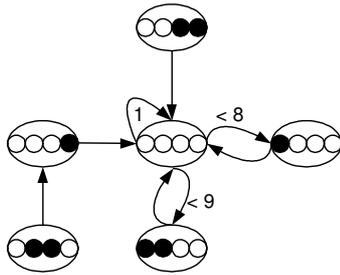
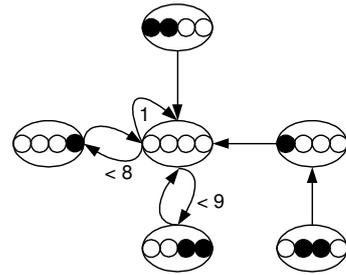


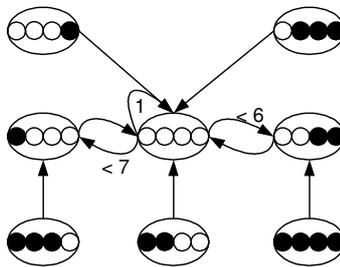
Figure 3.10: Mapping of sensory state to action (a) of an optimised P-controlled agent ($b = 0$, $\bar{s}\bar{r} = 0.7500$).



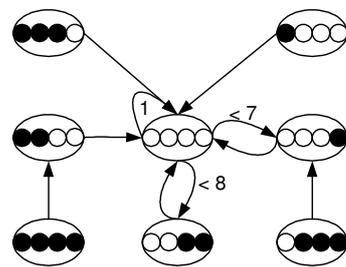
(a) Sensory state-transition diagram of a P-controlled agent facing a small object falling to the right.



(b) Sensory state-transition diagram of a P-controlled agent facing a small object falling to the left.



(c) Sensory state-transition diagram of a P-controlled agent facing a large object falling to the right.



(d) Sensory state-transition diagram of a P-controlled agent facing a large object falling to the left.

Figure 3.11: Sensory state-transition diagrams of an optimised P-controlled agent and a small object falling to the right (a) or to the left (b), and with a large object falling to the right (c) or to the left (d).

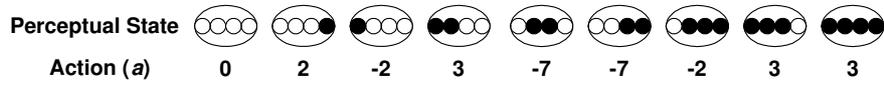
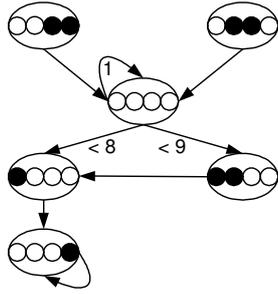
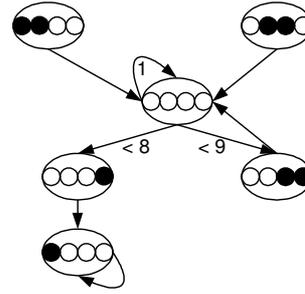


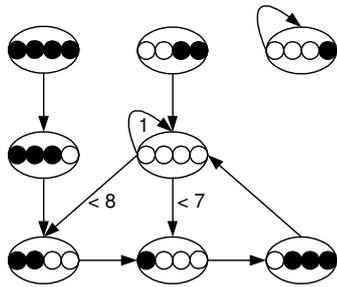
Figure 3.12: Mapping of sensory state to action (a) of an optimised MLP-controlled agent without blind sensors ($b = 0$, $\overline{s\overline{r}} = 0.8350$).



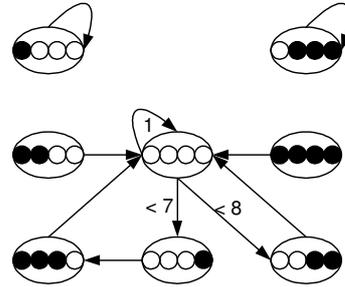
(a) Sensory state-transition diagram of an MLP-controlled agent facing a small object falling to the right.



(b) Sensory state-transition diagram of an MLP-controlled agent facing a small object falling to the left.



(c) Sensory state-transition diagram of an MLP-controlled agent facing a large object falling to the right.



(d) Sensory state-transition diagram of an MLP-controlled agent facing a large object falling to the left.

Figure 3.13: Sensory state-transition diagram of an optimised MLP-controlled agent and a small object falling to the right (a) or to the left (b), and with a large object falling to the right (c) or to the left (d).

environmental dynamics. For instance, a transition from state 0001 can result in a fixed point attractor 1000 to catch a small object falling to the right (figure 3.13(a)), or result in a state that is part of an n -cycle to avoid large objects falling to the left (figure 3.13(d)). The reactive agent relies on the consequence of its action, which is determined by the environmental dynamics, to act appropriately.

Our analysis of the SST diagrams reveals that reactive agents compensate their lack of internal memory by using the environment (i.e., the falling object) as an external memory. The behaviour of an agent corresponds to a route through the SST diagram. Each sensory state determines the subsequent sensory state in the route through the SST diagram, to arrive finally at the appropriate terminal sensory state(s). Hence, the environment (as represented by the sensory states) is used as an external memory. The use of external memory is most strikingly revealed by the attractor states that couple the dynamics of the agent to the dynamics of the object. Both P-controlled agents and MLP-controlled agents employ external memory as expressed by the behavioural attractors in the SST diagrams. The ability of MLP-controlled agents to perform non-linear SSA mappings enriches their ability to use the environment as an external memory by combining fixed-point attractors with n -cycle attractors.

By introducing SST diagrams we were able to reveal the behavioural strategies employed by reactive agents. Using the diagrams, our analysis showed that reactive agents use the environment as an external memory to compensate for their lack of internal memory. Our results provide a low-level understanding of similar observations on the use of external memory in natural agents (O'Regan, 1992; O'Regan and Noë, 2001b). In fact, what these SST diagrams show are O'Regan and Noë's (2001b) sensorimotor contingencies, which they define as "the structure of the rules governing the sensory changes produced by various motor actions" (*ibid*, p.941).

3.6 Discussion

Below, we summarise our results in subsection 3.6.1. Thereafter, we discuss the results in subsection 3.6.2.

3.6.1 Summary of results

Our results clearly contradict Beer's ((1996), p.424) statement that reactive agents "cannot organise their behavior according to sensory stimuli that are no longer present", since they show that reactive agents can cope with perceptual ambiguity in the context of active categorical perception. Both reactive controllers (P and MLP) perform categorisation above the level of chance, even when all sensory states are ambiguous to the agent, i.e., when $b = 3$. By exploiting the interaction with their environment, i.e., the consequences of their actions, reactive agents can organise their behaviour according to past sensory stimuli, and, as a result, cope with perceptual ambiguity. The improved performance of MLP-controlled agents as compared to P-

controlled agents indicates the contribution of a non-linear SSA mapping in reactive agents coping with perceptual ambiguity.

3.6.2 Discussion of results

Because categorisation is generally considered to be fundamental to cognition (Harnad, 1987; Pfeifer and Scheier, 1997; Tijsseling, 1998; Beer, 2003b; Harnad, 2003), generalising Nolfi’s (2002b) finding to a categorical perception task is relevant for research in adaptive behaviour. Our results indicate that, for a specific low-level task, this fundamental cognitive ability can be exhibited by reactive agents despite any perceptual ambiguity. This means that the ability to perform cognitive tasks, requiring integration of sensory information over time, depends only partially on the internal dynamics of an agent’s neurocontroller. In fact, the performance of reactive agents shows that representations can be exclusively external. As argued in section 2.3, they should be called representations, rather than presentations, because they are picked up, analysed, and processed by a perceptual system.

It is important to note that evolutionary selection occurs on the basis of complete sequences of sensorimotor behaviour, rather than on single perception-action steps (de Croon *et al.*, 2005b). Consequently, our results, and those of Nolfi (2002b), can be explained by the fact that perceptual ambiguity is defined locally in time, while behavioural success (catching and avoiding) is defined globally. A similar point is made by Izquierdo-Torres and Di Paolo (2005, pp.260-261), who claim that “a reactive controller in an embodied system doesn’t imply reactive behaviour: there is a difference between the local, instantaneous state definition of reactivity, and the behavioural definition.”

It should also be noted that the constancy in the environmental dynamics is exploited by the evolutionary algorithm to achieve an effective input-output mapping. Therefore, it can be argued that our results are due to the specificity of the environmental dynamics of ACP. To assess whether our results generalise over more complex environmental dynamics, we conducted two additional experiments in which we varied the horizontal and vertical velocities of the falling objects, respectively. These experiments and their results are reported in detail in appendix A. The results for both additional experiments show an overall decrease in performance as compared to the original results (for $b = 0$), which is due to the increased complexity of the task. More importantly, the relative pattern of performances remains largely unaffected by the variation of the velocities of falling objects. This is taken as an indication that our results generalise across moderate variations in the environmental dynamics. Whether our results generalise over larger variations remains to be established.

3.7 Chapter conclusions

From the study reported in this chapter we may draw four conclusions. First, we may conclude that reactive agents can cope with perceptual ambiguity in the context of active categorical perception. Second, we may conclude that Beer's (1996) statement is debatable, since our results show that reactive agents *can* organise their behaviour according to sensory stimuli that are no longer present using the environment as an external memory, which falls in line with Nolfi's (2002b) finding. Third, reactive agents incorporating a non-linear sensorimotor mapping are better able to deal with perceptual ambiguity in an active categorical perception task than those incorporating a linear mapping. Finally, we may conclude that SST diagrams provide insight into the behavioural strategies employed by reactive agents to deal with perceptual ambiguity, and their use of the environment as an external memory.

In summary, the findings above demonstrate that, as was argued in chapter 2, representations in situated systems can be internal and/or external. The performance of reactive agents showed that representations can be exclusively external. The SST-diagrams of the behaviour of these agents made explicit how such external representations depend on the perceptual system of the agent.

The behavioural strategies of reactive agents showed that representation can be exclusively external. In contrast, the analysis of behavioural strategies of non-reactive agents showed that non-reactive agents use internal representation. These findings demonstrate that, as was argued in chapter 2, representation in situated systems can be internal and/or external. The SST diagrams of the behaviour of reactive agents made explicit how external representation depends on the perceptual system of the agent, and that external representation is governed by the environmental dynamics of the model (i.e., ACP). The operationalisation of situated representation should allow internal representation and external representation. By doing so, SST diagrams can be adopted to study the externally represented information in reactive systems. Chapters 4 and 5 will focus on internal representation and external representation, respectively.

Chapter 4

Internal representation

In chapter 3 we demonstrated that representations in situated systems can be both internal and external. In this chapter¹, the aim is to study the nature of internal representation. As stated in chapter 1, according to the *physical symbol-system hypothesis* (Newell and Simon, 1985) of the classic AI, internal representation involves symbols (Newell and Simon, 1972). Furthermore, cognition is considered to be the manipulation of these symbols (Turing, 1950; Newell and Simon, 1972). While chapter 1 already mentioned the problems of symbolic models of cognition, and chapter 3 suggested that cognition does not depend on symbolic representation in several agent domains, it remains to be established whether situated systems can perform tasks that require symbol manipulation. Or as Matarić (1998, p.85) puts it: “How well will behavior-based systems scale-up to increasingly more cognitive problems, such as those involving symbolic reasoning[?]”. In this chapter, we will investigate whether a model that satisfies the conditions discussed in subsection 1.4.1 (i.e., a model that addresses situatedness, embodiment, and cognition, and is characterised by parsimony and transparency) can perform a task that requires symbol manipulation. We will provide some background on this investigation and previous investigations in section 4.1. At the end of that section we will give an overview of the remainder of this chapter.

4.1 Symbol manipulation in situated agents

Phaf and Wolters (1997, pp.294-295) state that (i) “[t]he simplest way for an organism to maintain the representation of an object is to keep the object present as sensory input” and (ii) internalised interaction with the environment constitutes

¹The experiments and results reported here were presented at the 5th International Workshop on Epigenetic Robotics in Nara, Japan (see van Dartel and Postma (2005)) and at the 17th Belgium-Netherlands Conference on Artificial Intelligence in Brussels, Belgium (see van Dartel and Postma (in press)). The author would like to thank his co-author and the publisher of the proceedings for their kind permission to reuse relevant parts of the article in this thesis.

symbolic representation. Hesslow (2002) suggests a mechanism to internalise interaction with the environment in his *simulation hypothesis*, which holds that conscious thought is based on the ability to simulate perception and behaviour internally (the ability to predict changes in sensory stimulation by internal sensorimotor models is often referred to as ‘forward modelling’, see, e.g., Grush (2004) for a discussion and Hoffmann and Möller (2004) for an implementation). Hesslow supports his simulation hypothesis with the observed evidence that (i) the activation patterns in sensory areas of natural brains are similar for imagined and actual perception (see, e.g., Kosslyn, Ganis, and Thompson (2001)), and (ii) motor areas in natural brains exhibit similar patterns of activation during imagined and actual behaviour (see, e.g., Jeannerod (1994)).

Ziemke, Jirenhed, and Hesslow (2005) show that situated agents can perform collision-free corridor following behaviour on the basis of a neurocontroller that allows internal simulation of perception and behaviour. This indicates that situated agents are able to exploit the ability to simulate perception and behaviour internally, which, as Phaf and Wolters (1997) state, constitutes symbolic representation. The preliminary experiments reported in appendix B show that situated agents with the ability to simulate perception and behaviour internally outperform agents that do not have this ability on the categorisation task described in chapter 3. This suggests that situated agents with the ability to simulate perception and behaviour internally may outperform agents that do not have this ability.

Manipulation of symbolic representation (i.e., symbol manipulation) is often associated with planning (Newell and Simon, 1972), a skill that is generally regarded as high-level cognition (Cooper, 2002). It is assumed that, to plan ahead in time, one needs to represent the current state of the task in symbols and extrapolate future states by manipulation of these symbols. The Tower of London task is a typical planning task (Shallice, 1982) that is a standard neuropsychological test to assess frontal lobe damage (Kolb and Wishaw, 1983), which impairs planning performance (see, e.g., Baddeley (1986)). Hesslow (2002, p.245) states that “simulating chains of behaviour is a plausible interpretation of the problem-solving process in tasks like the Tower of London”. We claim that the Tower of London task requires the manipulation of representations by means of internalised interaction with the environment. Following the physical symbol-system hypothesis, we refer to such manipulation of representation as symbol manipulation.

Since situated agents are able to benefit from the ability to simulate perception and behaviour internally (as was stated above), we expect situated agents to be able to perform the Tower of London task. To test this, we construct a situated Tower of London (sToL) model. Subsequently, we formulate the following main research question: *Can our situated model perform the Tower of London task?*

If so, the agents will provide a unique opportunity to investigate the problem statement formulated in chapter 1, i.e., “what is the nature of representation in situated systems?”, by analysing the symbol-manipulation process underlying the agents’ performance on the Tower of London task. To do so, we formulate two subsequent research questions: (1) which mechanism allows a situated agent to plan

ahead in time? and (2) how are the current and future states of the Tower of London represented in an agent?

The Tower of London task will be discussed in more detail in section 4.2. Section 4.3 describes the sToL model. The experiment conducted with the sToL model is described in section 4.4. The results of the experiment are reported in section 4.5. Two analyses of the behaviour of the agent in the sToL model will be conducted in section 4.6. In section 4.7, the performance of the agent will be compared to the performance of human subjects, the design of the sToL model will be compared with that of two other models (viz. (a) the model by Dehaene and Changeux (1997), and (b) the model by Polk *et al.* (2002)), and the sToL model will be compared with other models that incorporate internal simulation. In section 4.8 the symbol manipulation of the agent in the sToL model is discussed. Finally, chapter conclusions are given in section 4.9.

4.2 The Tower of London task

The Tower of London (ToL) task is often employed by psychologists to test a subject's ability to plan ahead in time (Shallice, 1982; Morris *et al.*, 1993; Dagher *et al.*, 1999; Lazeron *et al.*, 2000; Sikora *et al.*, 2002). In the ToL task, subjects are asked to change a given starting configuration of three coloured balls on three pegs to a certain goal configuration *in the least possible number of moves*. Figure 4.1 depicts the physical model that is used in the ToL task. The constraints of movement in the ToL task are the same as the constraints of movement in the Tower of Hanoi task

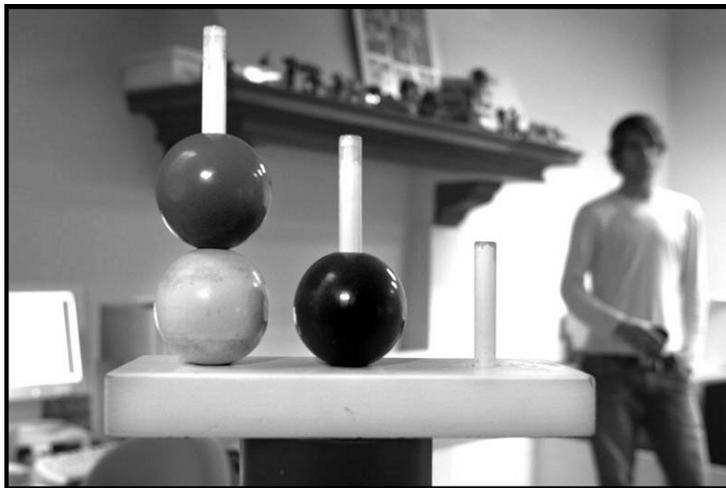


Figure 4.1: A Tower of London (ToL). Photo: Philip Driessen Fotografie.

(Lucas, 1882-3) (i.e., the *towers* of Hanoi task). However, there is no constraint on the order in which balls can be stacked, as is the case in the Tower of Hanoi task.

Figure 4.2 shows four possible states of the ToL labelled 0 to 3. Each state consists of a configuration of the three balls on the three pegs. Henceforth, we will use the terms *state* and *configuration* interchangeably. In figure 4.2, the balls,

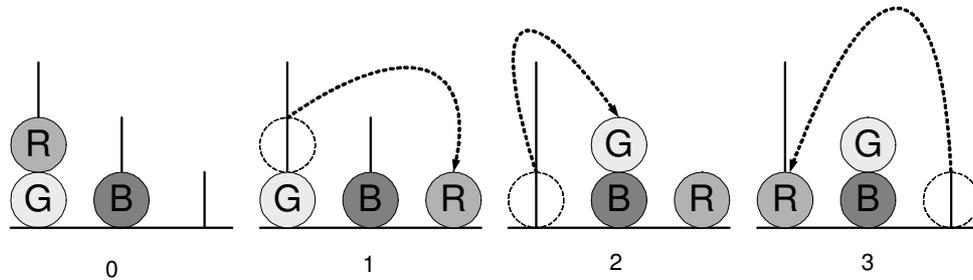


Figure 4.2: Illustration of the actions required to reach goal state 3 from the starting state 0 through intermediate states 1 and 2 (see the text for further details).

represented by shaded circles, are labelled according to their colour (R, G, and B, for red, green, and blue, respectively). Dashed circles represent previous ball positions and the arrows represent movements of balls. The figure illustrates how a subject can reach goal state 3 from starting state 0 by visiting two intermediate states (1 and 2) without violating the constraints of the ToL task.

In the ToL task, all possible goal configurations can be reached from any starting configuration. A typical ToL task consists of multiple ToL problems. A ToL problem is defined as finding the path with the least possible number of moves from a given starting configuration to a certain goal configuration. Successful completion of a ToL task and the number of successive moves needed to do so depend on a subject's ability to plan ahead in time. The ToL task is considered a typical high-level planning task, because for each problem in the task "successful completion requires the participant to 'look ahead' and solve the problem cognitively before actually moving the balls" (Bull, Espy, and Senn, 2004, p.743). To solve a ToL problem successfully, a subject is believed to decompose the problem into subproblems (Shallice, 1982; Cooper and Waldau, submitted).

The complexity of a ToL problem can vary considerably, depending on the starting configuration in relation to the goal configuration. An experimental test of twelve ToL problems defined by Shallice (1982) has become a standard neuropsychological test to assess frontal lobe damage (Kolb and Whishaw, 1983). In the test, the starting configuration is the same for all twelve problems, whereas the goal configuration varies. The starting configuration and goal configurations of Shallice's test are illustrated in table 4.1. There, states are represented by configurations of the three balls (R, G, and B) on the three pegs. A peg without any ball is represented by '-'. The

last column of the table lists the minimum number of moves required to solve the problems. Note that subsequent problems in the set have an equal or higher level of complexity, as follows from the minimal number of moves that a subject has to make in order to reach the goal state of that problem.

4.3 The situated Tower of London model

In order to study whether situated agents are capable of performing a task that requires symbol manipulation, we define the situated Tower of London (sToL) model. We describe the model in terms of the test (subsection 4.3.1), the agent's sensors and actuators (subsection 4.3.2), the agent's neurocontroller (subsection 4.3.3), the evolutionary algorithm (subsection 4.3.4), and the performance measures (subsection 4.3.5).

4.3.1 The test

In the sToL model, a test consisting of twelve ToL problems is used to determine whether our situated model can perform the ToL task. The twelve ToL problems in the test are the same as the ones used by Shallice (1982) (shown in table 4.1). The agent in the sToL model is trained to perform the twelve ToL problems in the test in the least possible number of moves. For practical reasons, the interaction between the agent and the ToL is simplified in terms of sensor and actuator definitions.

4.3.2 The agent's sensors and actuators

The agent in the sToL model is able to perceive the current configuration of the ToL and can act in response by moving a ball from one peg to another. The model is situated because the agent can observe the consequences of its own actions, and use these to learn how to reach the goal state of a ToL problem. The agent in the sToL model consists of a neurocontroller that receives information about the configuration of the three balls on the three pegs through its sensors. The agent's motor system can change the position of a ball according to the output of its neurocontroller and the constraints of the ToL task. The simplified interaction between the agent and the ToL is realised by encoding the configurations of the ToL and the actions taken by the agent in a straightforward way in its sensors and actuators. The encoding of the configuration of the ToL and of the actions produced by the agent is kept as simple as possible to adhere to the parsimony condition introduced in subsection 1.4.1 and to minimise the dimensionality of inputs and outputs.

The sensors

The sensor array of the agent contains eighteen sensors. Each sensor S_i ($i \in \{1, 2, \dots, 18\}$) senses the presence of a ball at a certain peg and position. To illustrate the encoding of the current configuration of the ToL in sensory activation, we

Starting state for all problems:

R
G B -

Problem	Goal state	Minimum # of moves
1	G - B R	2
2	G R B	2
3	B G R -	3
4	B G R -	3
5	G R B -	4
6	G B R -	4
7	B R - G	4
8	B - R G	4
9	B G R - -	5
10	G B R - -	5
11	G B R -	5
12	B R G	5

Table 4.1: Starting state and goal states of the twelve ToL problems employed by Shallice (1982; reproduced from Berg and Byrd (2002)), and the minimum number of moves in which each problem of the test can be solved.

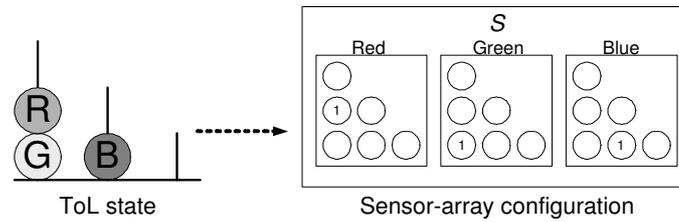


Figure 4.3: Example of encoding a ToL state (the current configuration of the ToL) in activation of the sensor array.

arranged the sensor array into three parts of six sensors. Each part corresponds to one of the colours (R, G, B) of the balls. The six sensors of each part correspond to the possible positions of balls on the pegs. Figure 4.3 illustrates the encoding of a particular current configuration of the ToL, i.e., the ToL state, into a sensory activation. In the figure, the ToL state is shown on the left, and the sensor array is shown on the right. The red ball is positioned on the left peg in the second position (starting from the bottom). In the sensor array, this is translated into an activity of one unit (represented by a ‘1’) of the corresponding sensor in the red part of the sensor array. Similarly, the positions of the green and blue balls are encoded as sensory activation in the green and blue parts of the sensor array, respectively. All other sensors are assigned zero activation (represented by empty circles in figure 4.3).

In addition to containing the ToL state, the sensor array may contain two additional states: the goal state and the “expected” state. The goal state is the objective of the particular ToL problem faced by the agent (see table 4.1). The expected state is generated by the internal simulation mechanism (to be described below). The goal state and the expected state are encoded in the same manner as the ToL state. The encodings of the ToL state, the goal state, and the expected state are added to yield a superimposed sensory activation pattern. Hence, an individual sensor S_i can have activation values ranging from zero units (no ball present on colour-position combination i in all three states) to three units (a ball is present on colour-position combination i in all three states). Figures 4.4 and 4.5 provide two additional illustrations of the encoding of two and three states into the sensor array, respectively. In figure 4.4, a ToL state (the current configuration of the ToL) and a goal state (the objective of the problem) are superimposed yielding the sensory activation pattern shown on the right. In figure 4.5, a ToL state (the current configuration of the ToL), a goal state (the objective of the problem), and an expected state (the state generated by the internal simulation mechanism to be described below) are superimposed resulting in the activation pattern shown in the sensor-array configuration on the right. The division of the sensor array into separate parts for each colour is done merely for our convenience; the agent treats all sensors on an equal basis.

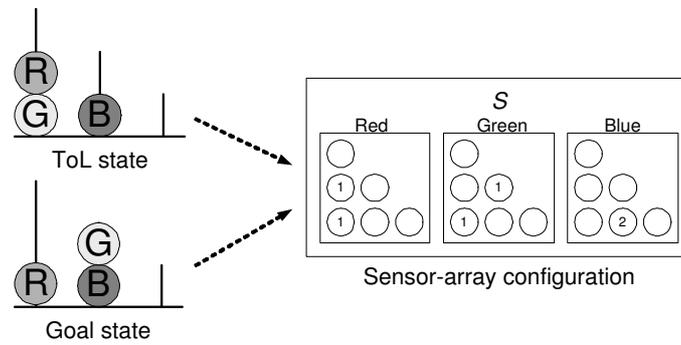


Figure 4.4: Example of encoding a ToL state and a goal state in activation of the sensor array.

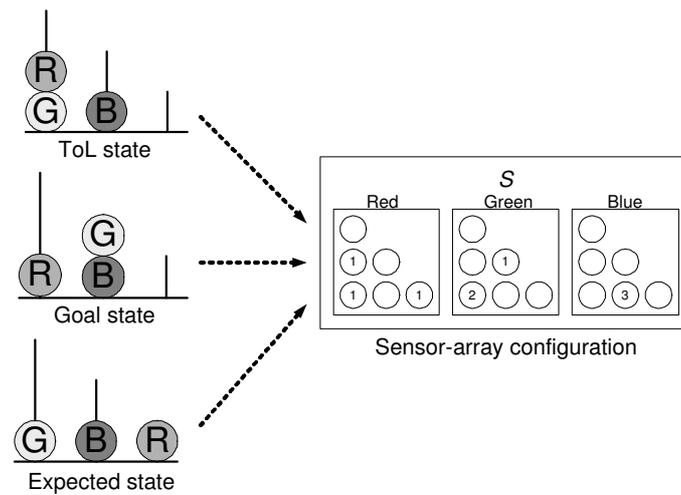


Figure 4.5: Example of encoding a ToL state, a goal state, and an expected state in activation of the sensor array.

Although more compact encodings are possible, we opted for the combined encoding of position and colour because it is commonly applied in neural models (see, e.g., Hinton, McClelland, and Rumelhart (1986)) and has been observed in natural brains (Schnapf, Kraft, and Baylor (1987)).

The actuators

Whereas the sensor array of the agent translates states into input for the neurocontroller, the actuator array of the agent translates the output of the neurocontroller into an action. The actuator array of the agent consists of six actuators. The first three actuators encode the colour of the ball to be moved (red, green, blue). The colour associated with the actuator with the highest level of activation is the colour of the ball to be moved. The last three actuators encode the peg to which the ball should be moved (left, middle, right). The peg associated with the actuator with the highest level of activation is the peg to which the ball should be moved. Figure 4.6 provides an illustration of encoding the colour of the ball to be moved and the peg to which it should be moved in the actuator-array configuration. In the figure,

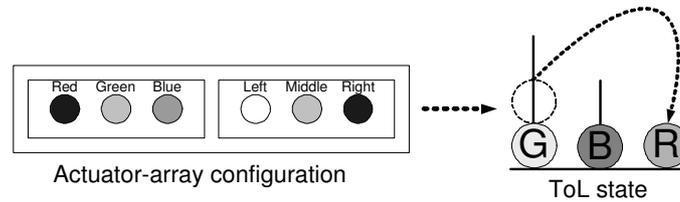


Figure 4.6: Example of encoding the colour of the ball to be moved, and the peg to which it should be moved in the actuator array.

the activation of the actuator array results in moving the red ball from the left peg to the right peg (the level of activation of an actuator is proportional to the level of shading). The agent performs the action indicated by the activation of the actuator array only if it meets the constraints of the ToL task. Otherwise, the action is not executed. As was the case for the sensors of the agent, a more compact encoding is possible for the actuators as well. For instance, in a preliminary version of the sToL model, the actions of the agent were encoded by two actuators; one for the colour of the ball to be moved and one for the peg to which the ball should be moved. By stimulating the actuators within certain predefined levels of activation, the agent executed a certain action. Although the type of encoding in the preliminary version of the sToL model better satisfies the parsimony condition stated in subsection 1.4.1, it was shown to lead to poor performances in comparison with the type of encoding adopted in the current version of the sToL model. As in the actuator-array configuration of the sToL model, in the human brain, position and colour are processed in separate areas as well (Palmer, 1999).

4.3.3 The agent's neurocontroller

In the sToL model, the agent can be controlled by either (i) the standard neurocontroller, or (ii) the neurocontroller with an internal simulation mechanism. Both neurocontrollers of the agent in the sToL model are non-reactive, because we expect, following the findings reported in chapter 3, that non-reactive agents are better capable of dealing with symbol-manipulation tasks than reactive agents. Below, both neurocontrollers are described. For reasons of readability we will henceforth call the standard neurocontroller by the name neurocontroller A and the neurocontroller with an internal simulation mechanism by the name neurocontroller B.

Neurocontroller A

Neurocontroller A of the agent in the sToL model consists of a simple recurrent neural network (RNN), i.e., an Elman network (Elman, 1990). The architecture of the RNN is similar to the RNN described in subsection 3.3.2. The eighteen input nodes of the RNN sample the activation of the sensor array described in subsection 4.3.2. The input is mapped onto h hidden nodes and an equal number of context nodes. The six output nodes of the RNN encode the actions in the actuator array described in subsection 4.3.2.

Neurocontroller B

In section 4.1, we claimed that the ToL task requires symbol manipulation on the basis of internalised interaction with the environment. Therefore, we equip the agent in the sToL model with a neurocontroller that enables internalised interaction with the environment. This neurocontroller, named neurocontroller B, has the same architecture as neurocontroller A, but is augmented with an internal simulation mechanism. The architecture of neurocontroller B is roughly the same as the architecture proposed by Jirnhed, Hesslow, and Ziemke (2001) and Ziemke *et al.* (2005).

The mechanism that enables the internal simulation consists of an additional output layer and a feedback connection from this additional output layer to the input layer of the neurocontroller. Figure 4.7 illustrates the architecture of the RNN with the internal simulation mechanism. The additional output layer employs the same encoding as the plain output layer. This means that the additional output layer also has six actuators of which the first three actuators encode the colour of the ball to be moved (red, green, blue), and the last three actuators encode the peg to which the ball should be moved (left, middle, right) (cf., figure 4.6). However, the output of the additional output layer is not used to reconfigure the ToL, but is internally processed to generate an expected state of the ToL. This expected state is generated in three steps. First, the additional output is interpreted as an action, determining which ball will be moved to which peg. Second, when the action meets the constraints of the ToL task, the action is applied to the expected state. If the action does not meet the constraints of the ToL task, no action is applied. Third, the

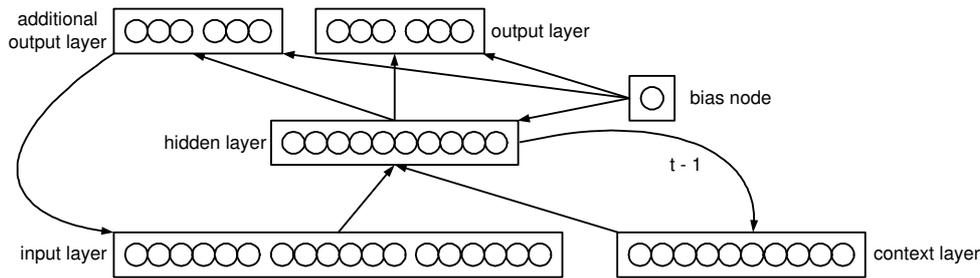


Figure 4.7: Architecture of neurocontroller B (i.e., the neurocontroller with an internal simulation mechanism).

updated expected state is superimposed on the activation caused by the ToL state and the goal state in the sensor array, to form a single sensory state, as illustrated in figure 4.5. This process continues until the plain output of the neurocontroller provides an action that satisfies the constraints of the ToL task.

4.3.4 The evolutionary algorithm

The behaviour of the agent in the sToL model depends on the weight values of its constituent neurocontroller. In the sToL model, we employ an evolutionary algorithm to optimise the performance of the agent on the ToL task. The evolutionary algorithm determines the weight values for all connections in the agent's neurocontroller.

The evolutionary algorithm used to optimise the agent's neurocontroller in the sToL model is similar to the one described in subsection 3.3.5. In the sToL model, the agent's neurocontroller is also represented in a genome by placing its weight values at arbitrary positions in a vector. All individuals within each generation consist of such a genome. In the first generation, all individuals are initialised by attributing a random value in the range $[-1,1]$ to each element in the genome. After initialisation, the fitness of each individual in the generation is determined. The fitness of an individual is determined by testing the success of the neurocontroller encoded by that individual on the set of ToL problems of table 4.1. The fitness F of an individual is calculated as:

$$F = ((C + S) \times 1000) - M \quad (4.1)$$

with C the number of balls positioned correctly upon termination of a problem, S the number of ToL problems that are solved, and M the number of moves that are made. In equation 4.1, C and S reward for successful behaviour, while M punishes for the number of moves made.

New generations are created in five steps. In the first step, all individuals from the last generation are copied to the new generation.

The second step is that, in the new generation, those individuals ranked 50 to 80 in the last generation are replaced by copies of the individuals ranked one to ten. For each of these 30 copies, chance determines which one of the ten best-ranked individuals of the last generation is selected to be copied to the new generation, with a chance of .30 for the best-ranked individual, .20 for the second-best-ranked individual, .10 for the individual ranked third, and .05 for the individuals ranked fourth to tenth.

The third step is that 200 mutations are performed over the new generation. A mutation consists of adding a random value to three randomly-selected elements in the genomes of three randomly-selected individuals for each of the connective layers that are encoded in the genomes. By repeating the mutations for each connective layer encoded in the genomes, the number of mutations is proportional to the size of the genomes. Hence, the number of mutations is proportionally larger when genomes that encode neurocontroller B are evolved than when genomes that encode neurocontroller A are evolved, because neurocontroller B has one connective layer more (the additional output layer) than neurocontroller A. The added random value is in the range $[-0.25, 0.25]$ for 66.67% of the mutations and in the range $[-1, 1]$ for the rest of the mutations.

The fourth step is that crossover is conducted on five randomly selected individuals. Crossover is conducted by replacing those elements that encode for the first half of the weights of each connective layer that is encoded in the genome with those same genome elements of another randomly selected individual. As in the evolutionary algorithm described in subsection 3.3.5, mutation and crossover are conducted over the complete new generation except for the best-ranked individual of the last generation, introducing elitism to the algorithm (Goldberg, 1986).

In the fifth step, the individuals ranked 80 to 100 in the last generation are replaced by individuals with randomly initialised genomes in the new generation. After a new generation has been created, the fitness of each individual is determined, and the process is repeated.

During the experiment described in section 4.4, the following settings for the evolutionary algorithm are maintained; the number of generations is set to 50,000 and each generation consists of 100 individuals.

By selecting the best-ranked individuals in each generation for reproduction, the evolutionary algorithm optimises the neurocontrollers encoded by individuals on the twelve ToL problems defined by Shallice (1982) shown in table 4.1. After the evolution, the neurocontroller encoded by the fittest individual of the last generation is selected to be the neurocontroller of the agent in the sToL model. The success of the evolutionary optimised agent is expressed in the performance measures described below.

4.3.5 The performance measures

Equation 4.1 indicates that the fitness of each individual in the evolution depends on both the number of solved problems (S in the right-hand side term) and the

number of moves used to solve those problems (M , also in the right-hand side term). Therefore, the performance of the agent in the sToL model can be expressed in two different measures, expressing either (1) the number of problems that an agent solved (S ; cf., Shallice (1982)), and (2) the number of moves it used to do so (M ; cf., Phillips *et al.* (1999)).

In the experiment that follows, the performance of the agent is expressed in terms of the following two new measures, that are defined to indicate S and M as ratios of S_{max} and M_{min} , respectively, in the range $[0, 1]$. The first measure is defined by S/S_{max} , the number of solved problems divided by the maximum number of problems (which equals 12 in the test). The second measure is defined by M_{min}/M , the minimal number of moves in which the agent could have solved the problems that it solved (48 when all problems in the test are solved, and fewer otherwise) divided by the number of moves it used to solve those problems. Each individual evolution in the experiment (described below) is replicated five times to obtain a good estimate of the optimised agent’s mean performance expressed in S/S_{max} and M_{min}/M . We represent the average values of S/S_{max} and M_{min}/M by \bar{S} and \bar{M} , respectively. For the experiment reported in the next section we indicate the performances of the agent on the ToL task in terms of the measures \bar{S} and \bar{M} .

4.4 Experiment

In the experiment, the agent in the sToL model is optimised to perform the test described in subsection 4.3.1 with either neurocontroller A or neurocontroller B (see subsection 4.3.3). For each neurocontroller, we indicate the performance of the optimised agent by the two new performance measures (\bar{S} and \bar{M}) described in subsection 4.3.5. The experiment consists of two parts.

First, we optimise the combined number of hidden nodes and context nodes, indicated by h , for both neurocontrollers (A and B). Starting from $h = 4$, the value for h is repeatedly increased by 2, until \bar{S} reaches a predefined criterium p_c . During the experiment we set $p_c \geq 0.95$.

Second, we determine the performance of the agent expressed in \bar{M} for each combined number of hidden nodes and context nodes, indicated by h , for which the agent is optimised. In the next section, the performances of the agent for both neurocontrollers (A and B) and all tested values for h are reported.

4.5 Results

For both neurocontrollers (A and B) the performances of the agent expressed in \bar{S} are shown in table 4.2 and figure 4.8. The table and figure show that, for both neurocontrollers (A and B), the agent reaches the performance criterium $p_c \geq 0.95$ when $h = 14$. As the error bars (showing the standard deviations) in figure 4.8 indicate, the performances expressed in \bar{S} of the agent with neurocontrollers A and

h	Neurocontroller A		Neurocontroller B	
	\bar{S}	sd	\bar{S}	sd
4	0.5333	0.1394	0.6000	0.0372
6	0.6833	0.1086	0.7500	0.1318
8	0.8833	0.1728	0.8167	0.1900
10	0.7833	0.1263	0.9167	0.1020
12	0.8833	0.1118	0.7667	0.0697
14	0.9500	0.0745	0.9667	0.0456

Table 4.2: The average ratio of solved problems \bar{S} and the corresponding standard deviations for neurocontroller A and neurocontroller B.

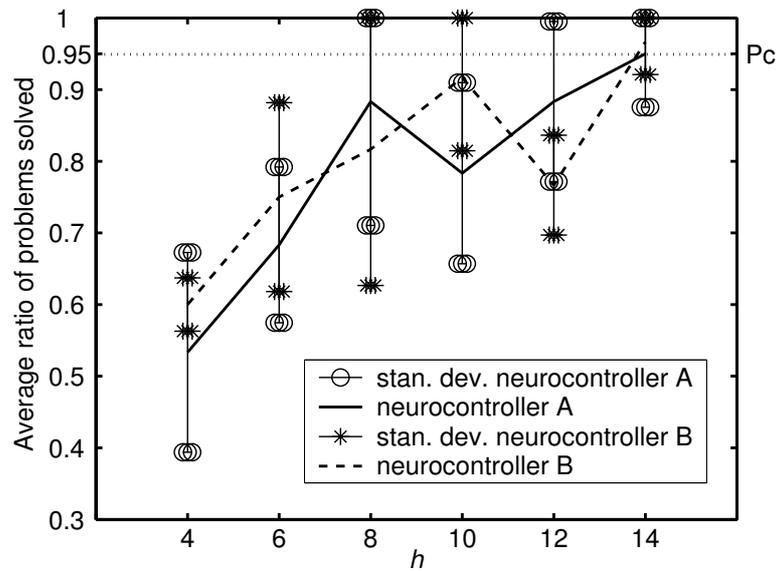


Figure 4.8: The average ratio of solved problems \bar{S} as a function of the number of hidden nodes and context nodes (h) for neurocontroller A (solid line) and neurocontroller B (dashed line).

B do not differ much for any value of h . An ANOVA on the mean performances obtained with the two neurocontrollers indicates that for any tested value of h , the mean performances expressed in \bar{S} do not differ significantly ($p > 0.05$).

The second performance measure \bar{M} is the average ratio of moves made by the optimised agent. For both neurocontrollers the performances of the agent expressed in \bar{M} are shown in table 4.3 and figure 4.9. The table and figure indicate that the agent performs the ToL task in the sToL model better when it is controlled by neurocontroller B than when it is controlled by neurocontroller A for any value of h . An ANOVA on the mean performances obtained with the two neurocontrollers indicates that the mean performances expressed in \bar{M} differ significantly ($p < 0.005$) for $h = 14$, but do not differ significantly ($p > 0.05$) for $h < 14$.

Taken together, the results show that the agent with neurocontroller B uses significantly fewer moves than with neurocontroller A when the ratio of problems solved \bar{S} is larger than or equal to 0.95. Apparently, when almost all twelve problems of the test are solved by the agent, it can do so more efficiently when it is controlled by neurocontroller B than when it is controlled by neurocontroller A.

4.6 Analyses of behaviour

The results reported above show that our situated model can perform the ToL task. The situated agent in the sToL model can perform the ToL task, whether it has the ability to simulate perception and behaviour internally or not. More importantly, the situated agent performs better when it is equipped with a neurocontroller with an internal simulation mechanism (neurocontroller B) than when it is equipped with a neurocontroller without such a mechanism (neurocontroller A). This suggests that internal simulation facilitates the planning performance on the ToL task. Since we are interested in the mechanism that allows situated agents to plan ahead in time (see section 4.1), and the results suggest that the agent with neurocontroller B can do so better than with neurocontroller A, we will concentrate on the behaviour of the agent with neurocontroller B in the analyses below.

To assess which mechanism allows the agent controlled by neurocontroller B to plan ahead in time, we perform two analyses: (i) an analysis of the produced behaviour of the optimised agent (subsection 4.6.1), and (ii) an analysis of the internally simulated behaviour of the optimised agent (subsection 4.6.2).

4.6.1 Produced behaviour

The first analysis is to examine the produced behaviour of the agent controlled by neurocontroller B in the problem space of the ToL task. Since the agent in the sToL model is non-reactive, the SST-diagrams introduced in chapter 3 are not useful to reveal the behavioural strategy of the agent (see section 3.5). Since the number of possible configurations of the ToL is relatively small (36), it is possible to illustrate the transitions between the states of the ToL in a diagram of the ToL problem space

h	Neurocontroller A		Neurocontroller B	
	\overline{M}	sd	\overline{M}	sd
4	0.8058	0.1159	0.8113	0.2324
6	0.8138	0.1360	0.8775	0.1261
8	0.7994	0.1015	0.8441	0.1698
10	0.8588	0.1358	0.8837	0.0864
12	0.7110	0.2128	0.7537	0.0988
14	0.4283	0.1354	0.7547	0.1501

Table 4.3: The average ratio of moves made \overline{M} and the corresponding standard deviations for neurocontroller A and neurocontroller B.

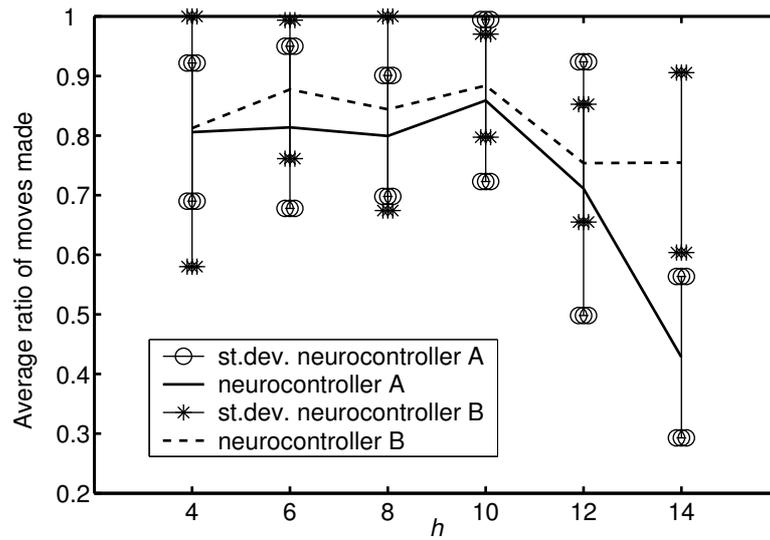


Figure 4.9: The average ratio of moves made \overline{M} as a function of the number of hidden nodes and context nodes (h) for neurocontroller A (solid line) and neurocontroller B (dashed line).

(cf., Berg and Byrd (2002)). In such a diagram the behavioural strategy of a subject performing the ToL task can be drawn as a sequence of boxes (states) and arrows (transitions).

In figure 4.10, the behavioural strategy of the agent is illustrated in a diagram of the ToL problem space. The numbered boxes in the figure represent ToL states (again, R, G, and B, indicate the colours of the balls, red, green, and blue, respectively). The black box represents the starting configuration (the state labelled 63) and the shaded boxes represent the twelve goal configurations (the states labelled 66, 55, 53, 54, 13, 44, 22, 46, 11, 21, 43, and 45) of the test described in subsection 4.3.1. While the arrows through the problem space indicate the permitted transitions between ToL states that constitute the behavioural strategy of the agent, all other permitted transitions between ToL states are indicated by dashed lines. To further relate the diagram in figure 4.10 to the test on which agents in the sToL model are optimised (described in subsection 4.3.1), we reproduce table 4.1 in table 4.4 and augment it with a column specifying the state numbers depicted in figure 4.10. The behavioural strategy illustrated in figure 4.10 gives rise to the two observations discussed below.

The first observation is that the behavioural strategy of the agent with neurocontroller B is sub-optimal. The trajectory through the problem space that is depicted in figure 4.10 shows that the agent behaves optimally for ten of the twelve problems in the test, but has sub-optimal behaviour for two problems in the test, viz. for problems 6 and 11, with states 44 and 43 as goal states, respectively. A shorter trajectory from the starting state to these goal states than the one indicated by the arrows exists. For example, the agent would solve problem 6 more efficiently, i.e., reach state 44 by a shorter trajectory through the problem space, if it would not enter state 53 (i.e., move the blue ball) after state 55 is visited for the first time, but would directly enter state 56 (i.e., move the green ball), as it does when state 55 is visited for the second time.

The second observation is that the agent with neurocontroller B associates multiple moves with single ToL states (see states 15, 55, and 63 in figure 4.10), i.e., maps multiple moves onto single ToL states. The agent is able to produce a different move each time it visits ToL states 15, 55, and 63 (see figure 4.10), because (i) a visit co-occurs with a different goal state, and (ii) the agent can exploit its ability to use the temporal context of a visit. In the first case, the agent is able to produce a different move each time it visits the same ToL state, because a different goal state during each visit results in a unique sensor activation by which the agent can distinguish each visit (e.g., each time the agent visits state 15). This unique sensor activation is the result of superimposing the encodings of the goal state onto the ToL state in the sensor array (see subsection 4.3.2). In the second case, the agent is able to produce a different move each time it visits the same ToL state, because the architecture of its neurocontroller enables the agent to use the temporal context of a visit.

The ability of the agent to use the temporal context of a visit arises from (a) the internal recurrence at the hidden layer of the neurocontroller and (b) the feedback connection from the additional output layer to the input layer of neurocontroller B.

Starting state for all problems:

Starting state for all problems:			
R			
G B -		63	
Problem	Goal state	State number	Minimum # of moves
1	G - B R	66	2
2	G R B	55	2
3	B G R -	53	3
4	B G R -	54	3
5	G R B -	13	4
6	G B R -	44	4
7	B R - G	22	4
8	B - R G	46	4
9	B G R - -	11	5
10	G B R - -	21	5
11	G B R -	43	5
12	B R G	45	5

Table 4.4: Starting state and goal states of the twelve ToL problems employed by Shallice (1982; reproduced from Berg and Byrd (2002)) accompanied by the state number (according to figure 4.10) and the minimum number of moves in which each problem of the test can be solved.

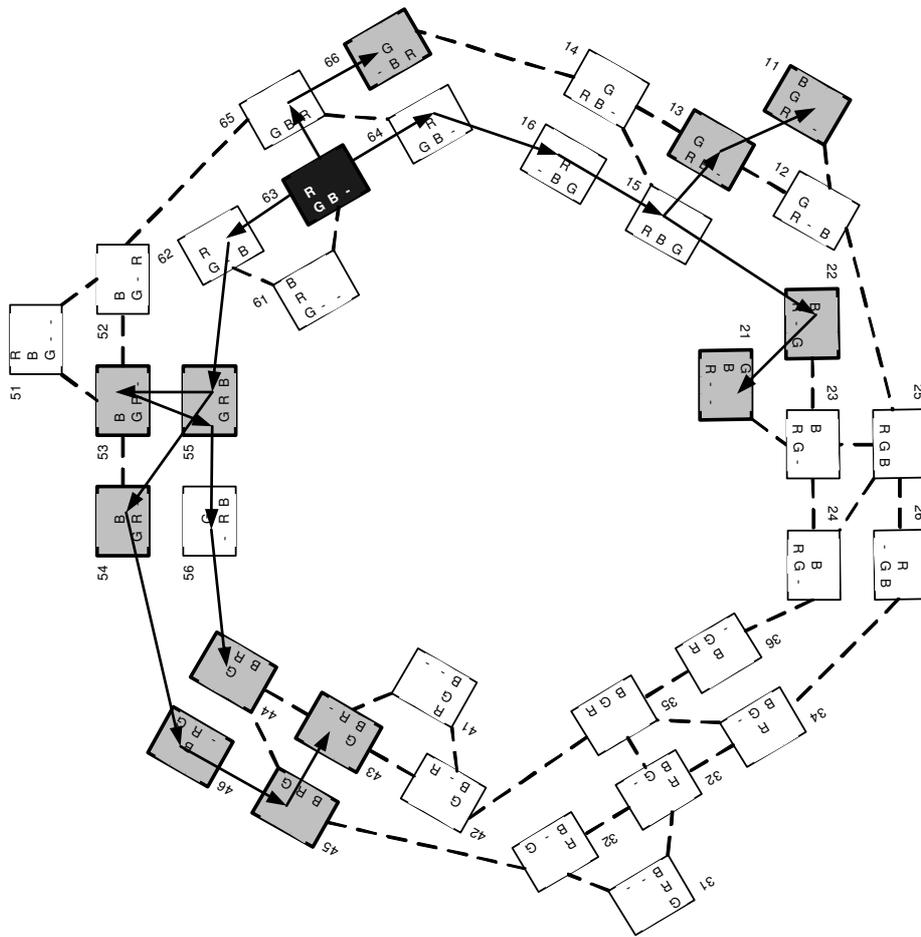


Figure 4.10: Trajectory through the ToL problem space for the optimised agent controlled by neurocontroller B. (The problem space representation is adopted from Berg and Byrd (2002).)

The internal recurrence at the hidden layer of the neurocontroller enables the agent to respond with a different move to the same sensory information. The activation generated in the context layer of the neurocontroller provides temporal context to the current situation on the basis of previously received sensory information (Elman, 1990; Bakker, 2004). The feedback connection from the additional output layer to the input layer of neurocontroller B enables the agent to respond with a different move to the same sensory information by generating an expected state (which is superimposed on the activity in the sensor array) that provides temporal context

to the current situation on the basis of previously received sensory information. By exploiting its ability to use the temporal context of a visit, the agent can even respond with a different move when the ToL state and the goal state are the same. This is the case when ToL state 55 is visited while the goal state is state 44. Then, the agent moves the blue ball to the first peg on the first visit and moves the green ball to the second peg on the second visit.

Since the agent maps multiple moves onto a single ToL goal-state combination (state 55 and state 44, respectively) only once, the agent seems to prefer one-to-one mappings of moves onto ToL goal-state combinations. This indicates that the agent takes advantage of its ability to determine the ToL states that it visits to select ToL goal-state combinations that are unique to the agent (cf., Nolfi (2002b)). Agents that do so are known to “select sensory patterns for which groups of patterns requiring different answers do not overlap too much” (Nolfi, 2002b, p.143). An internally simulating agent can even decrease this overlap by generating expected states. The analysis of the behaviour that is produced by the agent indicates that internalised interaction with the environment benefits planning performance on the ToL task by decreasing the overlap of sensory patterns that require different moves.

4.6.2 Internally simulated behaviour

The second analysis is a careful examination of the internally simulated behaviour of the optimised agent controlled by neurocontroller B. In order to perform the second analysis, we record the expected state of the optimised agent with neurocontroller B while performing the ToL task. A typical example of the sequences of expected states that are internally generated while solving a ToL problem is shown in table 4.5. In the left column (top to bottom), the table shows the sequence of ToL states that are visited by the agent while solving problem 9 of the test. The associated sequences of expected states that are generated by the internal simulation mechanism during each visit of a ToL state are shown in the right column (left to right) of the table. The arrows in the right column indicate the temporal order of the internally generated expected states. The number of expected states that are internally generated before the ToL state is changed by the agent varies per visited ToL state. The reason for this is that the neurocontroller does not always produce actions that satisfy the constraints of the ToL. The neurocontroller generates expected states until a valid action is produced. No expected states are generated after the goal state is reached, because the problem terminates upon reaching this state. The example shown in table 4.5 gives rise to the four observations discussed below.

The first observation is that an agent sometimes generates the same expected state repeatedly. For instance, the second expected state that is internally generated by the agent while the starting state is visited (see the first row of the second column) is the same as the first expected state that is internally generated. Apparently, the agent often simulates the same move multiple times to generate other simulated moves and, eventually, produce a move that changes the ToL.

The second observation is that the expected states that are internally generated

Problem 9:

ToL state	Expected state
R G B - (starting state)	R → R → R → B G B - → G B - → G B - → R G - -
R G B -	R → R → R G B - → G B - → G B -
R - B G	R G B → R G B -
R B G	B → B → R B G → B R - G → R - G → R - G
G R B -	G → B → B → G → B → R - B → G → R - - → R - - → R - B → G → R - -
	(cont.) B → B → B → B → B G → G → G → G → G R - - → R - - → R - - → R - - → R - -
B G R - - (goal state)	

Table 4.5: The ToL states that are visited (left column, top to bottom) and the expected states that are internally generated (right column, left to right) by the optimised agent with neurocontroller B while solving problem 9 of the test.

by the agent are often different from each other (i) between sequences of expected states, and (ii) within sequences of expected states. The difference between sequences of expected states is demonstrated by the difference between the expected states that are generated while the starting state is visited (see the first row of the second column) and the expected states that are generated while the third ToL state is visited (see the third row of the second column). The difference within sequences of expected states is demonstrated by the first expected state and the last expected state that are internally generated while the starting state is visited (see the first row of the second column). That the internally generated expected states differ within the sequences of expected states indicates that the agent can change the expected state that is internally generated independent from changing the ToL state, i.e., the

agent can make a simulated move in the absence of a real move.

The third observation is that the expected state that is internally generated by the agent may occur as a future ToL state. Table 4.5 shows that the goal state of problem 9 (the last ToL state in the sequence of visited ToL states shown in the left column) is already simulated internally before the last move to solve the problem is actually made. This seems to indicate that the last move is internally simulated before it is made. However, it takes the internal generation of eight more expected states, after the first time the goal state is simulated internally, before the last move is actually made and the problem is solved.

The fourth observation is that the expected states that are internally generated by the agent mostly match the ToL state that is currently visited or never occur as ToL states at all. For instance, the first expected state that is internally generated while the starting state is visited (the first state in the first row of the second column) matches the starting state of the problem, while the last expected state that is internally generated while the starting state is visited (the last state in the first row of the second column) never occurs as a ToL state at all. Apparently, the internal simulation mechanism does not just feed back previous ToL states or predict future ToL states.

These four observations indicate that the optimised agent with neurocontroller B in the sToL model is able to simulate future moves internally before actually performing them, but only occasionally does so. Although the expected state that is internally generated by the agent may occur as a future ToL state, it is debatable whether the optimised agent with neurocontroller B exploits its ability to internally simulate moves to predict future states in order to plan ahead in time. This is debatable for two reasons. The first reason is that the expected state only occasionally matches a future ToL state, but mostly matches the currently visited ToL state or never occurs as a ToL state at all. The second reason is that when the expected state does match a future state, the agent continues to internally generate expected states before it actually produces the move to change the ToL into the future state. This suggests that the internal simulation mechanism may merely serve the function of building up sufficient activation in the neurocontroller to produce a certain move, rather than to simulate a future move.

Since the internally generated expected state changes independently from changes in the ToL state, the success of the agent's behaviour appears to depend on its ability to map moves onto the expected state (besides onto the ToL state and the goal state). Therefore, the expected state is crucial for the success of the agent on the ToL task, i.e., for its ability to plan ahead in time.

Newell and Simon (1972) regard symbol manipulation to be the processing of symbols to produce an output that benefits an objective. Although, the results reported in section 4.5 indicate that the generation of expected states benefits the agent's objective (i.e., reaching the goal state in the least possible number of moves), the expected state may or may not be regarded as a symbol. The expected state may be regarded as a symbol, because it symbolises a configuration of the ToL, but it may not be regarded as such, because it is distributed over several sensors

(Fodor and McLaughlin, 1990). Therefore, it is debatable whether the generation of expected states may be considered to be symbol manipulation. Disregarding whether an expected state is a symbol, the expected state can be attributed a symbolic value, since, as the observations reported above show, the generation of an expected state can be regarded to simulate external events (changes in the ToL) internally. We claim that the internal simulation of behaviour should be considered a situated variant of symbol manipulation.

In summary, the analysis of the behaviour that is internally simulated by the agent indicates that the ability of the agent in the sToL model to generate expected states is crucial to plan ahead in time. Furthermore, the analyses led us to claim that internal simulation of behaviour should be considered a situated variant of symbol manipulation.

4.7 Evaluation by comparison

The performance of the agent on the ToL task in the sToL model is compared to the performance of human subjects on the ToL task in subsection 4.7.1. In subsection 4.7.2, the design of the sToL model is compared to that of two other situated models of the ToL task. Thereafter, in subsection 4.7.3, we compare the sToL model to other models that incorporate internal simulation.

4.7.1 Comparison with human subjects

To evaluate the performance of the agent in the sToL model, we compare the performance of the agent with that of human subjects. For this purpose, we plot the average number of moves made by human subjects on Shallice's (1982) test reported by Owen *et al.* (1990) with that of the optimised agent with $h = 14$.

In figure 4.11 the average number of moves made by human subjects on Shallice's (1982) test reported by Owen *et al.* (1990) is plotted together with that of the optimised agent with neurocontroller A. The figure shows that human subjects use fewer moves to solve the ToL problems in the test than the agent with neurocontroller A does. The error bars in the figure (showing the standard error) indicate that the difference between the performance of humans and that of the agent with neurocontroller A is considerable, since the error bars of both performances (human and agent) do not overlap. Unfortunately, only the mean performances (and standard errors) of human subjects on the ToL task reported by Owen *et al.* (1990) are available for comparison (Owen, personal correspondence), so that an ANOVA cannot be performed.

In figure 4.12 the average number of moves made by human subjects on Shallice's (1982) test reported by Owen *et al.* (1990) is plotted together with that of the optimised agent with neurocontroller B. The figure shows that human subjects use fewer moves to solve the ToL problems in the test than the agent with neurocontroller B does, except for problems with a minimal number of four moves. On those problems,

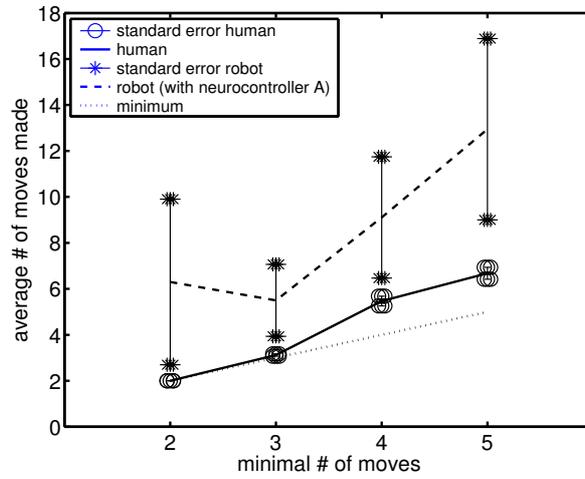


Figure 4.11: The average number of moves made by human subjects (reported by Owen *et al.* (1990)) and by the agent with neurocontroller A over the minimal number of moves in which a problem can be solved.

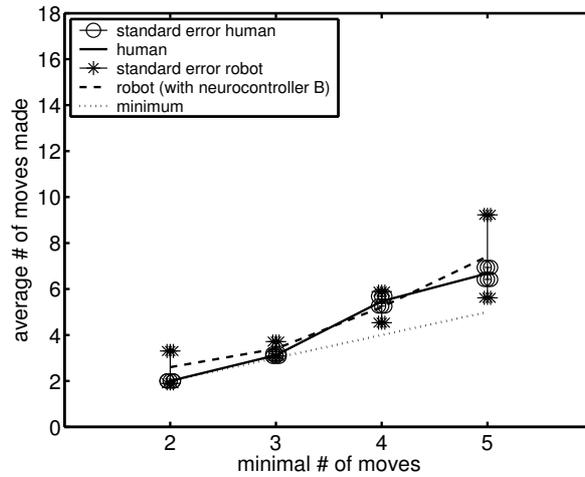


Figure 4.12: The average number of moves made by human subjects (reported by Owen *et al.* (1990)) and by the agent with neurocontroller B over the minimal number of moves in which a problem can be solved.

human subjects perform slightly better than the agent with neurocontroller B. The error bars (showing the standard error) in figure 4.12 indicate that the performance of humans and that of the agent with neurocontroller B do not differ significantly, since the error bars of the human performance completely overlap with those of the agent performance.

Taken together, figures 4.11 and 4.12 indicate that the performance of human subjects on the ToL task is dissimilar to that of the agent with neurocontroller A, but is similar to that of the agent with neurocontroller B. Apparently, the performance on the ToL task of an agent with the ability to simulate perception and behaviour internally more closely resembles human performance on the ToL task than the performance of an agent without the ability to do so.

4.7.2 Comparison with other situated models of the ToL task

An adequate comparison of performance between our sToL model and other situated models applied on the ToL task is not feasible, because all other situated models applied on the ToL task reported in the literature use tests that contain a different number of ToL problems and ToL problems of a different complexity than the test adopted in our sToL model. Below, we evaluate our sToL model by comparing the design of the model with that of two other situated models applied on the ToL task. First, we briefly describe the model reported by Dehaene and Changeux (1997), henceforth referred to as DC97, and compare it to our sToL model. Second, we briefly describe the model reported by Polk *et al.* (2002), henceforth referred to as PSLF02, and compare it to our sToL model.

The DC97 model (Dehaene and Changeux, 1997) is designed to mimic the planning behaviour of patients with prefrontal cortex damage (when the model is lesioned). The model consists of a neurocontroller that receives the current configuration of the ToL as input and produces a gesture toward one of the three balls of the ToL as output. Furthermore, DC97 is hierarchically organised, with different levels for planning, operating, producing gestures, and evaluation. Since, in contrast to the sToL model described in section 4.3, the DC97 model is hierarchically organised, and produces gestures instead of directly changing the ToL by the neurocontroller's output, the DC97 model is biologically and psychologically more plausible than our sToL model. However, the DC97 model does not satisfy the parsimony condition discussed in subsection 1.4.1, which indicates that our sToL model may be better suitable for debates on situated cognition.

The PSLF02 model (Polk *et al.*, 2002) is designed to test the hypothesis that the dorsolateral prefrontal cortex in the human brain represents internally generated subgoals while performing high-level cognitive tasks. To do so, in PSLF02, attractor networks are trained to map ToL states, goal states, and internally generated subgoals onto moves. Although, in contrast to our sToL model, in PSLF02 attractor networks are trained instead of simple recurrent neural networks, PSLF02 is similar to our sToL model in that it allows internal simulation of perception and behaviour. It can perform internal simulation, because it can feed back internally

generated ToL states from the subgoal module to the module representing the current ToL state. Experiments with PSLF02 show that the effect of lesions in the dorsolateral prefrontal cortex on the performance of patients on the ToL task, can be modelled by lesioning the module representing the subgoals in PSLF02. This supports our finding (reported in section 4.5) that internal simulation of perception and behaviour enhances performance on the ToL task. However, Polk *et al.* (2002) did not analyse the generated subgoals and the feedback from the subgoal module to the module representing the ToL state. To be able to do so, a model has to satisfy the transparency condition discussed in subsection 1.4.1. Since our sToL model is designed to satisfy this condition (see subsection 1.4.2), it allows the analysis of internally generated states and feedback (as is done in subsection 4.6.2), and is likely a better model to do so than PSLF02.

4.7.3 Comparison with other internal simulation models

Other situated models that incorporate internal simulation of interaction with the environment were studied by Tani and Nolfi (1999), Jirnhed *et al.* (2001), Hoffmann and Möller (2004), and Ziemke *et al.* (2005). The approaches of Tani and Nolfi (1999), Jirnhed *et al.* (2001), and Hoffmann and Möller (2004) differ in many ways from our approach. The most important difference is that their agents were explicitly trained to predict future states. In contrast, an agent with neurocontroller B in the sToL model is trained to exhibit appropriate behaviour. This leads to agents that simulate moves internally that are not necessarily predictions. Ziemke *et al.* (2005) also trained agents to behave appropriately. However, they trained agents on a task of collision-free corridor following, which is not a typical planning task for the study of situated symbol manipulation.

4.8 Discussion

In section 4.1 we claimed that the ToL task requires symbol manipulation on the basis of internalised interaction with the environment. This claim is supported by the results reported in section 4.5, which show that the optimised agent with neurocontroller B (i.e., with the ability to internalise its interaction with the world) outperforms the optimised agent with neurocontroller A (i.e., without the ability to internalise its interaction with the world) on the ToL task in the sToL model. Furthermore, the claim is supported by both analyses reported in section 4.6 as follows.

The first analysis (see section 4.6.1) indicates that internalised interaction with the environment benefits planning performance on the ToL task by decreasing the overlap of sensory patterns that require different moves. This analysis shows how internalised interaction with the environment benefits planning performance on the ToL task.

The second analysis (see section 4.6.2) indicates that internalised interaction

increases the ability of the agent in the sToL model to plan ahead in time by enabling the agent to map moves onto the expected state (besides onto the ToL state and the goal state). Furthermore, the analyses led us to claim that internal simulation of behaviour should be considered a situated variant of symbol manipulation.

In summary, the results from the experiment support our claim that the ToL task requires symbol manipulation on the basis of internalised interaction with the environment. Furthermore, the first analysis shows how internalised interaction with the environment benefits planning performance on the ToL task, and the second analysis shows that internalised interaction should be considered a situated variant of symbol manipulation.

Following the claim that the ToL task requires symbol manipulation on the basis of internalised interaction with the environment, in section 4.1, we formulated a main research question and two additional research questions. These research questions are answered in the next section.

4.9 Chapter conclusions

The results reported and discussed above show that our situated model can perform the ToL task, which answers the main research question formulated in section 4.1. The analyses of the behaviour of the optimised agent with neurocontroller B provide the answers to the two subsequent research questions formulated in section 4.1: (1) which mechanism allows a situated agent to plan ahead in time? and (2) how are the current and future states of the Tower of London represented in an agent?

The answer to the first question of the subsequent research questions is that the mechanism that allows a situated agent in the sToL model to plan ahead in time is the ability to perform symbol manipulation by internalised interaction with the environment. The ability to perform symbol manipulation by internalised interaction with the environment facilitates planning by decreasing the overlap of sensory patterns that require different moves.

The answer to the second question of the subsequent research questions is that the current and future states of the ToL are both represented by the mapping of moves to the activation of the sensor array of the agent with neurocontroller B. A current state of the ToL can be externally represented by the mapping of moves to the activation of the sensor array. A future state of the ToL can be internally represented by the mapping of simulated moves to the activation of the sensor array.

On the basis of the findings reported above, we may conclude that the ability to perform (situated) symbol manipulation by internal simulation allows situated agents to plan ahead in time. Furthermore, we may conclude that representation of both the current and future states of the environment occurs by mapping moves to the activation of the sensor array of the agent received from the environment and the internally generated expected state, respectively. To incorporate internal representation into the operationalisation of the notion of situated representation, the operationalisation should allow internal simulation of perception and behaviour.

Chapter 5

External representation

In chapter 3 it was demonstrated that representations in situated systems can be internal and/or external. Then, in chapter 4, we focused on the nature of internal representation. In this chapter¹, the aim is to study the nature of external representation in an environment in which there are no deterministic rules governing the environment. The agent-environment interaction is analysed by two different types of analysis (microscopic and macroscopic); they are briefly discussed in section 5.1. In section 5.2, the foraging agent model introduced in chapters 1 and 2 is discussed in more detail. Section 5.3 discusses the experiments conducted with the model and reports the results found. In section 5.4 and section 5.5 we conduct a macroscopic analysis and microscopic analysis, respectively. Both analyses investigate how the externally represented knowledge is accessed and used by the situated agent. Section 5.6 discusses the outcome of our investigations. In section 5.7, we provide our conclusions.

5.1 The basics of macroscopic analysis

In order to understand adaptive behaviour, Beer (2003b) suggested that one should start by investigating minimal models, as is a common approach in physics. In recent years, in-depth analysis of such models revealed many new insights into the processes underlying cognition (see, e.g., Beer (1995; 2000; 2003b), Nolfi (1998;

¹This chapter is an updated version of van Dartel M.F., Postma E.O., van den Herik H.J., and de Croon G. (2004) Macroscopic analysis of robot foraging behaviour. *Connection Science* 16(3), pp.169-181, and is in part reprinted from the Proceedings of the 4th Workshop Dynamic Perception, Würtz R.P. and Lappe M. (Eds.), van Dartel M.F., and Postma E.O. *Analysing Situated Behaviour from a Macroscopic Perspective*, pp 165-170, Copyright (2002), with permission from IOS Press. The experiments and results reported here were presented at the 4th Workshop Dynamic Perception in Bochum, Germany (see van Dartel and Postma (2002)) and at the 14th Belgium-Netherlands Conference on Artificial Intelligence in Leuven, Belgium (see van Dartel, Postma, and van den Herik (2002)). The author would like to thank his co-authors and the publishers for their kind permission to reuse relevant parts of the articles in this thesis.

2002a), Slocum *et al.* (2000), van Dartel *et al.* (2005), de Croon, Nolfi, and Postma (in press)). So far, analysis of single agents has been performed at a microscopic level, in which the focus is on the successful behaviour of individual agents. Although such microscopic analyses have led to explanatory insights and testable predictions for individual agents in specific environments, the insights and predictions are difficult to generalise (Guillot and Meyer, 2001). For example, in a study by Nolfi (2002a), a microscopic analysis of the sensorimotor behaviour of a few successful agents on a self-localisation task revealed that success on that specific task depends on the ability to integrate information over different time scales. Moreover, microscopic analyses reveal properties of agent behaviour on various tasks. At the same time the microscopic approach raises questions about the generality of the properties revealed. In the case of Nolfi's study, it remains to be established how general the strategy of integrating information over different time scales is in relation to self-localisation. We believe that for adequate handling of such issues of generality another approach is needed.

In this chapter we demonstrate that *macroscopic* analysis might be such an approach. Applied to robotics, macroscopic analysis analyses the average properties of many robots/agents in similar environments. So far, macroscopic analysis has only been applied to robotics in the context of multi-agent systems (see, e.g., Lerman *et al.* (2001), Lerman and Galstyan (2003), Jones and Matarić (2003)). In order to improve the generality of findings in single-agent systems, macroscopic analysis has to be conducted over many instances of the same agent on the same task.

For two reasons, macroscopic analysis might be adequate to handle issues of generality in evolutionary robotics. The first reason is that macroscopic analysis ignores individual differences by analysing the average properties of many agents in similar environments. Therefore, macroscopic analysis may be more suitable to reveal universal properties, i.e., properties characteristic for a class of agent-environment systems. The second reason is that the application of macroscopic analysis in statistical physics has revealed many universal properties; for instance on DNA sequences, heartbeat rates, and weather variations (Havlin *et al.*, 1999; Stanley *et al.*, 2001; Beekman, Sumpter, and Ratnieks, 2001).

Hence, we propose that macroscopic analysis of single-agent systems might improve the generalisability of contemporary research in robotics. To evaluate the effectiveness of macroscopic analysis in the context of single agent systems we pose the straightforward research question introduced in section 1.4.3: *Can macroscopic analysis reveal a universal property of adaptive behaviour in an agent model?*

In the experiments that follow, we optimise foraging behaviour of neural-network-controlled agents using an evolutionary algorithm. Then, we perform a macroscopic analysis on the foraging behaviour of the optimised agents. We expect that, by averaging over many behavioural interactions, macroscopic analysis will reveal properties that would remain obscured when only microscopic analysis were applied. Hence, we investigate whether macroscopic analysis will predict a universal property of foraging behaviour that can be confirmed by microscopic analysis. Additionally, we perform a microscopic analysis to examine whether a property revealed by the macroscopic analysis can be recognised in the microscopic behaviour of the agents.

5.2 The foraging model

The foraging model is described in terms of the environment (subsection 5.2.1), the agent (subsection 5.2.2), and the evolutionary algorithm by which the behaviour of the agent is optimised (subsection 5.2.3).

5.2.1 The environment

The environment is defined as an $L \times L$ square lattice with periodic boundary conditions (i.e., the environment is defined on a torus) containing n food elements. Randomly distributed dots over the L^2 locations represent the food elements. An agent collects food by walking over the food elements. Whenever a food element is collected, it is removed from the environment and replaced by a new one at a random location. In this way, the number of food elements remains constant. Figure 5.1 is an illustration of the environment with randomly distributed food elements (dots), and in the environment an agent (circle).

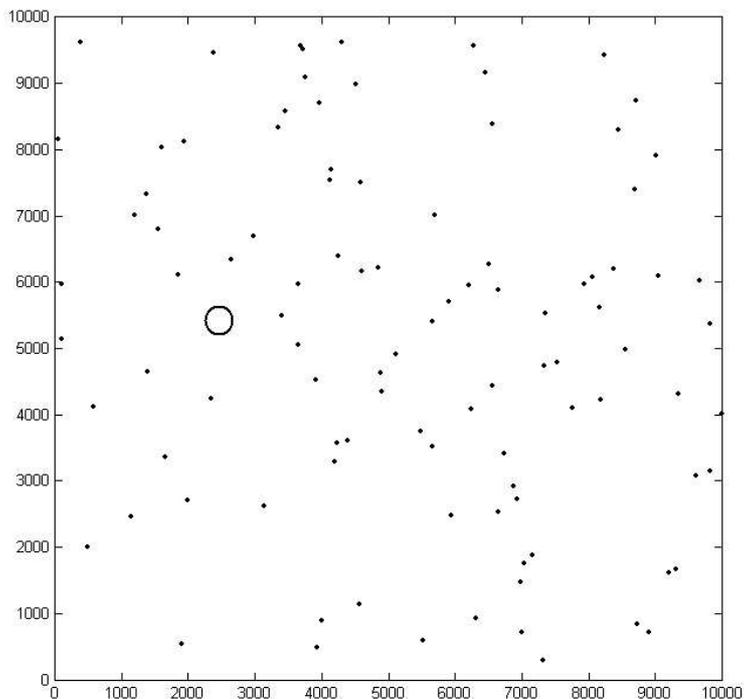


Figure 5.1: Illustration of an agent (circle) in the experimental environment with randomly distributed food elements (dots). The values on the x-axis and y-axis are spatial coordinates ($0 \leq x, y < L = 10,000$).

5.2.2 The agent

The agent performing the foraging task is controlled by a neural network and is defined in terms of a sensor mechanism and a brain. Both are described below.

Sensor mechanism

The sensor mechanism of the agent contains a sensor and an internal Gaussian noise source. The sensor detects the nearest food element within its circular field of view with radius r . It only processes the nearest food element within the field of view; the sensor is orientation sensitive. Defining the orientation of the agent by α and the orientation of the nearest food element by β , the sensor activation I (i.e., the input) is given by the normalised one-dimensional Von Mises basis function (Gutowski, 2001).

$$I = \frac{e^{k\cos(\alpha-\beta)}}{e^k} \quad (5.1)$$

where k is a positive constant proportional to the width of the basis function. The Von Mises basis function is the spherical analogue of the Gaussian basis function. The normalisation constant e^k ensures that the maximal value of the right-hand side term equals 1 when $\alpha = \beta$.

The internal Gaussian noise source (zero mean, standard deviation sd) is to be considered as a second input: it models the intrinsic noise of neural systems (cf., Zippelius and Engel (2003)). A food element is collected when the distance between the food element and the agent equals $0.1r$.

Brain

The brain (or neurocontroller) of the agent is a recurrent neural network with two input nodes (a food-detecting input and a noise source), h hidden nodes, and two output nodes. The input nodes are connected to the hidden nodes and to the output nodes. The hidden nodes have neural inertia by means of self-connections (Nolfi, 2002a). This means that the activity of the hidden nodes can be sustained for an extended period of time. The decay in activation is inversely proportional to the weight on the self-connection of the node. The initial network configuration is shown in figure 5.2. Each connection can be switched on or off during the evolutionary process, while retaining its weight value (cf., Spronck, Sprinkhuizen-Kuyper, and Postma (2001)). Initially, all weights are assigned random values symmetrically distributed around zero on the interval $[-rw, rw]$, with $rw > 0$. The transfer function for the hidden nodes is the sigmoid \tanh function that maps onto the interval $\langle -1, 1 \rangle$. The two output nodes control the agent's relative direction and step size, respectively. The output of the direction node is multiplied by π . A modulo operation restricts the direction to the interval $\langle -\pi, \pi \rangle$. The transfer function of the step-size output node is a semi-linear function $l = f(u)$ that maps negative values to zero and positive values u to the interval $\langle 0, uL/2 \rangle$, with $L = 10,000$.

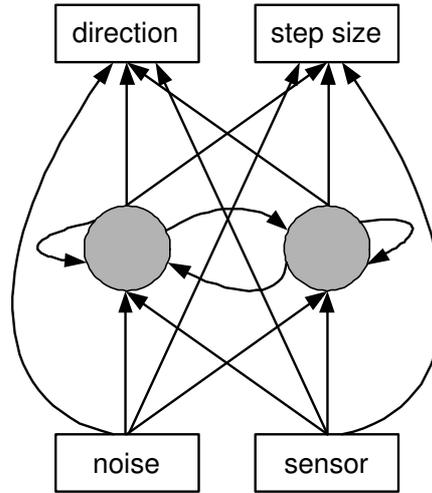


Figure 5.2: The initial network configuration of agents entering the evolutionary optimisation process.

5.2.3 The evolutionary algorithm

The real-valued weights of the neural network controlling the agent are optimised for foraging efficiency using the standard evolutionary-computation techniques of reproduction, crossover, and mutation (Goldberg, 1986). The connections of the agent's neurocontroller are represented in a genome as pairs of real and Boolean values. For each pair, the real value represents the weight of the connection and the Boolean value represents the presence or absence of the connection (cf., Spronck *et al.* (2001)). We employ the same evolutionary algorithm as described in subsection 3.3.5, except for the fitness function F , which is defined as follows.

$$F = \frac{1}{T} \left(\sum_{t=1}^T c(t) - \lambda \sum_{t=1}^T l(t) \right) \quad (5.2)$$

where t is an index for individual simulation steps ($t = 1, 2, \dots, T$) with T denoting the total number of steps, $c(t)$ is a function that returns 1 if a food element is collected at step t and 0 otherwise, $l(t) = f(u, t)$ is the step length of the agent at step t . The first term between the brackets favours food collection. The second term punishes long steps. The strength of the second term is set by λ .

5.3 Experiments and results

We performed a large series of experiments to optimise the foraging behaviour of agents. The goal of the experiments was to collect sufficient data of efficient foraging behaviour as required for macroscopic analysis. The foraging behaviour of 1000 agents was optimised. The optimal behaviour of each agent was the result of applying the evolutionary algorithm over 1000 generations with a population size of 1000 agents. All experiments were performed with the following parameter values: h (the number of hidden units) = 2, r (the radius of the circular field of view) = 1, L (the width and height of the environment) = 10,000, n (the number of food elements) = 100, rw (the parameter that defines the range in which weights are assigned values) = 0.5, k (the positive constant that is proportional to the width of the basis function) = 20, sd (the standard deviation of the Gaussian noise source) = 0.5, T (the total number of steps) = 10,000, and λ (the parameter that determines the impact of punishment for long steps) = 0.00001.

The experiments yielded a population of 1000 optimised foragers, with $-0.0045 \leq F \leq 0.0065$. After observing the behaviour of the optimised agents, we distinguished three types of foraging behaviour: (1) systematic behaviour (e.g., foraging in straight lines or circles), (2) random-walk behaviour, and (3) Lévy-flight behaviour. Systematic strategies were employed by 719 of the optimised agents, and considered sub-optimal because these strategies were employed by the optimised agents with the lowest fitness ($F < 0.0$). These 719 agents were left out of all further analyses. The other 281 optimised agents either employed the random-walk strategy shown in figure 5.3 or exhibited the Lévy-flight strategy shown in figure 5.4. Although most of these 281 fittest optimised agents exhibited the random-walk behaviour, the best ones exhibited Lévy-flight behaviour.

A characteristic feature of the fittest agents (see figure 5.4) is that their local random-walk behaviour is occasionally interrupted by large jumps. As a result, the area covered by these agents is much larger than the area covered by agents that only exhibit the random-walk behaviour (hence, there is a large difference between the depicted ranges of the x-axis and the y-axis in figures 5.3 and 5.4). The sudden jumps are known as Lévy flights (Viswanathan *et al.*, 1999; Gutowski, 2001; Viswanathan *et al.*, 2001). Foragers adopting a Lévy-flight strategy outperform the agents adopting a random-walk strategy in terms of foraging efficiency. The advantage of Lévy flights is that the probability of returning to a previously-visited site is smaller than it is for pure random-walk behaviour (Shlesinger and Klafter, 1985).

5.4 A macroscopic analysis

In this section a macroscopic analysis of the foraging behaviour of the optimised agents is performed. We expect that, by averaging over many behavioural interactions, macroscopic analysis reveals properties of foraging behaviour that would remain obscured when only microscopic analysis was applied.

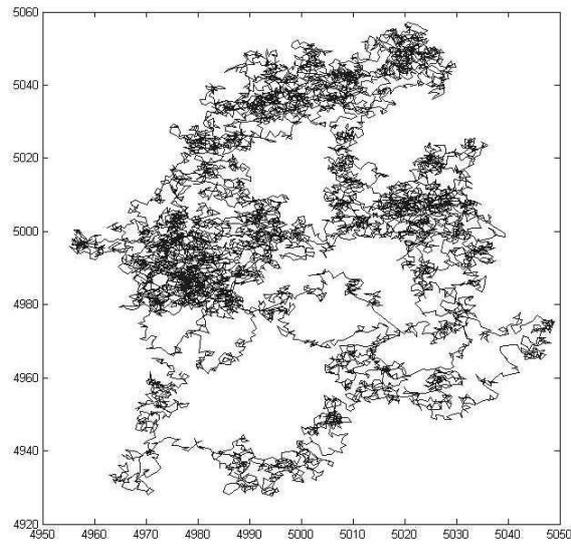


Figure 5.3: Illustration of random-walk foraging. The path consist of 10,000 steps. The values on the x-axis and y-axis are spatial coordinates with $0 \leq x, y < L = 10,000$, of which only the ranges $4950 \leq x < 5050$ and $4920 \leq y < 5060$ are depicted.

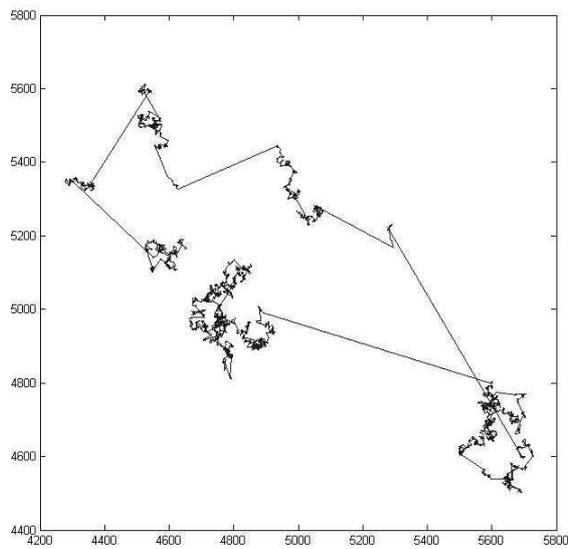


Figure 5.4: Illustration of Lévy-flight foraging. The path consists of 10,000 steps. The values on the x-axis and y-axis are spatial coordinates with $0 \leq x, y < L = 10,000$, of which only the ranges $4200 \leq x < 5800$ and $4400 \leq y < 5800$ are depicted.

Inspired by the success of macroscopic analysis in statistical physics, we conduct a macroscopic analysis on the efficient foraging behaviour of the optimised agents. In our context, macroscopic analysis involves the following three main steps: (1) the collection of a very large number of instantiations of the system or behaviour under study, (2) summarising the data in terms of a distribution or descriptors thereof, and (3) the identification of a property that is characteristic for (the behaviour of the) system. Our macroscopic analysis focuses on the quantification of the difference between random walks and Lévy flights in terms of a single parameter μ . The first step of the macroscopic analysis consists of the measurement of step lengths for a rather large number of foraging agents. Then, in the second step we summarise the data so obtained by a descriptor (μ) of the shape of the step-length distribution (Viswanathan *et al.*, 1999). Finally, in the third step, we identify the value of μ to be characteristic for the foraging behaviour of the agents. Concentrating on the probability of large step lengths, the tail of the step-length distribution scales (cf., Viswanathan *et al.* (2001)) according to,

$$P(l) = \frac{l^{-\mu}}{Z} \quad (5.3)$$

with $P(l)$ representing the probability of a step of length l , and Z a normalising constant. The parameter μ is proportional to the slope of the tail of the step-length distribution. In particular for a Gaussian function that generates random-walk behaviour, the parameter μ is larger than 3.0. Lévy-flight behaviour is associated with $1.0 < \mu \leq 3.0$. These values of μ yield ‘fatter’ tails. In our agent, the step-length distribution is generated from the step lengths produced by the output node.

To perform a macroscopic analysis we create step-length histograms of the foraging behaviour of optimised agents exhibiting both the random-walk and Lévy-flight behaviours. Figure 5.5 shows an example of such a step-length histogram. Subsequently, we analyse the tails of the log-transformed histograms by fitting a linear regression line through the data points. The slope of the line is an estimate of the value of μ that underlies the behaviour of both the random-walk behaviour and the Lévy-flight behaviour of agents. Figure 5.6 displays the regression line for an agent that exhibits the Lévy-flight behaviour shown in figure 5.4. The slope of the regression line of the log data is approximately equal to -2.0 (i.e., $\mu \approx 2.0$). In terms of equation 5.3 this corresponds to $P(l) = l^{-2}/Z$.

For each of the 281 fittest foragers a histogram (such as shown in figure 5.5) is created from several runs of T steps each. Subsequently, the value of μ is determined for each histogram. Figure 5.7 plots the fitness of the 281 fittest foragers as a function of μ . The values of μ ranged from $\mu \approx 2.0$ to $\mu \approx 4.3$. Figure 5.7 shows an increase of fitness for agents with a distribution of step lengths that can be characterised by a value of μ that is closer to 2.0. In other words, the best-performing foragers, viz. those employing a Lévy-flight strategy, can be associated with $\mu \approx 2.0$. Of the 281 fittest foragers, 61 agents employ a Lévy-flight strategy with $\mu \geq 2.0$. A range of animals exhibits efficient foraging behaviour that is characterised by Lévy flights with $\mu \approx 2.0$: albatrosses, foraging bumblebees, deer, and amoebas (Viswanathan *et al.*, 1999, 2001).

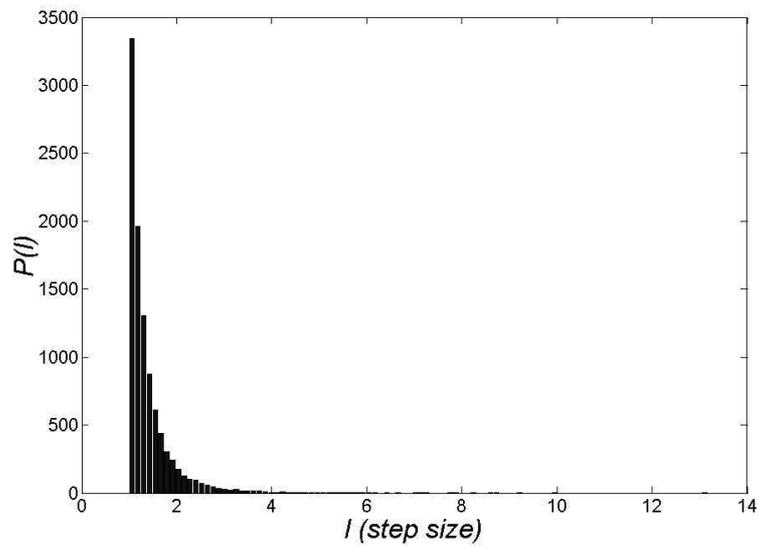
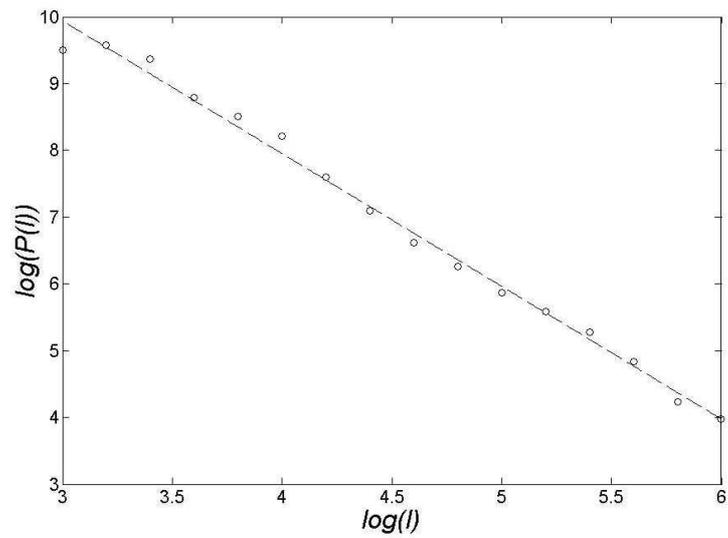


Figure 5.5: Histogram of step lengths.

Figure 5.6: Log-log plot of the tail ($l \gg 0$) of the histogram. The slope of the regression line is ≈ -2.0 .

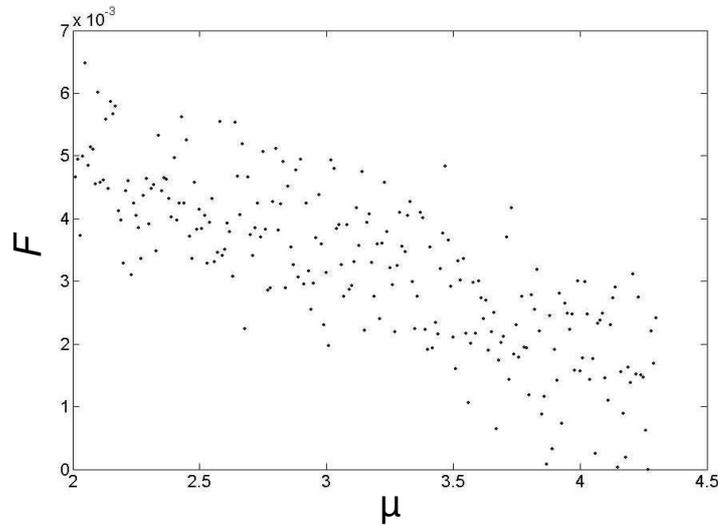


Figure 5.7: The fitness F as a function of the step-size distribution parameter μ for the 281 fittest foragers obtained through evolutionary optimisation of foraging behaviour.

We provisionally conclude that the macroscopic analysis revealed a property, viz. Lévy-flight behaviour as characterised by $\mu \approx 2.0$, that is universal to optimal foraging behaviour.

5.5 A microscopic analysis

In addition to the macroscopic analysis, a microscopic analysis is performed to investigate whether the universal property of foraging behaviour revealed by macroscopic analysis can be confirmed by microscopic analysis. Our microscopic analysis focuses on a successful foraging agent, i.e., an agent that exhibits Lévy-flight behaviour. The SST-diagrams introduced in chapter 3 cannot be applied to the behavioural strategies of the foraging agents, since the environment of the foraging task is not discrete and the neurocontroller of the agents is non-reactive. Therefore, the analysis that aims at explaining the mechanism that generates the efficient foraging behaviour is not conducted by SST-diagrams, but by inspecting the optimised weights of successful agents. Three interesting examples of such agents are shown in figure 5.8: the food-detecting sensor of these agents is disconnected. As can be seen in the figure, the internal noise fully determines the behaviour of these agents. In what follows, we will focus on the functioning of the agent with the neurocontroller that is illustrated in the middle of figure 5.8. Our analysis of the functioning of this par-

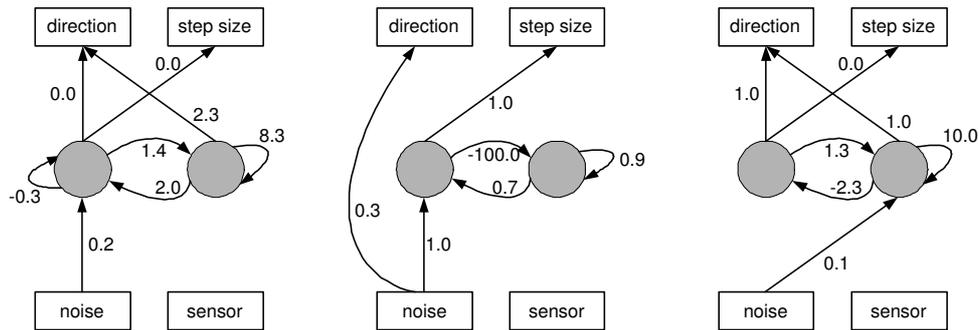


Figure 5.8: Three examples of the network configurations of successful foraging agents. Arrows indicate connections and the associated numbers indicate the weight values of the connections. A connection with an associated number of 0.0 is switched on, but has a very small (< 0.05) weight value.

ticular agent reveals the following mechanism to underlie the Lévy-flight behaviour. The noise source excites the left hidden node randomly. Whenever the excitation exceeds the threshold value of the node, it is activated and generates a large step (i.e., a flight). At the same time, its activation excites the right hidden node that, in turn, inhibits the left node. The neural inertia of the right node sustains the inhibition for an extended period. As a result, the large steps occur infrequently, as required for the efficient Lévy-flight strategy. The neural inertia realises a memory that acts as a counting mechanism. Next to these extreme examples there were other examples (some with connected food-detecting sensors) of efficient foraging agents.

We provisionally conclude that the microscopic analysis can explain the universal property revealed by macroscopic analysis at a microscopic level. It is important to note that, without macroscopic analysis, microscopic analysis could have revealed Lévy-flight behaviour to be characteristic for optimal foraging, but could not have revealed the exact relation between foraging success and the property μ , and therefore could not have revealed the universal property of optimal foraging behaviour $\mu \approx 2.0$.

5.6 Discussion

Below, we discuss the macroscopic and microscopic analyses. In section 5.6.1, we provide an example of how macroscopic analysis can be applied to complement microscopic analysis in existing research in evolutionary robotics. In section 5.6.2, we discuss our macroscopic and microscopic analyses of foraging behaviour in relation to the framework of scientific research as described by Cohen (1995).

5.6.1 Applying macroscopic analysis

To illustrate how macroscopic analysis might generalise findings from existing studies in evolutionary robotics, we adopt the results of a study performed by Nolfi (2002a). First, we briefly discuss the study and its findings. Then, we explain how macroscopic analysis can be applied to generalise those findings.

Nolfi (2002a) performed a microscopic analysis on agents that had to accomplish a task of self-localisation while driving around in an environment. He investigated a neurocontroller with nodes that change at different time rates. The changing rate, i.e., the neural inertia, of individual nodes was optimised by an evolutionary algorithm. The experiments with the self-localisation task resulted in agents with both fast changing nodes and slow changing nodes. While the fast changing nodes detected regularities occurring at short time scales, the slow changing nodes detected regularities at longer time scales. The first finding of the study is that the self-localisation task requires an agent to be capable of detecting regularities on different time scales. The second finding is that the evolvable changing rates offer the capability of detecting regularities on different time scales to the studied agent.

A macroscopic analysis might be a valuable tool to make both findings easier to generalise. The first finding can be generalised by a macroscopic analysis that relates the self-localisation task to the distribution of neural changing rates of successful agents. It suggests that successful agents should possess both fast and slow changing nodes to be able to extract regularities on short and long time scales. Hence, the above-mentioned distribution might reveal a universal property of self-localisation, i.e., a property characteristic of the distribution of changing rates. To generalise the second finding, a similar macroscopic analysis can be performed on agents performing another task that poses different requirements on the detection of regularities in time. If another task only requires exploitation of regularities on short time scales, then the second finding suggests that evolution results in agents with nodes that change fast. Hence, we expect the macroscopic analysis of such a task to result in a distribution of changing rates that is different from the one that is characteristic of the self-localisation task.

In summary, the macroscopic analysis that we suggest might be able to reveal a universal property of self-localisation.

5.6.2 Relation to the framework of scientific research

According to Cohen (1995, p.4) the task of science is to provide theories to answer the following three basic research questions. (1) How will a change in the agent's structure affect their behaviour given a task and an environment? (2) How will a change in an agent's task affect its behaviour in a particular environment? (3) How will a change in an agent's environment affect its behaviour on a particular task?

In the context of our study, the agent's structure corresponds to the neurocontroller and sensor mechanism, the behaviour is defined by the relative direction and step size provided by the two output nodes, the task is foraging, and the environment consists of randomly distributed food elements.

The first basic research question is addressed in our study by evolving the agent's neurocontroller to enhance the foraging behaviour. Both macroscopic and microscopic analyses answer this question, albeit in different ways. The macroscopic analysis answers this question by relating the distribution of step-lengths produced by the agent's neurocontroller to the foraging behaviour observed (see section 5.4). The microscopic analysis answers this question by relating the structure of the neurocontroller to the observed behaviour of the agent (see section 5.5).

The second basic research question is not addressed in our study because we do not vary the task of foraging.

Finally, the third basic research question is addressed by analysing the way in which an agent copes with the change in distribution of food elements in the environment that results from the agent's own foraging. The macroscopic analysis can supply the answer to this question on its own, whereas the microscopic analysis depends on the result from the macroscopic analysis to answer it. The macroscopic analysis answers this question by revealing the universal property of $\mu \approx 2.0$ as the optimal way to cope with the changing distribution of food elements (see section 5.4). The microscopic analysis cannot answer the third basic research question on its own, because the analysis of a single (or sometimes a few) behaviour(s) cannot detect a universal property that is revealed by the macroscopic analysis. The reason is that the universal property is a characteristic of the average agent-environment interaction, rather than a characteristic of a single (or a few) agent-environment interaction(s). Only after the macroscopic analysis had revealed the answer to the third basic research question (viz. Lévy flights as a characteristic of efficient foraging) the mechanism causing the behaviour at a microscopic level (viz. the neural inertia of the analysed neurocontroller) could be determined. So, macroscopic analysis enabled us to answer the third basic research question.

To perform the task of science, viz. provide theories to answer the three basic research questions (Cohen, 1995), macroscopic analysis is preferred over microscopic analysis. Figure 5.9 illustrates this by depicting the space of basic research questions. Understanding progresses from descriptions, through prediction, towards causal explanations. Many early microscopic analyses of agents comprise merely descriptions of agent behaviour (see, e.g., Brooks (1986), Steels (1994)). More recently, microscopic analyses have become more focused on the underlying mechanisms (see, e.g., de Croon *et al.* (in press)). Our contribution is to move from specific models towards more general models along the generalisation axis (see figure 5.9) by employing macroscopic analysis. Evidently, additional macroscopic analyses of agent behaviour are needed to arrive at more general models. As in many other domains where macroscopic analysis has been successfully applied (Stanley *et al.*, 2001), it may lead to a conceptual framework for universality in complex behavioural systems. Future investigations may reveal to what extent other insights into statistical physics can be transferred to robotics.

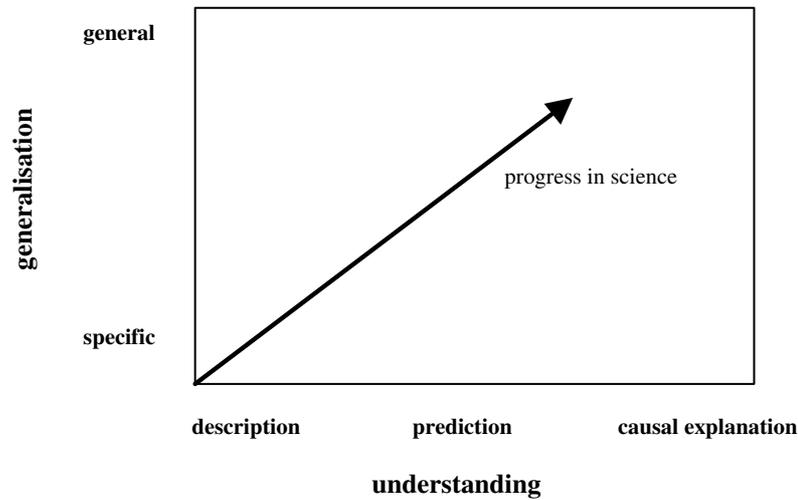


Figure 5.9: The two-dimensional space of versions of the basic research questions (Redrawn from Cohen, 1995).

5.7 Chapter conclusions

Using macroscopic analysis, we revealed a universal property of foraging behaviour in evolutionary optimised agents, viz. Lévy flights as characterised by $\mu \approx 2.0$. By doing so, we have shown that macroscopic analysis of an agent model can reveal a universal property of adaptive behaviour.

From the analyses and discussion we may draw the following three conclusions. (1) By combining the two provisional conclusions that followed from our macroscopic and microscopic analyses of efficient foraging behaviour (see sections 5.4 and 5.5, respectively) we arrive at the conclusion that macroscopic analysis may predict a universal property that can be explained at the microscopic level by microscopic analysis. (2) From the discussion in subsection 5.6.1, on an example of macroscopic analysis applied to existing research in evolutionary robotics, we conclude that macroscopic analysis may complement microscopic analysis in the study of adaptive behaviour. (3) From the discussion in subsection 5.6.2, on how our macroscopic and microscopic analyses of foraging behaviour relate to the framework of scientific research described by Cohen (1995), we conclude that macroscopic analysis may be preferred over microscopic analysis, owing to its power to reveal universal properties.

Given these three conclusions, and the success of macroscopic analysis in statistical physics and other disciplines (Stanley *et al.*, 2001), we expect it to generate

novel insights into the universal properties of adaptive behaviour in agent models and natural systems, which is crucial to the understanding of the interaction between an agent and its environment.

Our analyses revealed that optimised foraging agents are tuned to the statistics of the environment to optimise food intake. The foraging agents perform well on average by relying on the average interaction with the environment. The agent shown in figure 5.8 does not receive any sensor information. The only coupling with the environment is through the fitness value. Whereas the agents in ACP rely on their sensors to ensure a tight coupling with the environment (see chapter 3), in the foraging agents the coupling is indirect. The fittest agents are most successful in foraging by matching their stochastic behavioural dynamics to the statistics of the environment. Still, the main question is: What does our macroscopic analysis of foraging agents reveal about the nature and development of situated representations?

The experiment reported above shows that external representation may reside in the average properties of the interaction with the environment. Agents in the model of foraging do not represent individual food elements (or their physical locations) by element-specific interaction, but represent the uniform distribution of those food elements in their average interaction with the environment. The operationalisation of situated representation should allow representation by the average interaction with the environment.

Chapter 6

The operationalisation of situated representation

In chapter 2 we elaborated on the challenge posed by the incorporation of interaction in a theory of mind. Thereafter, in chapters 3, 4, and 5, we empirically investigated this challenge by the formalisation of mental processes within the theory of situated cognition.

The experiments and analysis performed in chapter 3 revealed that, in situated agents, representation can be understood as residing both in the brain and in the environment. Agents in the model of active categorical perception are able to distinguish small from large objects by representing the size of an object in the interaction with the environment. Successful agents in the model evolve object-specific interaction for each type of object.

Subsequently, chapter 4 showed that internal simulation of perception and behaviour may realise internal representation. The agent in the situated Tower of London (sToL) model is optimised to solve Tower of London (ToL) problems. In the optimised agent, the sensory activation consists of the current state of the ToL task, the goal state of the ToL problem, and a ToL state that is internally generated. So, to respond correctly to a certain ToL state, the agent combines internally represented and externally represented knowledge. Therefore, ToL states are partly represented by state-specific interaction of the agent with the ToL, and partly represented by internal simulation of such interaction.

Then, chapter 5 revealed that external representation may reside in the average properties of the environment. The best-performing agents in the model of foraging adopt a Lévy-flight strategy. Agents in the model of foraging do not represent individual food elements (or their physical locations) by element-specific interaction, but they represent the uniform distribution of those food elements by interaction that is specific for such a distribution.

These findings indicate that knowledge is represented in the situated systems

studied in chapters 3, 4, and 5, on the basis of each systems' interaction with its environment. The situated systems studied in these chapters represent the type of object in the ACP model (in chapter 3), the ToL state, the goal state, and the expected state, in the sToL model (in chapter 4), and the distribution of food elements in the situated model of foraging (in chapter 5), by learning to interact in a way that is specific to each type, state, or distribution, respectively. This implies that situated systems represent entities (i.e., anything with separate and distinct existence), e.g., a type, a state, or a distribution, by *entity-specific interaction*.

A system performs *entity-specific interaction* when it interacts with its environment in a way that is characteristic for the entity. For instance, when a human's eye is fixated on a certain position on the surface of a tomato, the sensory input that is received from the red round surface of the tomato will determine a shift of the eye towards another position. Thereafter, the same will happen again, eventually yielding a sequence of eye movements and sensory inputs. In the case of a banana, the sensory input received from the yellow slightly-bent surface of the banana during the first fixation results in a different shift of the eye, eventually yielding a sequence of eye movements and sensory inputs that is very different from the one yielded by a tomato. In this way, every entity yields a unique sequence of interaction, although some entities may yield more similar interaction (e.g., two tomatoes) than other entities (e.g., a tomato and a banana).

The investigations reported in chapters 3 and 5 show that situated systems can represent entities by evolving the ability to perform entity-specific interaction. The investigation in chapter 4 shows that situated systems can represent entities by combining entity-specific interaction that is performed with internally simulated entity-specific interaction.

On the basis of our findings we formulate the following new operationalisation of (the notion of) representation.

For an entity to be adequately represented in a system, it is implied that the system is able to perform and/or simulate internally the entity-specific interaction with the environment.

According to the new operationalisation, a system can represent an entity by being able to (i) perform the entity-specific interaction with the environment, (ii) simulate internally the entity-specific interaction with the environment, or (iii) both perform *and* simulate internally the entity-specific interaction with the environment. This means that a representation is always positioned somewhere on the continuum between external representation and internal representation (see figure 2.5), either at the extreme ends (in the cases (i) and (ii)) or somewhere in between (in the case (iii)). We consider an entity to be part of its environment. For instance, according to the new operationalisation of representation, humans represent tomatoes by being able to (i) interact in a tomato-specific way with tomatoes, (ii) simulate such specific interaction internally, or (iii) do both. The new operationalisation enriches debates in the cognitive sciences on situated systems with a working definition of the notion of representation that clearly indicates when something is a representation.

Bickhard (1999) states that “the fundamental form of representation, from which all others are derivative” is “naturally emergent in the organization of interactive systems” (*ibid*, p.435). The new operationalisation defines situated representation as such ‘interactive representation’ (Bickhard, 1999; 2004) and derives two basic forms; external representation and internal representation.

The new operationalisation of representation differs from the classic notion of representation in artificial intelligence as formulated by Newell (1980, p.156), viz. that something is a representation when “an entity X designates an entity Y relative to a process P, if, when P takes X as input, its behaviour depends on Y” (see chapter 1). We will refer to this classic notion of representation as the operationalisation of *non-situated* representation and refer to the new operationalisation as the operationalisation of *situated* representation.

The operationalisation of situated representation offers a number of advantages over the operationalisation of non-situated representation. Four of these advantages are discussed in section 6.1. Furthermore, the operationalisation of situated representation may have some implications for the fields of artificial intelligence, cognitive neuroscience, and cognitive psychology. These implications are discussed in section 6.2.

6.1 Advantages of the new operationalisation

The advantages of the operationalisation of situated representation over the operationalisation of non-situated representation concern (i) external representation and internal representation (subsection 6.1.1), (ii) the representation debate (subsection 6.1.2), (iii) situated accounts of cognition and awareness (subsection 6.1.3), and (iv) the symbol grounding problem (subsection 6.1.4). In subsection 6.1.5 we will relate both operationalisations to each other.

6.1.1 External representation and internal representation

The first advantage of the operationalisation of situated representation over the operationalisation of non-situated representation is that it distinguishes external representation and internal representation, and relates the two by stating that the latter is the internalisation of the first (a view that Dartnall (2005) refers to as ‘internalism’). We will discuss this advantage below.

In section 2.3, we argued that an operationalisation of situated representation should account for representation at any location on the continuum between external representation and internal representation (see figure 2.5), because, for all cognitive tasks, knowledge is represented somewhere on that continuum. In addition, we stated in section 2.3 that the symbol-system hypothesis of classic AI does not relate external representation to internal representation. The symbol-system hypothesis does not do so, because the operationalisation of non-situated representation, to

which the hypothesis adheres, does not distinguish between external representation and internal representation.

The operationalisation of situated representation distinguishes between external representation and internal representation, and relates them as follows. According to the operationalisation of situated representation formulated above, a representation can be absent, or lie somewhere on the continuum between external representation and internal representation. If a system encounters an entity with which it has no prior experience, the representation of the entity is absent. This implies that the interaction with the entity will not be specific to the entity, and, as a result, it will not be recognised by the system. If a system encounters an entity with which it has little prior experience, the representation of the entity is external. This implies that the object will bring about entity-specific interaction, and the entity will be recognised by the system. If a system encounters an entity with which it has considerable experience, the representation of the entity is internal. This implies that the system can simulate the entity-specific interaction internally, and it can reason about the entity in the entity's absence. An internal representation does not replace a corresponding external representation, but enhances the latter by internalising it. However, an internal representation may exist in the absence of a corresponding external representation.

6.1.2 The representation debate

The second advantage of the operationalisation of situated representation over the operationalisation of non-situated representation is that it causes progress in the so-called representation debate. The new operationalisation causes progress in two ways: (i) the new operationalisation unites the view that cognition is based on internal models of the world with the Gibsonian view that perception is direct, and (ii) the new operationalisation generalises Shepard and Chipman's (1970) second-order isomorphism to situated representations. In the following, we will discuss these two ways in which the new operationalisation causes progress in the representation debate.

The first way in which the new operationalisation causes progress in the representation debate concerns the part of the debate that has to do with the existence of mental representation. Marr's (1982) *computational theory of perception* and Gibson's (1979) *direct perception theory* represent two contrasting views on the existence of mental representation. In Marr's (1982) computational theory of perception, the world is mentally represented by a detailed three-dimensional internal model of the world, that is reconstructed from and matched to two-dimensional sensory data (Gordon, 1989). In contrast, according to Gibson's (1979) direct perception theory, the world is perceived through invariances in the stimulation from the environment, i.e., the perception of the world is *direct*, without the mediation of an internal model (Gordon, 1989), i.e., without the use of mental representations. The latter view has also been advocated by embodied robotics (see, e.g., Brooks (1990;1991)) and dynamic systems theory (see, e.g., Thelen and Smith (1994)).

The new operationalisation agrees with the Gibsonian view that perception is *direct* when it occurs on the basis of external representation. However, the new operationalisation also allows perception by the use of an internal model of the world, when it occurs on the basis of internal representation. Therefore, it unites the view that cognition is based on internal models of the world with the Gibsonian view that perception is direct.

The second way in which the new operationalisation causes progress in the representation debate concerns the part of the debate that has to do with the form of a mental representation. In this part of the debate, the central issue is whether representations are isomorphic to reality, i.e., whether they are in some way analogues to what they represent (see, e.g., Haselager, de Groot, and van Rappard (2003)). In Marr's view (described above), there is a one-to-one mapping of the physical world onto the mental world. Therefore, in his view, mental representations of entities are isomorphic to the physical entities that they stand in for. Shepard and Chipman (1970) introduced a more moderate position, by suggesting that the isomorphism between the physical world and the mental world has a second-order nature rather than a first-order nature. In a first-order isomorphism, the structure of the mental world is isomorphic to that of the physical world. In a second-order isomorphism, the relations among internal representations of the mental world are isomorphic to the corresponding relations in the physical world. Although Marr's view regarding isomorphism is still commonly excepted (especially in the study of mental imagery (Pylyshyn, 2003a)), the idea of a second-order isomorphism between the physical world and the mental world is also very popular (see, e.g., Edelman (1998), Lacroix *et al.* (2004), Lacroix, Postma, and Murre (in press), Steyvers, Shiffrin, and Nelson (2004)).

Our operationalisation of situated representation suggests a new view regarding the isomorphism between the mental world and the physical world, that can be regarded as a situated extension of Shepard and Chipman's (1970) ideas. For example, the behaviour of agents in the model of active categorical perception, described in chapter 3, suggests that their representations are second-order isomorphic. In the active categorical perception model, an agent represents the type of an object by the specific interaction that the agent has with the object (see, e.g., the agent behaviour illustrated in figure 3.9(a)). Small objects are caught by an agent, because they yield the same characteristic agent-object interaction at different locations in the environment. Similarly, large objects are avoided by an agent, because they yield the same characteristic agent-object interaction at different locations in the environment. The second-order isomorphism in the representation of an object is in the similarity between the interaction that an agent has with each type of object. Moreover, representation by the agents in the active categorical perception model is second-order isomorphic, because the agents store relations in the physical world (the interaction between an agent and an object) to represent the size of an object, rather than first-order isomorphic. The latter would be the case when the structure of the physical world were represented directly, e.g., when the size of an object were represented by the number of cells it occupied in the grid. The behaviour of

agents in the model of active categorical perception illustrates that the new operationalisation generalises Shepard and Chipman's (1970) second-order isomorphism to situated representations.

6.1.3 Situated accounts of cognition and awareness

The third advantage of the operationalisation of situated representation over the operationalisation of non-situated representation is that it provides the conceptual progress concerning the notion of representation that situated accounts of cognition and awareness require. Below, we will discuss this advantage.

Accounts of cognition and awareness that incorporate the situatedness of cognitive systems (see, e.g., Thelen (1995), O'Regan and Noë (2001b), Beer (2003b), Noë (2004)) are often accused of undermining the importance of representation by supporters of the operationalisation of non-situated representation (see, e.g., Markman and Dietrich (2000b), Cohen (2001), van Gulick (2001), Brookes (2001), Bullock (2003), Edelman (2003), Markman and Brendl (2005), Simons and Rensink (2005)). However, the notion of representation is never actually dismissed in accounts of cognition and awareness (O'Regan and Noë, 2001a; Beer, 2003a; Noë, 2004). (Although de Pinedo and Noble (2003) argue that they should do so.) Noë (2004) and O'Regan (2001b) dismiss the operationalisation of non-situated representation, but do not supply any evidence against the existence of representation. The notion of representation, according to Noë (2005), just needs empirical and conceptual progress. The operationalisation of situated representation provides such conceptual progress. To illustrate this conceptual progress, we borrow an example from Noë (2004, p.217):

“Take a tomato out. Look at it. Yes, you have a sense that the facing side of the tomato is all there, all at once. But if you are careful you will admit that you don't actually experience every part of its visible surface all at once. Your eyes scan the surface, and direct your attention to this and that.”

We will discuss this example in terms of (i) the operationalisation of non-situated representation, and (ii) the operationalisation of situated representation, to show how the latter operationalisation provides the conceptual progress that the first operationalisation lacks.

First, in terms of the operationalisation of non-situated representation, the experience of the tomato is caused by the activation of a mental representation of a tomato. This representation fills in those parts of the tomato that are not actually experienced. However, it seems rather illogical for a perceiver to activate a mental representation of a tomato to act as if the tomato were not there, while it is readily available to the perceiver (Noë, 2004). Furthermore, experiments on scene perception and change blindness suggest that human perceivers do not construct detailed internal models of a scene (Noë, Pessoa, and Thompson, 2000; O'Regan and Noë, 2001b).

Second, in terms of the operationalisation of situated representation, the experience of the tomato is caused by the interaction with the tomato. Those parts of the tomato that are not actually experienced are represented by the ability to perform and/or internally simulate tomato-specific interaction. Since those parts of the tomato that are not actually experienced are readily available in this way, they are not perceived as ‘gaps’ in our experience. So, the operationalisation of situated representation makes way for Noë’s account of what happens in the example. According to Noë (2004), the experience of the tomato is caused by ‘virtual representation’ of the tomato, which means that the content is *accessible* rather than *represented* (by non-situated representations). This example indicates that the operationalisation of situated representation progresses the notion of representation to the point that it conceptually meets accounts of cognition and awareness that incorporate the situatedness of cognitive systems.

6.1.4 The symbol grounding problem

The fourth advantage of the operationalisation of situated representation over the operationalisation of non-situated representation is that the representation accounted for by the new operationalisation is grounded in the interaction with the world and/or in the internal simulation of that interaction. In what follows, we will describe this advantage.

The operationalisation of non-situated representation yields the symbol grounding problem (Harnad, 1990). The symbol grounding problem is the problem of how symbols get their meaning. Harnad (1990, p.335) summarises the problem as follows: “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?”. The operationalisation of non-situated representation yields the symbol grounding problem by allowing symbolic representation that is not grounded. The operationalisation of situated representation allows internal simulation to be interpreted as a situated variant of symbol manipulation (see chapter 4). This situated variant of symbol manipulation is not affected by the symbol grounding problem, because the manipulated symbols are grounded in the coordination of the sensory and motor system of the cognitive system (cf., Thelen and Smith (1994)). We illustrate this by means of the sToL model described in chapter 4. In the sToL model, sensory states are grounded in the interaction of the agent with the ToL, because the agent can observe the consequences of those sensory states and their corresponding actions (cf., Pfeifer and Scheier (1997; 1999)). Similarly, the situated variant of symbol manipulation of agents in the sToL model, i.e., the expected state generated by the internal simulation mechanism, is grounded, because an agent can observe the consequence of an action that is co-determined by the expected state. Therefore, situated representations are not affected by the symbol grounding problem.

6.1.5 The relation between the operationalisations

So far, in this section, we have focused on the advantages of the operationalisation of situated representation over that of non-situated representation. However, the operationalisation of non-situated representation is not without its merits (see chapter 1), and is still popular in the cognitive sciences (see, e.g., Markman and Dietrich (2000a; 2000b), Markman and Brendl (2005)). Moreover, the operationalisation of non-situated representation is popular in the cognitive sciences because it explains cognition in terms of symbol manipulation (Newell and Simon, 1972). The operationalisation of situated representation allows symbol manipulation by internal simulation of perception and behaviour, which can occur in the absence of actual interaction with the world.

The new operationalisation dismisses the claim of the physical symbol-system hypothesis (Newell and Simon, 1985) that a “physical symbol system [...] has the necessary and sufficient means for general intelligent action” (Simon, 1969, p.23, see also chapter 1), if the symbols in that symbol system are not learned through interaction with the world. Situated representations should not be ‘hand coded’, but should emerge and change through interaction with the world (Morrison, 1998). The operationalisation of situated representation claims that the physical symbol-system hypothesis is tenable for symbol systems whose symbolic representations are learned through interaction with the world. By generating behaviour through mere internal simulation of perception and behaviour, situated representations can be put to use in the absence of actual interaction with the environment, as classic symbol systems do.

Restricting symbol systems to the use of situated symbolic representation may make it harder to formalise the symbol manipulation they perform. However, at the same time, the level to which these systems are biologically and psychologically plausible increases. We believe that the operationalisation of non-situated representation should be replaced by the new operationalisation, because the latter holds on to the strong point of the first (symbol manipulation), while making progress regarding its weaknesses (discussed above).

6.2 Implications of the new operationalisation

The operationalisation of situated representation formulated at the beginning of this chapter, may have implications for the fields of artificial intelligence, cognitive neuroscience, and cognitive psychology. These implications are discussed in subsections 6.2.1, 6.2.2, and 6.2.3, respectively.

6.2.1 Artificial Intelligence

Since representation is a central concept in Artificial Intelligence (AI) (see chapter 1), the new operationalisation may have implications for the field of AI. First, we will discuss whether such implications exist concerning the goal of AI. Second, we

will discuss whether such implications exist concerning the distinction in AI between symbol systems and situated systems.

The goal of AI is to create intelligent machines (McCarthy, 2004), i.e., machines that can think. Following Turing (1950), the question “Can machines think?” should be replaced by the question whether machines can behave in a way that is indistinguishable from that of humans. Symbol systems are regarded the best means to this end (Newell and Simon, 1972). However, the Turing approach to the creation of intelligent machines has been widely criticised on the ground that machines can generate human-like behaviour without having any understanding of what they are doing (Searle, 1980). This criticism resulted in another route to the creation of intelligent machines; creating machines that understand what they do. This proved to be an objective that is hard to reach within the paradigm of the physical symbol-system hypothesis, because of the symbol grounding problem (Harnad, 1989; 1990). Subsection 6.1.4 pointed out that, by adhering to the operationalisation of situated representation, representations are grounded in the interaction with the world. So, the new operationalisation enables AI to create machines that understand what they do. By doing so, the new operationalisation may have the implication that AI can achieve its goal; create intelligent machines.

In AI, a distinction exists between symbol systems (see section 1.1) and situated systems (see section 1.2). Symbol systems have been shown to perform complex abstract planning problems relatively well (see, e.g., Hsu (2004)), but have also been shown to perform poorly on problems in the real world (see, e.g., Nilsson (1984)). In contrast, situated systems have been shown to be able to cope with the complex dynamics of the real world (see, e.g., Arkin (1998)), but are known to perform poorly on tasks that require much planning (Matarić, 2002). The solution seems to lie in a hybrid approach that combines symbols and situatedness in a single system. Although such hybrid systems already exist in AI (see Arkin (1998), for an overview), the symbols and situatedness of such systems are modelled in separate modules of the system, which might explain why they have not been very successful to date. In contrast, the operationalisation of situated representation suggests a new type of hybridization in which the symbols themselves are situated (see subsection 6.1.5). By doing so, the new operationalisation may have the implication for AI that it enables the creation of systems that combine the ability of situated systems to cope with the complex dynamics of the real world, with the ability of symbol systems to perform complex abstract planning problems relatively well.

6.2.2 Cognitive neuroscience

Cognitive neuroscience studies the relation between the brain and the mind by recording and analysing brain activity. It assumes that input from the world is represented in the brain in such a way that it can be extracted by measuring activity in the brain. For instance, O’Toole *et al.* (2005) successfully located brain activities that correlate with the presence of various types of objects, such as faces, houses, cats, chairs, shoes, scissors, and bottles. Brain activity that correlates with the

presence of an object is regarded by O’Toole *et al.* (2005) to *represent* an object. Attributing symbolic meaning to activity in the brain in this way is common in cognitive neuroscience, and is in accordance with the operationalisation of non-situated representation. Below, we will discuss whether the operationalisation of situated representation has an implication for the field of cognitive neuroscience, regarding the attribution of symbolic meaning to activity in the brain.

Although cognitive neuroscience has provided insight into the relation between activity in the brain and cognitive events, the field is criticised for not doing anything more than that; locating activity in the brain that correlates with cognitive events (Revonsuo, 2001). The attribution of symbolic meaning to brain activity by cognitive neuroscience results in the symbol grounding problem (see subsection 6.1.4). An illustration of the symbol grounding problem in cognitive neuroscience is the ‘binding problem’ (Roskies, 1999). According to Roskies (1999, p.7), the canonical example of the binding problem is that “one sort of visual feature, such as an object’s shape, must be correctly associated with another feature, such as its location, to provide a unified representation of that object”. A unified representation is required because the operationalisation of non-situated representation that is adhered to in cognitive neuroscience requires one. The operationalisation of non-situated representation requires a unified representation, because it holds that one entity (in this case the activity in the brain) stands in for another entity (in this case the object). In the operationalisation of situated representation, such a unified representation is not a prerequisite for experience. In the new operationalisation, brain activity has the mere role of mediating the interaction with the world, disregarding its spatial distribution over the brain. Although activity in the brain has to lead up to convergence at the side of the motor system, it does not have to be ‘bounded’. Since the operationalisation of situated representation does not require any ‘binding’ of activity in the brain, adhering to the new operationalisation has the implication that it frees the field of cognitive neuroscience from the binding problem.

6.2.3 Cognitive psychology

As in AI, representation is also a central concept in cognitive psychology (Markman, 1998). In general, cognitive psychological models adhere to the operationalisation of non-situated representation. Therefore, the new operationalisation may have implications for the fundamental notions of cognitive psychology, i.e., perception, action, memory, and cognition. In what follows, we focus on whether such implications exist for models of memory.

The vast majority of cognitive models, treat memory as an open-loop system (see figure 2.1 in section 2.1), a container that stores detailed and complete information that can be retrieved at a later time (see Conway (2003), for an overview). However, the new operationalisation stresses the fact that cognitive systems are closed loop (see figure 2.2 in section 2.1), implying that action and perception are an integral part of memory. The radically different conception of memory implied by the new operationalisation gives rise to the development of models of memory that are situ-

ated. These models differ from existing models in two main ways. First, they take realistic stimuli such as natural images as input (see, e.g., Lacroix *et al.* (2004)), rather than simplified abstractions of these stimuli. Second, they translate current and past sensory inputs into actions, the effects of which affect the input (see, e.g., de Croon, Postma, and van den Herik (2005a)). As a consequence, memory can only be understood in terms of this translation rather than in terms of a store for detailed and complete representations. Evidently, the same applies to other faculties such as perception and action, and to cognition in general. Therefore, the new operationalisation has the implication that it may translate into novel hypotheses and theories of cognition in cognitive psychology.

Chapter 7

Conclusion

In section 1.3 we formulated the problem statement: *What is the nature of representation in situated systems?* This problem statement was translated into two research questions.

(i) *To what extent can we identify where the knowledge resides that is used by a situated system to perform a certain task?*

and

(ii) *How is this knowledge accessed and used by a situated system when performing a certain task?*

In what follows, we will answer these two research questions and conclude on the problem statement, respectively.

First, the investigations reported in the chapters 3, 4, and 5, show that we can identify where the knowledge resides that is used by a situated system to perform a certain task to the extent that we can reveal the coordination between the sensory and motor system(s) of a system and relate it to the environmental dynamics. If this coordination can be revealed and related to the environmental dynamics, then the knowledge that is used by a situated system to perform a certain task can be analysed. Hence, its situated representations can be studied. We found knowledge that is used by a situated system to perform a certain task to reside in (i) the attractors in the interaction with the environment in chapter 3, (ii) the internal simulation of interaction with the environment in chapter 4, and (iii) the average properties of the interaction with the environment in chapter 5.

Second, the investigations reported in chapters 3, 4, and 5, show that the knowledge which a situated system uses to perform a certain task is accessed and used by performing and/or simulating internally the specific interaction with the environment by which the knowledge is represented. Moreover, we found that the knowledge

which a situated system uses to perform a certain task is accessed and used by:

- (i) exploiting the attractors in the interaction with the environment (chapter 3),
- (ii) simulating interaction with the environment internally (chapter 4), and
- (iii) exploiting the average properties of the interaction with the environment (chapter 5).

In the first case, agents in the active categorical perception model evolve attractors in the interaction with the environment during evolution, that result in adequate categorisation when the interaction is performed. In the second case, the agent in the situated Tower of London model evolves interaction with the environment and internal simulation of such interaction, which is combined to solve Tower of London problems adequately. In the third case, the best-performing agents in the foraging model evolve interaction with the environment, that completely depends on the average properties of the interaction with the environment to perform foraging adequately.

From these findings we may conclude that in emphasising the role of interaction for cognition in the theory of situated cognition (Clancey, 1997) the operationalisation of situated representation is essential. In a situated system, representation is as strongly rooted in the environment as the system itself. Therefore, the answer to the problem statement formulated in section 1.3 is that *representation is situated in nature*.

Finally, we demonstrate the power of the operationalisation of situated representation by returning to the ‘sign pole’ example stated in chapter 1. In the ‘sign pole’ example we stated that, according to Newell (1980), a sign pole is a representation of a speed limit, because a sign pole designates a speed limit relative to the process of slowing down to the speed indicated by the sign pole. Furthermore, we stated that if the sign pole in the example were replaced by a rule in someone’s mind concerning maximum speed learned in a traffic school, this rule would be a *mental* representation of the speed limit. According to the operationalisation of situated representation a speed limit is adequately represented when someone is able to perform and/or simulate internally the interaction with the environment that is specific for the speed limit. This specific interaction may consist of (i) the interaction with the environment, e.g., pushing and releasing the gas pedal of the car according to the change in speedometer reading, (ii) the interaction with the environment that is simulated internally, e.g., anticipation on the consequences of speeding, and (iii) the average properties of the interaction with the environment, e.g., the chance of having an accident or of getting fined for speeding.

References

- Agre, P. E. (1995). Computational research on interaction and agency. *Artificial Intelligence*, Vol. 72, pp. 1–52. [4]
- Aloimonos, J. Y., Weiss, I., and Bandopadhyay, A. (1988). Active vision. *International Journal of Computer Vision*, Vol. 2, pp. 333–356. [22]
- Arkin, R. C. (1998). *Behavior-Based Robotics*. MIT Press, Cambridge, MA. [93]
- Bach-y-Rita, P. (1995). *Nonsynaptic Diffusion Neurotransmission and Late Brain Reorganization*. Demos Publications, New York, NY. [4]
- Baddeley, A. D. (1986). *Working Memory*. Oxford University Press, Oxford, NY. [42]
- Bajcsy, R. (1988). Active perception. *Proceedings of the IEEE, Special issue on Computer Vision* (eds. H. Li and J. R. Kender), Vol. 76 (8), pp. 996–1005. [22]
- Bakker, B. and van der Voort van der Kleij, G. (2000). Trading off perception with internal state: Reinforcement learning and analysis of Q-Elman networks in a Markovian task. *Proceedings of the International Joint Conference on Neural Networks* (eds. C. S.-I. Amari, L. Giles, M. Gori, and V. Piuri), Vol. 3, pp. 213–218. [34]
- Bakker, B. (2004). *The State of Mind: Reinforcement Learning with Recurrent Neural Networks*. Ph.D. thesis, Leiden University, Leiden, NL. [59]
- Balkenius, C. (1995). *Natural Intelligence in Artificial Creatures*. Ph.D. thesis, Lund University Cognitive Studies, Lund, SE. [14, 19]
- Bechtel, W. (1998). Representations and cognitive explanations: Assessing the dynamicist’s challenge in cognitive science. *Cognitive Science*, Vol. 22(3), pp. 295–318. [5]
- Beekman, M., Sumpter, D. J. T., and Ratnieks, F. L. W. (2001). Phase transitions between disordered and ordered foraging in Pharaoh’s ants. *Proceedings of the National Academy of Sciences*, Vol. 98, pp. 9703–9706. [70]

- Beer, R. D. (1990). *Intelligence as Adaptive Behavior: An Experiment in Computational Neuroethology*. Academic Press, San Diego, CA. [4, 7]
- Beer, R. D. (1995). Computational and dynamical languages for autonomous agents. *Mind as Motion: Dynamics, Behavior, and Cognition*. (eds. R. Port and T. J. van Gelder), Chapter 5, pp. 121–148. MIT Press, Cambridge, MA. [4, 7, 69]
- Beer, R. D. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. *From Animals to Animats 4. Proceedings of the 4th International Conference on Simulation of Adaptive Behavior* (eds. P. Maes, M. J. Matarić, J.-A. Meyer, J. Pollack, and S. Wilson), Vol. 4, pp. 421–429. [4, 7, 16, 19, 21, 22, 34, 38, 40]
- Beer, R. D. (2000). Dynamical approaches to cognitive science. *Trends in Cognitive Sciences*, Vol. 4, pp. 91–99. [4, 7, 69]
- Beer, R. D. (2003a). Arches and stones in cognitive architecture. *Adaptive Behavior*, Vol. 11(4), pp. 299–305. [90]
- Beer, R. D. (2003b). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, Vol. 11(4), pp. 209–243. [2, 4, 7, 8, 16, 17, 19, 20, 21, 22, 39, 69, 90]
- Berg, W. K. and Byrd, D. L. (2002). The Tower of London spatial problem-solving task: Enhancing clinical and research implementation. *Journal of Clinical and Experimental Neuropsychology*, Vol. 24, pp. 586–604. [57]
- Bickhard, M. H. and Richie, D. M. (1983). *On the Nature of Representation: A Case Study of James Gibson's Theory of Perception*. Praeger Publishers, New York, NY. [4]
- Bickhard, M. H. (1997). The emergence of representation in autonomous agents. *Epistemological Issues of Embodied AI* (ed. E. Prem), Vol. 28(6), pp. 489–498. Taylor and Francis, London, UK. [5]
- Bickhard, M. H. (1999). Interaction and representation. *Theory and Psychology*, Vol. 9(4), pp. 435–458. [87]
- Bickhard, M. H. (2004). The dynamic emergence of representation. *Representation in Mind: New Approaches to Mental Representation* (eds. H. Clapin, P. Staines, and P. Slezak), pp. 71–90. Elsevier, New York, NY. [4, 5, 87]
- Boden, M. (1990). *The Creative Mind: Myths and Mechanisms*. Weidenfeld and Nicolson, London, UK. [3]
- Brachman, R. J., Levesque, H. J., and Reiter, R. (1991). Introduction to the special volume on knowledge representation. *Knowledge Representation* (eds. R. J. Brachman, H. J. Levesque, and R. Reiter), pp. 1–3. MIT Press, Cambridge, MA. [1]

- Braitenberg, V. (1986). *Vehicles: Experiments in Synthetic Psychology*. MIT Press, Cambridge, MA. [13, 14]
- Brookes, J. (2001). Experience, attention, and mental representation. *Behavioural and Brain Sciences*, Vol. 24(5), pp. 978–979. (A commentary on: A sensorimotor account of vision and visual consciousness. by J.K. ORegan and A. Noë). [90]
- Brooks, R. A. (1986). A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, Vol. 2, pp. 14–23. [81]
- Brooks, R. A. (1990). Elephants Don't Play Chess. *Robotics and Autonomous Systems*, Vol. 6(1-2), pp. 3–15. [3, 5, 88]
- Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, Vol. 47(1-3), pp. 139–159. [1, 5, 13, 88]
- Brower, J. V. (1958). Experimental studies of mimicry in some North American butterflies. *Evolution*, Vol. 12, pp. 32–47. [20]
- Bull, R., Espy, K. A., and Senn, T. E. (2004). A comparison of performance on the Towers of London and Hanoi in young children. *Journal of Child Psychology and Psychiatry*, Vol. 45(4), pp. 743–754. [44]
- Bullock, S. (2003). Making room for representation. *Adaptive Behavior*, Vol. 11(4), pp. 279–280. (A commentary on: The dynamics of active categorical perception in an evolved model agent. by R. D. Beer). [90]
- Chambers, D. and Reisberg, D. (1985). Can mental images be ambiguous? *Journal of Experimental Psychology: Human Perception and Performance*, Vol. 11, pp. 317–328. [15, 20]
- Chase, W. G. and Simon, H. A. (1973a). The mind's eye in chess. *Visual Information Processing* (ed. W. G. Chase). Academic Press, New York, NY. [3]
- Chase, W. G. and Simon, H. A. (1973b). Perception in chess. *Cognitive Psychology*, Vol. 4, pp. 55–81. [3]
- Chrisley, R. and Ziemke, T. (2002). Embodiment. *Encyclopedia of Cognitive Science* (ed. L. Nadel). Nature Publishing Group / Macmillan Reference, London, UK. [6]
- Churchland, P. S. and Sejnowski, T. (1992). *The Computational Brain*. MIT Press, Cambridge, MA. [3]
- Clancey, W. J. (1995). A boy scout, Toto, and a bird: How situated cognition is different from situated robotics. *The "Artificial Life" Route to "Artificial Intelligence": Building Situated Embodied Agents* (eds. L. Steels and R. Brooks), pp. 227–236. Lawrence Erlbaum Associates, Mahwah, NJ. [11]

- Clancey, W. J. (1997). *Situated Cognition: On Human Knowledge and Computer Representations*. Cambridge University Press, New York, NY. [1, 3, 4, 5, 6, 22, 98]
- Clark, A. (1997). *Being There: Putting Brain, Body, and World Together Again*. MIT Press, Cambridge, MA. [1, 5]
- Clark, A. (2001). *Mindware: An Introduction to the Philosophy of Cognitive Science*. Oxford University Press, New York, NY. [12]
- Cliff, D. and Noble, J. (1997). Knowledge-based vision and simple visual machines. *Philosophical Transactions of the Royal Society: Biological Sciences*, Vol. 352(1358), pp. 1165–1175. [5]
- Cliff, D., Harvey, I., and Husbands, P. (1993). Explorations in evolutionary robotics. *Adaptive Behavior*, Vol. 2(1), pp. 71–104. [7]
- Cohen, P. R. (1995). *Empirical Methods for Artificial Intelligence*. MIT Press, Cambridge, MA. [79, 80, 81, 82]
- Cohen, J. (2001). Whither visual representations? Whither qualia? *Behavioural and Brain Sciences*, Vol. 24(5), pp. 980–981. (A commentary on: A sensorimotor account of vision and visual consciousness. by J.K. ORegan and A. Noë). [90]
- Connell, J. (1989). *A Colony Architecture for an Artificial Creature*. Ph.D. thesis, MIT AI Lab, Cambridge, MA. [13]
- Conway, M. A. (ed.) (2003). *Cognitive Models of Memory*. MIT Press, Cambridge, MA. [94]
- Cooper, A. N. and Shepard, R. N. (1973). The time required to prepare for a rotated stimulus. *Memory and Cognition*, Vol. 1, pp. 246–250. [15]
- Cooper, R. and Waldau, R. C. (submitted). Development of problem solving strategies: Strategy shift in the Tower of London task. *Developmental Science*. Retrieved: 7 December 2004, from COGENT Online: <http://cogent.psyc.bbk.ac.uk/publications/cooperwaldau99.html>. [44]
- Cooper, R. (2002). *Modelling High-Level Cognitive Processes*. Lawrence Erlbaum Associates, Mahwah, NJ. [8, 42]
- Crick, F. and Mitchison, G. (1983). The function of dream sleep. *Nature*, Vol. 304(14), pp. 111–114. [3]
- Cruse, H. (2003). The evolution of cognition - a hypothesis. *Cognitive Science*, Vol. 27(1), pp. 135–155. [14]
- Cummins, R. (1989). *Meaning and Mental Representation*. MIT Press, Cambridge, MA. [5]

- Cummins, R. (1996). *Representations, Targets and Attitudes*. MIT Press, Cambridge, MA. [5]
- Dagher, A., Owen, A. M., Boecker, H., and Brooks, D. J. (1999). Mapping the network for planning: A correlation PET activation study with the Tower of London task. *Brain*, Vol. 122, pp. 1973–1987. [43]
- Dartnall, T. (2005). Does the world leak into the mind? Active externalism, 'internalism' and epistemology. *Cognitive Science*, Vol. 29(1), pp. 135–143. [87]
- Davis, R., Shrobe, H. E., and Szolovits, P. (1993). What is a knowledge representation? *AI Magazine*, Vol. 14(1), pp. 17–33. [1]
- de Croon, G., Postma, E. O., and van den Herik, H. J. (2005a). Sensory-motor coordination in gaze control. *Proceedings of the 7th European Workshop on Evolutionary Computation in Image Analysis and Signal Processing* (eds. F. Rothlauf, J. Branke, S. Cagnoni, D. W. Corne, R. Drechsler, Y. Jin, P. Machado, E. Marchiori, J. Romero, G. D. Smith, and G. Squillero). Springer-Verlag, New York, NY. Springer Lecture Notes in Computer Science Vol. 3449. [95]
- de Croon, G., van Dartel, M. F., and Postma, E. O. (2005b). Evolutionary learning outperforms reinforcement learning on non-Markovian tasks. *Workshop on Memory and Learning Mechanisms in Autonomous Robots* (eds. E. Tuci and D. Kim). [29, 39]
- de Croon, G., Nolfi, S., and Postma, E. O. (in press). Toward pro-active embodied agents: On the importance of neural mechanisms suitable to process information in time. *Complex Engineering Systems* (eds. D. Braha, A. Minai, and Y. Bar-Yam). Perseus Books Group, Boulder, CO. [70, 81]
- de Groot, A. D. (1946). *Het Denken van den Schaker*. Noord-Hollandsche Uitgevers Maatschappij, Amsterdam, NL. [3]
- de Groot, A. D. (1965). *Thought and Choice in Chess*. Mouton, Den Haag, NL. [3]
- de Pinedo, M. and Noble, J. (2003). Mind/body problems? Turn to Beer. *Adaptive Behavior*, Vol. 11(4), pp. 289–290. (A commentary on: The dynamics of active categorical perception in an evolved model agent. by R. D. Beer). [90]
- Dehaene, S. and Changeux, J. P. (1997). A hierarchical neuronal network for planning behavior. *Proceedings of the National Academy of Sciences*, Vol. 94, pp. 13293–13298. [8, 43, 65, 129]
- Dreyfus, H. L. (2002a). A phenomenology of skill acquisition as the basis for a Merleau-Pontian non-representationalist cognitive science. Technical report, University of California, Berkeley. Retrieved 20 April 2005, from homepage: <http://ist-socrates.berkeley.edu/~hdreyfus/pdf/MerleauPontySkillCogSci.pdf>. [3]

- Dreyfus, H. L. (2002b). Intelligence without representation: Merleau-Ponty's critique of mental representation. *Phenomenology and Cognitive Science*, Vol. 1(4), pp. 367–383. [3]
- Easton, R. D. (1992). Inherent problems of attempts to apply sonar and vibrotactile sensory aid technology to the perceptual needs of the blind. *Optometry and Vision Science*, Vol. 69, pp. 3–14. [4]
- Edelman, S. (1998). Representation is representation of similarities. *Behavioral and Brain Sciences*, Vol. 21, pp. 449–498. [89]
- Edelman, S. (2003). But will it scale up? Not without representations. *Adaptive Behavior*, Vol. 11(4), pp. 273–275. (A commentary on: The dynamics of active categorical perception in an evolved model agent. by R. D. Beer). [90]
- Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, Vol. 14, pp. 179–211. [24, 50, 59]
- Field, H. H. (1978). Mental representation. *Erkenntnis*, Vol. 13(1), pp. 9–61. [1]
- Field, H. H. (1994). Mental representation. *Mental Representation: A Reader* (eds. S. P. Stich and T. A. Warfield). Basil Blackwell, Oxford, UK. [1]
- Fodor, J. A. and McLaughlin, B. (1990). Connectionism and the problem of systematicity: Why Smolensky's solution doesn't work. *Cognition*, Vol. 35, pp. 183–204. [63]
- Fodor, J. A. (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. MIT Press, Cambridge, MA. [1]
- Fodor, J. A. (2000). *The Mind Doesn't Work That Way: The Scope and Limits of Computational Psychology*. MIT Press, Cambridge, MA. [3]
- Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Houghton-Mifflin, Boston, MA. [4, 88]
- Goldberg, D. E. (1986). *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley Publishing Company, Reading, MI. [29, 30, 52, 73]
- Goldenberg, E., Garcowski, J., and Beer, R. D. (2004). May we have your attention: Analysis of a selective attention task. *From Animals to Animals 8. Proceedings of the 8th International Conference on the Simulation of Adaptive Behaviour* (eds. S. Schaal, A. Ijspeert, A. Billard, S. Vijajamumar, J. Hallam, and J.-A. Meyer). MIT Press, Cambridge, MA. [7]
- Goodwin, C. J. (1995). *Research in psychology: Methods and Design*. Wiley, New York, NY. [6]

- Gordon, I. E. (1989). *Theories of Visual Perception*. Wiley and Sons, Chichester, UK. [88]
- Gray, W. D. and Fu, W. (2004). Soft constraints in interactive behavior: The case of ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head. *Cognitive Science*, Vol. 28, pp. 359–382. [20]
- Grossberg, S. (ed.) (1987). *The Adaptive Brain*, Vol. 1-2. Elsevier Science, Amsterdam, NL. [123]
- Grush, R. (2004). The emulation theory of representation: Motor control, imagery, and perception. *Behavioural and Brain Sciences*, Vol. 27, pp. 377–442. [42]
- Guillot, A. and Meyer, J.-A. (2001). The animat contribution to cognitive systems research. *Cognitive Systems Research*, Vol. 2, pp. 157–165. [70]
- Gutowski, M. (2001). Lévy flights as an underlying mechanism for global optimization algorithms. *Math-ph/0106003*. [72, 74]
- Harnad, S. (1987). *Categorical Perception: The Groundwork of Cognition*. Cambridge University Press, New York, NY. [3, 8, 17, 39]
- Harnad, S. (1989). Minds, machines and searle. *Journal of Theoretical and Experimental Artificial Intelligence*, Vol. 1, pp. 5–25. [93]
- Harnad, S. (1990). The symbol grounding problem. *Physica D*, Vol. 42, pp. 335–346. [91, 93]
- Harnad, S. (2003). Cognition is categorization. Retrieved 21 July 2004, from Cogprints Web site: <http://cogprints.ecs.soton.ac.uk/archive/00003027/01-catconf.html>. [39]
- Hartcup, G. (1979). *Camouflage: A History of Concealment and Deception in War*. David and Charles, Newton Abbot, Devon. [20]
- Harvey, I., Husbands, P., Cliff, D., Thompson, A., and Jakobi, N. (1997). Evolutionary robotics: The sussex approach. *Robotics and Autonomous Systems*, Vol. 20, pp. 205–224. [7]
- Haselager, W. F. G., de Groot, A. D., and van Rappard, J. F. H. (2003). Representationalism versus anti-representationalism: A debate for the sake of appearance. *Philosophical Psychology*, Vol. 16(1), pp. 5–23. [5, 89]
- Havlin, S., Buldyrev, S. V., Bunde, A., Goldberger, A. L., Ivanov, P., Peng, C.-K., and Stanley, H. E. (1999). Scaling in nature: From DNA through heartbeats to weather. *Physica A*, Vol. 273, pp. 46–69. [70]
- Hesslow, G. (2002). Conscious thought as simulation of behaviour and perception. *Trends in Cognitive Sciences*, Vol. 6, pp. 242–247. [13, 42]

- Hinton, G. E., McClelland, J. L., and Rumelhart, D. E. (1986). Distributed Representations. *Parallel Distributed Processing* (eds. Rumelhart and McClelland), Vol. 1. MIT Press, Cambridge, MA. [49]
- Hoffmann, H. and Möller, R. (2004). Action selection and mental transformation based on a chain of forward models. *Proceedings of the 8th Conference on Simulation of Adaptive Behavior* (eds. S. Schaal, A. Ijspeert, A. Billard, S. Vijayakumar, J. Hallam, and J.-A. Meyer), pp. 213–222. MIT Press, Cambridge, MA. [42, 66]
- Horst, S. (2003). The computational theory of mind. *Stanford Encyclopedia of Philosophy* (ed. E. N. Zalta). Stanford University, Stanford, CA. <http://plato.stanford.edu>. [3]
- Hsu, F.-H. (2004). *Behind Deep Blue: Building the Computer that Defeated the World Chess Champion*. Princeton University Press, Princeton, NJ. [93]
- Hutchins, E. (1995). How a cockpit remembers its speed. *Cognitive Science*, Vol. 19, pp. 265–288. [20]
- Izquierdo-Torres, E. and Di Paolo, E. (2005). Is an embodied system ever purely reactive? *Proceedings of the 8th European Conference on Artificial Life (ECAL 2005)* (eds. M. Capcarrere, A. A. Freitas, P. J. Bentley, C. G. Johnson, and J. Timmis), pp. 252–261. Springer, Berlin, DE. [39]
- Jeannerod, M. (1994). The representing brain: Neural correlates of motor intention and imagery. *Behavioral and Brain Sciences*, Vol. 17, pp. 187–245. [42]
- Jirenhed, D.-A., Hesslow, G., and Ziemke, T. (2001). Exploring internal simulation of perception in mobile robots. *Proceedings of the 4th European Workshop on Advanced Mobile Robots* (eds. K. O. Arras, C. Balkenius, A. Baerfeldt, W. Burgard, and R. Siegwart), pp. 107–113. Lund, Sweden. [50, 66]
- Jones, C. and Matarić, M. J. (2003). Adaptive division of labor in large-scale multi-robot systems. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1969–1974. [70]
- Juarrero, A. (1999). *Dynamics in Action: Intentional Behavior as a Complex System*. MIT Press, Cambridge, MA. [4]
- Kay, L. (1984). Electronic aids for blind persons: an interdisciplinary subject. *IEEE Proceedings, part A*, Vol. 131. [4]
- Keijzer, F. A. (2001). *Representation and Behavior*. MIT Press, Cambridge, MA. [5]
- Keijzer, F. A. (2002). Representation in dynamical and embodied cognition. *Cognitive Systems Research*, Vol. 3, pp. 275–288. [5]

- Kirsh, D. and Maglio, P. (1992). Reaction and reflection in Tetris. *Artificial Intelligence planning systems: Proceedings of the 1st Annual International Conference on AI Planning Systems (AIPS92)* (ed. J. Hendler). Morgan Kaufmann, San Mateo, CA. [20]
- Kirsh, D. (1995). Complementary strategies: Why we use our hands when we think. *Proceedings of the 17th Annual Conference of the Cognitive Science Society* (eds. J. D. Moore and J. F. Lehman), Lawrence Erlbaum, Hillsdale, NJ. [20]
- Kolb, B. and Wishaw, I. (1983). Performance of schizophrenic patients on tests sensitive to left or right frontal temporal or parietal function in neurological patients. *Journal of Nervous and Mental Disease*, Vol. 171, pp. 435–443. [42, 44]
- Kortmann, L. J. (2003). *The Resolution of Visually Guided Behaviour*. Ph.D. thesis, Universiteit Maastricht, Maastricht, NL. [4]
- Kosslyn, S. M., Ganis, G., and Thompson, W. L. (2001). Neural foundations of imagery. *Nature Reviews Neuroscience*, Vol. 2, pp. 635–642. [42]
- Lacroix, J. P. W., Murre, J. M. J., Postma, E. O., and Herik, H. J. van den (2004). The natural input memory model. *Proceedings of the Cognitive Science Conference (CogSci 2004)*, pp. 773–778. [20, 89, 95]
- Lacroix, J. P. W., Postma, E. O., and Murre, J. M. J. (in press). Modeling recognition memory using the similarity structure of natural input. *Cognitive Science*. [89]
- Laird, J. E., Newell, A., and Rosenbloom, P. S. (1987). SOAR: An architecture for general intelligence. *Artificial Intelligence*, Vol. 33(3), pp. 1–64. [3]
- Langton, C. G. (ed.) (1995). *Artificial Life: An Overview*. MIT Press, Cambridge, MA. [9]
- Lazeron, R. H., Rombouts, S. A., Machielsen, W. C., Scheltens, P., Witter, M. P., Uylings, H. B., and Barkhof, F. (2000). Visualizing brain activation during planning: The tower of London test adapted for functional MR imaging. *AJNR American Journal of Neuroradiology*, Vol. 21(8), pp. 1407–1414. [43]
- Lerman, K. and Galstyan, A. (2003). Macroscopic analysis of adaptive task allocation in robots. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1951–1956. [70]
- Lerman, K., Galstyan, A., Martinoli, A., and Ijspeert, A. J. (2001). A macroscopic analytical model of collaboration in distributed robotic systems. *Artificial Life*, Vol. 7(4), pp. 375–393. [70]

- Lorenz, K. (1973). *Foundations of Ethology*. Springer-Verlag, New York, NY. [19]
- Lucas, É. (1882-3). *Récréations Mathématiques*. Gauthier-Villars, Paris. [44]
- Luria, A. R. (1968). *The Mind of a Mnemonist*. Harvard University Press, Cambridge, MA. [14, 15]
- Markman, A. B. and Brendl, C. M. (2005). Constraining theories of embodied cognition. *Psychological Science*, Vol. 16(1), pp. 6–10. [90, 92]
- Markman, A. B. and Dietrich, E. (2000a). Extending the classical view of representation. *Trends in Cognitive Sciences*, Vol. 4, pp. 70–75. [5, 92]
- Markman, A. B. and Dietrich, E. (2000b). In defense of representation. *Cognitive Psychology*, Vol. 40, pp. 138–171. [1, 5, 90, 92]
- Markman, A. B. (1998). *Knowledge Representation*. Lawrence Erlbaum Associates, Mahwah, NJ. [94]
- Marr, D. (1982). *Vision*. W. H. Freeman and Sons, New York, NY. [88]
- Matarić, M. J. (1998). Behavior-based robotics as a tool for synthesis of artificial behavior and analysis of natural behavior. *Trends in Cognitive Science*, Vol. 2(3), pp. 82–87. [41]
- Matarić, M. J. (2002). Situated robotics. *Encyclopedia of Cognitive Science* (ed. L. Nadel). Nature Publishing Group / Macmillan Reference, London, UK. [93]
- Matlin, M. W. (1994). *Cognition*. Harcourt Brace Publishers, New York, NY, 3 edition. [6]
- Maturana, H. R. and Varela, F. J. (1980). *Autopoiesis and Cognition: The Realization of the Living*. D. Reidel, Dordrecht, NL. [4]
- McCarthy, J. (2004). What is Artificial Intelligence. Retrieved 12 March 2005, from homepage: <http://www-formal.stanford.edu/jmc>. [93]
- Miller, B. L. and Goldberg, D. E. (1995). Genetic Algorithms, Tournament Selection, and the Effects of Noise. *Complex Systems*, Vol. 9, pp. 193–212. [29]
- Mobus, G. E. (1999). Foraging search: Prototypical intelligence. *Proceedings of the 3rd International Conference on Computing Anticipatory Systems (CASYS'99)* (ed. D. Dubois), Vol. 517(1), pp. 592–605. American Institute of Physics, Liege, BE. [8]
- Morén, J. and Balkenius, C. (2000). A computational model of emotional learning in the amygdala. *From Animals to Animats 6. Proceedings of the 6th International Conference on the Simulation of Adaptive Behaviour* (eds. J.-A. Meyer, A. Berthoz, D. Floreano, H. L. Roitblat, and S. W. Wilson). MIT Press, Cambridge, MA. [3]

- Morris, R. G., Ahmed, S., Syed, G. M., and Toone, B. K. (1993). Neural correlates of planning ability: Frontal lobe activation during the Tower of London test. *Neuropsychologia*, Vol. 31, pp. 1367–1378. [43]
- Morrison, C. T. (1998). *Situated Representation: Solving the Handcoding Problem with Emergent Structured Representation*. Ph.D. thesis, Binghamton University, Binghamton, NY. [5, 11, 13, 92]
- Newell, A. and Simon, H. A. (1972). *Human Problem Solving*. Prentice-Hall, Englewood Cliffs, NJ. [11, 41, 42, 62, 92, 93]
- Newell, A. and Simon, H. A. (1976). Computer science as empirical enquiry: Symbols and search. *Communications of the Association for Computing Machinery*, Vol. 19, pp. 113–126. [3, 4, 5]
- Newell, A. and Simon, H. A. (1985). Computer science as empirical inquiry: Symbols and search. *Mind Design* (ed. J. Haugeland), pp. 35–66. MIT Press, Cambridge, MA. [3, 4, 5, 41, 92]
- Newell, A. (1980). Physical symbol systems. *Cognitive Science*, Vol. 4(4), pp. 135–183. [1, 3, 5, 87, 98]
- Nilsson, N. J. (1984). Shakey the robot. *Technical Report 323, AI Center*. SRI International, Menlo Park, CA. [3, 12, 93]
- Noë, A. and Thompson, E. (2002). Introduction to Vision and Mind. *Vision and Mind: Selected Readings in the Philosophy of Perception* (eds. A. Noë and E. Thompson), Chapter 1. MIT Press, Cambridge, MA. [14]
- Noë, A., Pessoa, L., and Thompson, E. (2000). Beyond the grand illusion: What change blindness really teaches us about vision. *Visual Cognition*, Vol. 7, pp. 93–106. [90]
- Noë, A. (2004). *Action in Perception*. MIT Press, Cambridge, MA. [4, 14, 90, 91]
- Noë, A. (2005). What does change blindness teach us about consciousness? *Trends in Cognitive Sciences*, Vol. 9(5), p. 218. [90]
- Nolfi, S. and Floreano, D. (2000). *Evolutionary Robotics: The Biology, Intelligence and Technology of Self-Organizing Machines*. MIT Press, Cambridge, MA. [4, 6, 7, 14]
- Nolfi, S. and Marocco, D. (2001a). Active perception: A sensorimotor account of object categorization. *From Animals to Animats 7. Proceedings of the 7th International Conference on Simulation of Adaptive Behavior* (ed. B. Hallam). MIT Press, Cambridge, MA. [4, 7]

- Nolfi, S. and Marocco, D. (2001b). Evolving robots able to integrate sensory-motor information over time. *Theory in Biosciences*, Vol. 120(3-4), pp. 287–310. [7, 32]
- Nolfi, S. and Parisi, D. (1999). Exploiting the power of sensory-motor coordination. *Advances in Artificial Life: Proceedings of the 5th European Conference on Artificial Life* (eds. D. Floreano, J. Nicoud, and F. Mondada), pp. 73–182, Springer-Verlag, New York, NY. [7, 21, 32]
- Nolfi, S. (1998). Evolutionary robotics: Exploiting the full power of self-organization. *Connection Science*, Vol. 10, pp. 167–184. [69]
- Nolfi, S. (2002a). Evolving robots able to self-localize in the environment: The importance of viewing cognition as the result of processes occurring at different time scales. *Connection Science*, Vol. 14, pp. 231–244. [70, 72, 80]
- Nolfi, S. (2002b). Power and limits of reactive agents. *Neurocomputing*, Vol. 42, pp. 119–145. [7, 14, 19, 21, 22, 39, 40, 60]
- Norman, D. A. (1993). Cognition in the head and in the world: An introduction to the special issue on situated action. *Cognitive Science*, Vol. 17(1), pp. 1–6. [12, 14, 15]
- Nosofsky, R. M. (1987). Attention and learning processes in the identification and categorization of integral stimuli. *Journal of Experimental Psychology: Learning Memory and Cognition*, Vol. 13, pp. 87–109. [3]
- O'Regan, J. K. and Noë, A. (2001a). Acting out our sensory experience: Authors' Response to Commentary. *Behavioral and Brain Sciences*, Vol. 24(5), pp. 1011–1021. [90]
- O'Regan, J. K. and Noë, A. (2001b). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, Vol. 24(5), pp. 939–1031. [4, 13, 14, 20, 21, 38, 90]
- O'Regan, J. K. (1992). Solving the 'real' mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology*, Vol. 46, pp. 461–488. [14, 20, 21, 38]
- O'Toole, A. J., Jiang, F., Abdi, H., and Haxby, J. V. (2005). Partially distributed representations of objects and faces in ventral temporal cortex. *Journal of Cognitive Neuroscience*, Vol. 17, pp. 580–590. [93, 94]
- Owen, A. M., Downes, J. J., J.Sahakian, B., Polkey, C. E., and Robbins, T. W. (1990). Planning and spatial working memory following frontal lobe lesions in man. *Neuropsychologia*, Vol. 28(10), pp. 1021–1034. [63]

- Palmer, S. E. (1999). *Vision Science: Photons to Phenomenology*. The MIT Press, Cambridge, MA. [49]
- Pfeifer, R. and Scheier, C. (1997). Sensory-motor coordination: The metaphor and beyond. *Robotics and Autonomous Systems*, Vol. 20, pp. 157–178. [39, 91]
- Pfeifer, R. and Scheier, C. (1999). *Understanding Intelligence*. MIT Press, Cambridge, MA. [1, 4, 6, 13, 14, 91]
- Phaf, R. H. and Wolters, G. (1997). A constructivist and connectionist view on conscious and unconscious processes. *Philosophical Psychology*, Vol. 10, pp. 287–307. [41, 42]
- Phillips, L. H., Wynn, V., Gilhooly, K. J., Sala, S. D., and Logie, R. H. (1999). The role of memory in the Tower of London task. *Memory*, Vol. 7(2), pp. 209–231. [53]
- Pitt, D. (2002). Mental representation. *Stanford Encyclopedia of Philosophy* (ed. E. N. Zalta). Stanford University, Stanford, CA. <http://plato.stanford.edu>. [3]
- Polk, T. A., Simen, P., Lewis, R., and Freedman, E. (2002). A computational approach to control in complex cognition. *Cognitive Brain Research*, Vol. 15(1), pp. 71–83. [43, 65, 66, 130]
- Port, R. and van Gelder, T. J. (1995). *Mind as Motion: Explorations in the Dynamics of Cognition*. MIT Press, Cambridge, MA. [4, 14]
- Postma, E. O., van Dartel, M. F., and Kortmann, L. J. (2001). Recognition by fixation. *Proceedings of the 13th Belgium-Netherlands Conference on Artificial Intelligence* (eds. B. Kröse, M. D. Rijke, G. Schreiber, and M. van Someren), pp. 425–432. [20]
- Putnam, H. (1960). *Minds and Machines*. New York University Press, New York, NY. [3]
- Pylyshyn, Z. (2003a). Mental Imagery: In search of a theory. *Behavioral and Brain Sciences*, Vol. 25(2), pp. 157–237. [89]
- Ramachandran, V. S. and Gregory, R. L. (1991). Perceptual filling in of artificially induced scotomas in human vision. *Nature*, Vol. 350, pp. 699–702. [21]
- Ratcliff, R., Shiffrin, R. M., and Clark, S. (1990). The list-strength effect: I. Data and discussion. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, Vol. 16, pp. 163–178. [3]
- Revonsuo, A. (2001). Can functional brain imaging discover consciousness in the brain? *Journal of Consciousness Studies*, Vol. 8(3), pp. 3–23. [94]

- Rochester, N., Holland, J. H., Haibt, L. H., and Duda, W. L. (1956). Tests on a cell assembly theory of the action of the brain, using a large digital computer. *Transactions on Information Theory*, Vol. 2, pp. 80–93. [123]
- Rosenschein, S. J. and Kaelbling, L. P. (1996). A situated view of representation and control. *Computational Theories of Interaction and Agency* (eds. P. E. Agre and S. J. Rosenschein), pp. 515–540. The MIT Press, Cambridge, MA. [11]
- Roskies, A. L. (1999). The binding problem. *Neuron*, Vol. 24, pp. 7–9. [94]
- Schank, R. C. and Abelson, R. P. (1975). *Scripts, Plans, Goals, and Understanding*. Lawrence Erlbaum Associates, Mahwah, NJ. [11]
- Schank, R. C. (1990). What is AI anyway? *The Foundation of Artificial Intelligence* (eds. D. Partridge and Y. Wilks), pp. 3–13. Cambridge University Press, Cambridge, MA. [1]
- Schnapf, J. L., Kraft, T. W., and Baylor, D. A. (1987). Spectral sensitivity of human cone photoreceptors. *Nature*, Vol. 325, pp. 439–441. [49]
- Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, Vol. 3(3), pp. 417–457. [93]
- Sejnowski, T. J. and Rosenberg, C. R. (1987). Parallel networks that learn to pronounce english text. *Complex Systems*, Vol. 1, pp. 145–168. [3]
- Shallice, T. (1982). Specific impairments of planning. *Philosophical Transactions of the Royal Society London*, Vol. B 298, pp. 199–209. [42, 43, 44, 45, 52, 53, 63]
- Shepard, R. and Chipman, S. (1970). Second order isomorphism of internal representation: Shapes of states. *Cognitive Psychology*, Vol. 1, pp. 1–17. [88, 89, 90]
- Shiffrin, R. M. (1988). Attention. *Stevens' Handbook of Experimental Psychology* (eds. R. C. Atkinson, R. J. Herrnstein, G. Lindzey, and R. D. Luce), pp. 739–811. Wiley, New York, NY, 2 edition. [3]
- Shiffrin, R. M., Ratcliff, R., and Clark, S. (1990). The list-strength effect: II. Theoretical mechanisms. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, Vol. 16, pp. 179–195. [3]
- Shlesinger, M. F. and Klafter, J. (1985). Lévy walk vs. Lévy flights. *On Growth and Form* (eds. H. E. Stanley and N. Ostrowsky). Martinus Nijhoff Publishers, Dordrecht, NL. [74]
- Sikora, M. D., Haley, P., Edwards, J., and Butler, R. W. (2002). Tower of London test performance in children with poor arithmetic skills. *Developmental Neuropsychology*, Vol. 21(3), pp. 243–254. [43]

- Simon, H. A. and Gilmarin, K. (1973). A simulation of memory for chess positions. *Cognitive Psychology*, Vol. 5, pp. 29–46. [3, 129]
- Simon, H. A. (1969). *The Sciences of the Artificial*. MIT Press, Cambridge, MA. [3, 92]
- Simons, D. J. and Rensink, R. A. (2005). Change blindness: Past, present, and future. *Trends in Cognitive Sciences*, Vol. 9(1), pp. 16–20. [90]
- Skinner, B. F. (1938). *The Behavior of Organisms: An Experimental Analysis*. Appleton-Century, New York, NY. [15]
- Slocum, A. C., Downey, D. C., and Beer, R. D. (2000). Further experiments in the evolution of minimally cognitive behavior: From perceiving affordances to selective attention. *From Animals to Animats 6. Proceedings of the 6th International Conference on Simulation of Adaptive Behavior* (eds. J.-A. Meyer, A. Berthoz, D. Floreano, H. Roitblat, and S. Wilson), pp. 430–439, MIT Press, Cambridge, MA. [4, 7, 70]
- Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences*, Vol. 11(1), pp. 1–23. [3, 5]
- Spronck, P. H. M., Sprinkhuizen-Kuyper, I. G., and Postma, E. O. (2001). Evolutionary learning of a neural robot controller. *Proceedings of the International Conference on Computational Intelligence for Modelling, Control and Automation* (ed. M. Mouhammadian), pp. 510–518, IOS Press, Amsterdam, NL. [72, 73]
- Stanley, H. E., Amaral, L. A. N., Gopikrishnan, P., Ivanov, P. C., Keitt, T. H., and Plerou, V. (2001). Scale invariance and universality: Organizing principles in complex systems. *Physica A*, Vol. 281, pp. 60–68. [70, 81, 82]
- Steels, L. (1994). Emergent functionality in robotic agents through on-Line evolution. *Proceedings of Alife IV* (eds. R. Brooks and P. Maes), pp. 8–16, MIT Press, Cambridge, MA. [81]
- Steyvers, M., Shiffrin, R. M., and Nelson, D. L. (2004). Word association spaces for predicting semantic similarity effects in episodic memory. *Experimental Cognitive Psychology and its Applications* (ed. A. Healy). American Psychological Association, Washington, DC. [89]
- Tani, J. and Nolfi, S. (1999). Learning to perceive the world as articulated: An approach for hierarchical learning in sensory-motor systems. *Neural Networks*, Vol. 12(7-8), pp. 1131–1141. [66]
- Thelen, E. and Smith, L. B. (1994). *A Dynamic Systems Approach to the Development of Cognition and Action*. MIT Press, Cambridge, MA. [3, 5, 13, 14, 88, 91]

- Thelen, E. (1995). Motor development: A new synthesis. *American Psychologist*, Vol. 50(2), pp. 79–95. [4, 32, 90]
- Tijsseling, A. (1998). *Connectionist Models of Categorization: A Dynamical Approach to Cognition*. Ph.D. thesis, University of Southampton, Southampton, UK. [39]
- Tinbergen, N. (1951). *The Study of Instinct*. Oxford University Press, New York, NY. [19]
- Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, Vol. 55, pp. 189–208. [15]
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, Vol. 59, pp. 433–460. [41, 93]
- van Dartel, M. F. and Postma, E. O. (2002). Analysing situated behaviour from a macroscopic perspective. *Proceedings of the 4th Workshop on Dynamic Perception* (eds. R. P. Würtz and M. Lappe), pp. 165–170. IOS Press. [69]
- van Dartel, M. F. and Postma, E. O. (2005). Symbol manipulation by internal simulation of perception and behaviour. *Proceedings of the 5th Workshop on Epigenetic Robotics* (eds. C. Balkenius, L. Berthouze, F. Kaplan, J. Konczak, H. Kozima, G. Metta, J. Nadel, G. Sandini, G. Stojanov, and H. Yano), pp. 121–124. [41]
- van Dartel, M. F. and Postma, E. O. (in press). Symbol manipulation by internal simulation of perception and behaviour. *Proceedings of the 17th Belgium-Netherlands Conference on Artificial Intelligence*. [41]
- van Dartel, M. F., Postma, E. O., and van den Herik, H. J. (2002). Universal properties of adaptive behaviour. *Proceedings of the 14th Belgium-Netherlands Conference on Artificial Intelligence* (eds. H. Blockeel and M. Denecker), pp. 59–66. [69]
- van Dartel, M. F., Sprinkhuizen-Kuyper, I. G., Postma, E. O., and van den Herik, H. J. (2003). Reactive agents and perceptual ambiguity. *Proceedings of the 15th Belgium-Netherlands Conference on Artificial Intelligence* (eds. T. Heskes, P. Lucas, L. Vuurpijl, and W. Wiegerinck), pp. 75–82. [19]
- van Dartel, M. F., Postma, E. O., and van den Herik, H. J. (2004a). Categorisation through internal simulation of perception and behaviour. *Proceedings of the 16th Belgium-Netherlands Conference on Artificial Intelligence* (eds. L. Schomaker, N. Taatgen, and R. Verbrugge), pp. 187–194. [121]
- van Dartel, M. F., Postma, E. O., van den Herik, H. J., and de Croon, G. (2004b). Macroscopic analysis revealing a universal property. *Connection Science*, Vol. 16(3), pp. 169–181. [4]

- van Dartel, M. F., Sprinkhuizen-Kuyper, I. G., Postma, E. O., and van den Herik, H. J. (2005). Reactive agents and perceptual ambiguity. *Adaptive Behavior*, Vol. 13(3), pp. 227–242. [4, 70]
- van Gelder, T. J. (1995). What might cognition be, if not computation? *Journal of Philosophy*, Vol. 91, pp. 345–381. [3, 5]
- van Gulick, R. (2001). Still room for representations. *Behavioural and Brain Sciences*, Vol. 24(5), pp. 1007–1008. (A commentary on: A sensorimotor account of vision and visual consciousness. by J.K. ORegan and A. Noë). [90]
- Viswanathan, G. M., Buldyrev, S. V., Havlin, S., Luz, M. G. E. da, Raposo, E. P., and Stanley, H. E. (1999). Optimizing the success of random searches. *Nature*, Vol. 401, pp. 911–914. [74, 76]
- Viswanathan, G. M., Afanasyev, V., Buldyrev, S. V., Havlin, S., Luz, M. G. E. da, Raposo, E. P., and Stanley, H. E. (2001). Lévy flights search pattern of biological organisms. *Physica A*, Vol. 295, pp. 85–88. [74, 76]
- Wiener, N. (1948). *Cybernetics, or Control and Communication in the Animal and the Machine*. John Wiley and Sons, Sussex, UK. [4]
- Winograd, T. (1990). Thinking machines: Can there be? Are we? *The Foundation of Artificial Intelligence*, pp. 167–189. Cambridge University Press, Cambridge, MA. [1]
- Zhang, J. and Norman, D. A. (1994). Representations in distributed cognitive tasks. *Cognitive Science*, Vol. 18, pp. 87–122. [16]
- Zhang, J. (1997). The nature of external representations in problem solving. *Cognitive Science*, Vol. 21(2), pp. 179–217. [15, 16]
- Ziemke, T. (2003). What’s that thing called embodiment? *Proceedings of the 25th Annual Meeting of the Cognitive Science Society* (eds. R. Alterman and D. Kirsh), pp. 1134–1139. Lawrence Erlbaum, Mahwah, NJ. [6]
- Ziemke, T. (2004). Embodied AI as science: Models of embodied cognition, embodied models of cognition, or both? *Embodied Artificial Intelligence* (eds. F. Iida, R. Pfeifer, L. Steels, and Y. Kuniyoshi), Vol. 3139 of *Lecture Notes in Computer Science*, pp. 27–36. Springer, Berlin, DE. [6]
- Ziemke, T., Jirnhed, D.-A., and Hesslow, G. (2005). Internal simulation of perception. A minimal neurobotic model. *Neurocomputing*, Vol. 68, pp. 85–104. [42, 50, 66, 126]
- Zippelius, A. and Engel, A. (2003). Statistical mechanics of neural networks. *The Handbook of Brain Theory and Neural Networks* (ed. M. Arbib). MIT Press, Cambridge, MA. [72]

Appendix A

Environmental complexity of Acp

To investigate whether the results reported in chapter 3 generalise over more complex environmental dynamics, we conducted two additional experiments in which we varied the horizontal and vertical velocities of the falling objects in ACP. Both experiments and their results are reported in section A.1. A discussion and conclusion follow in section A.2.

A.1 Experiments and results

The following two subsections report the experiments and the results found with variable horizontal (subsection A.1.1) and variable vertical (subsection A.1.2) velocities of falling objects in ACP. In both these additional experiments, only agents without blind sensors ($b = 0$) were optimised.

A.1.1 Variable horizontal velocities

In the first experiment, ACP's environmental dynamics were made more complex by varying the *horizontal* velocities at which objects moved. Whereas, in our original experiments, objects moved two grid cells leftward or rightward per time step, in our first additional experiment objects moved 1, 2, 3, or 4 cells leftward or rightward per time step. The size of the movement was kept constant during a trial, but varied between trials.

The average success rate of agents with all three types of neurocontroller is shown in figure A.1 for ACP with objects falling at variable horizontal velocities.

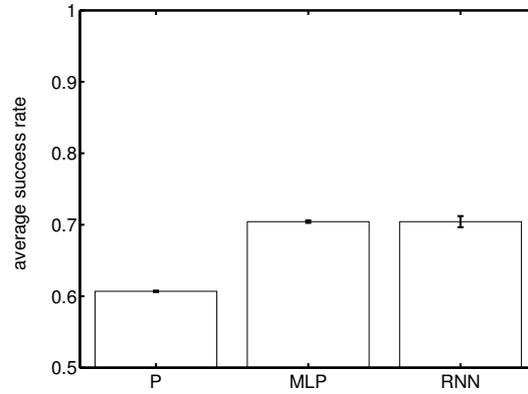


Figure A.1: Average success rate of the best-performing agents on the active categorical perception task with objects moving at variable horizontal velocities. Each bar represents the averaged performance of agents with a specific type of neurocontroller (P, MLP, RNN).

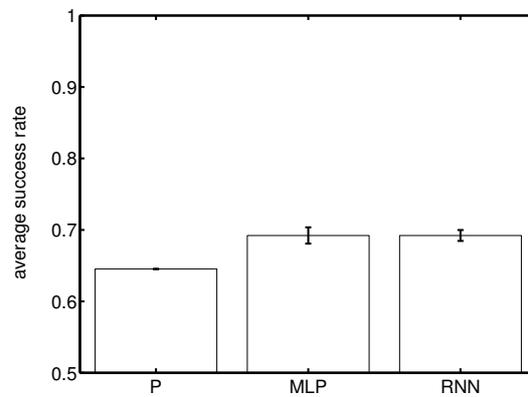


Figure A.2: Average success rate of the best-performing agents on the active categorical perception task with objects moving at variable vertical velocities. Each bar represents the averaged performance of agents with a specific type of neurocontroller (P, MLP, RNN).

A.1.2 Variable vertical velocities

In the second additional experiment, ACP's environmental dynamics were made more complex by varying the *vertical* velocities at which objects moved. Whereas objects moved one grid cell per time step in our original experiment, in the second additional experiment we varied the number of grid cells moved per time step from 1 to 3. Again, the velocity within a trial was kept constant, whereas it varied between the trials.

The average success rate of agents with all three types of neurocontroller is shown in figure A.2 for ACP with objects falling at variable vertical velocities.

A.2 Discussion and conclusion

The results for both additional experiments show an overall decrease in performance as compared to the original results for $b = 0$ in figure 3.8. We assume that this is due to the increased complexity of the task. However, the relative pattern of performances remains largely unaffected by the variation of the velocities of falling objects. So, we may conclude that this indicates that the results reported in chapter 3 generalise across moderate variations in the environmental dynamics.

Appendix B

Categorisation through internal simulation

ACP was used to perform the categorisation experiments reported in chapter 3. Below, we introduce a revised version of ACP, which we named ACP^{*}, to test whether cognitive systems benefit from the ability to simulate perception and behaviour internally in the categorisation task of ACP¹. In section B.1, ACP^{*} is discussed. The experiments conducted with ACP^{*} and the results found are reported in section B.2. Finally, in section B.3, the results are discussed and concluded upon.

B.1 Acp^{*}

To be able to test whether cognitive systems benefit from the ability to simulate perception and behaviour internally in ACP, ACP is extended with three mechanisms to form ACP^{*}: (1) an output-input feedback mechanism, to enable agents to simulate perception and behaviour internally (subsection B.1.1), (2) an occlusion mechanism, to occlude agents from environmental input for a variable period of time (subsection B.1.2), and (3) a normalisation mechanism, to normalise the input received by the agent (subsection B.1.3).

B.1.1 Output-input feedback mechanism

To enable agents to simulate perception and behaviour internally in ACP^{*}, all three types of ACP neurocontroller (P, MLP, and RNN) are extended with an output-input feedback mechanism that consists of an additional output node for each input

¹The experiments and results reported here were presented at the 16th Belgium-Netherlands Conference on Artificial Intelligence, Groningen, The Netherlands (see van Dartel, Postma, and van den Herik (2004a)). The author would like to thank his co-authors and the editors of the proceedings for their kind permission to reuse relevant parts of the article in this thesis.

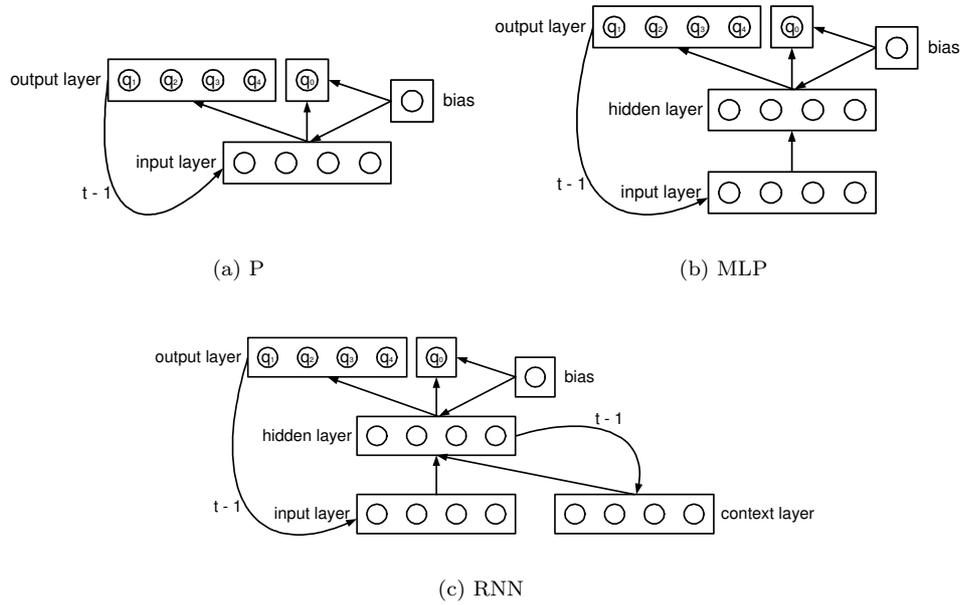


Figure B.1: Topology of three types of extended neurocontrollers (P, MLP, and RNN). The additional output nodes are labelled q_1 to q_4 .

node. The output nodes are assigned indices that range from zero to the number of input nodes. The original ACP output node is assigned the index zero, the additional output nodes are assigned the indices 1 to the number of input nodes. The patterns that these extra output nodes produce constitute internal input patterns. At time step t , the input nodes are excited by the superposition of the external input ('the environmental input') and the internal input produced at time step $t-1$. To implement the output-input feedback mechanism, equation 3.2 is replaced by

$$I(x, t) = G_t(x, y_{max} - t) + \tanh(O(q, t - 1)) \quad (\text{B.1})$$

with $O(q, t - 1)$ the q -th element of the output pattern (i.e., the internal input pattern) at time step $t - 1$, where q is defined as $q = x - x_s + 1$ with x_s the x position of the first sensor. Figure B.1 shows how the output-input feedback mechanism is embedded in the topology of the three types of neurocontrollers (P, MLP, and RNN).

Each of the additional output nodes' activation is initialised by assigning a value of 0 at $t = 0$, as was done for all other nodes' activation of neurocontrollers in ACP*. By optimising the weights between the output nodes with $q > 0$ and the input nodes, the internal input patterns constitute a recurrency much like the recurrent connections on the hidden nodes of the RNN do. However, recurrent connections from

output to input nodes allow agents to predict future external inputs, i.e., simulate perception and behaviour internally, while recurrent connections on hidden nodes do not.

B.1.2 Occlusion mechanism

To test the ability of agents to simulate perception and behaviour internally with ACP*, we extend the model with an occlusion mechanism. The mechanism occludes the environmental input for a predefined period of time ot . We assume that the occlusion mechanism encourages agents to simulate the future sensor states (i.e., environmental input). Occlusion always occurs at the last ot consecutive time steps between $t = 0$ and $t = t_{max}$. The occlusion mechanism is implemented by replacing equation B.1 by equation B.2 for $t > y_{max} - ot$.

$$I(x, t) = \tanh(O(q, t - 1)) \quad (\text{B.2})$$

B.1.3 Normalisation mechanism

Whereas no normalisation of input occurred in ACP, in ACP* activation of the input nodes is normalised by

$$I_{norm}(x, t) = I(x, t) * (2 / \sum_{x=1}^{x_{max}} I(x, t)) \quad (\text{B.3})$$

If $\sum_{x=1}^{x_{max}} I(x, t) = 0$, then equation B.3 is ignored and the activation of all input nodes is set to 0.5. The normalisation mechanism keeps the summed normalised input constant (cf., Grossberg (1987)) at a value of 2 irrespective of the source of input (internal, external, or internal and external), which enhances the biological plausibility of the model (Rochester *et al.*, 1956).

B.2 Experiments and results

For the experiments described below, agents are equipped with four sensors that directly pass their binary activation on to the input nodes. The boundaries of grid G_t are set to $x_{max} = 20$ and $y_{max} = 10$. As in the experiments reported in chapter 3, the evolutionary algorithm is applied for 20,000 generations with a population size of 100 agents. Through optimisation of a neurocontroller's weights the evolutionary algorithm can optimise the input-output mapping, including the output-input feedback mechanism encoded in the weights between the hidden nodes and the output nodes with $q > 0$. These weights are initialised by assignment of random values, as was done with all other weights of neurocontrollers in ACP*.

Experiments with ACP* are conducted for all three types of neurocontrollers and for three conditions in which output-input feedback was: (1) not present (no feedback), (2) only present during occlusion (feedback during occlusion), or (3) present at all time steps (feedback at all times).

In the first condition (no feedback) equation 3.2 is used for all time steps t to update the input nodes. In the second condition (feedback during occlusion) equation 3.2 is replaced by equation B.1 for time steps $t > t_{max} - ot$ to update the input nodes (hence, at $t = 0$ agents are only initialised in G_t and no input is received). In the third condition (feedback at all times) equation 3.2 is replaced by equation B.1 for all time steps t to update the input nodes.

All nine experimental conditions, i.e., the three types of neurocontroller and the three output-input feedback conditions, are tested for all possible values of ot , $0 \leq ot \leq t_{max}$. Each experiment is repeated 5 times, over which the success rates of the best-performing agents are averaged. To compare the three feedback conditions, we plotted the average success rates of agents in each condition over the different occlusion times for each type of neurocontroller (see figures B.2, B.3, and B.4). Standard deviations (*sd*) were computed for all data points, but appear as single lines in figures B.2, B.3, and B.4, because $sd < 0.001$ for all data points.

Figure B.2 shows the average success rate of P-controlled agents for the three feedback conditions. For most occlusion times, feedback during occlusion leads to the best-performing agents on the active categorical perception task, and feedback at all times leads to the second-best-performing agents. Figure B.3 shows the average success rate of MLP-controlled agents for the three feedback conditions. Again, feedback during occlusion leads to the best-performing agents on the active categorical perception task, and feedback at all times leads to the second-best-performing agents for most occlusion times. Figure B.4 shows the same results for the RNN-controlled agents. For most occlusion times, no feedback and feedback during occlusion are preferred over feedback at all times. For RNN-controlled agents there seems to be no general preference for feedback during occlusion over no feedback. Four relations between figures B.2, B.3, and B.4 can be observed. First, when $ot = 0$, agents with feedback during occlusion perform equally well compared to agents that do not use the output-input feedback mechanism at all. The reason is that the feedback is not operative for this value of ot . Second, when $ot = 9$ all agents perform at chance level, i.e., an average success rate of 0.5. For this value of ot , the occlusion time extends over the entire simulation episode, i.e., no external input is received. Third, the main variations in average success rates of the different agents occur for intermediate occlusion times. Fourth, the average success rate does not vary smoothly with the occlusion time for any of the agents. For instance, there is an increase in performance for most agents when $ot = 1$ compared to when $ot = 0$, while an increase in ot results in a decrease in the average success rate. This can be attributed to how the problem space relates to variables in ACP^* , such as the occlusion time and the neural structure of an agent. Although it would be interesting to reveal the exact relation between the problem space and the variables in ACP^* , such an analysis is beyond the scope of our current research goal.

The main result of the experiments is that for both the P and MLP-controlled agents there is a general preference, although small, for feedback (either during occlusion or at all times) above no feedback at all (see figures B.2 and B.3). This general preference is not present in the average success rates of RNN-controlled agents

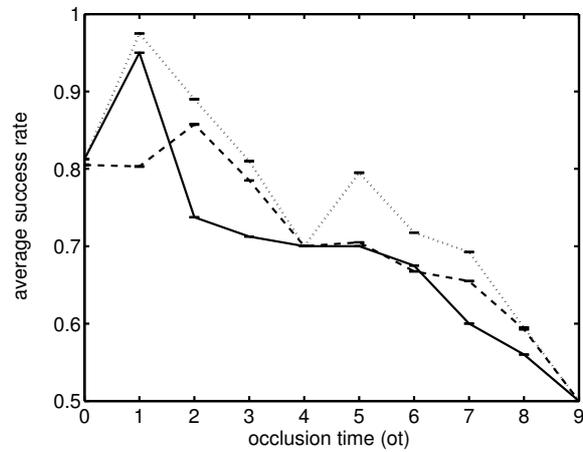


Figure B.2: Average success rate of P-controlled agents as a function of occlusion time ot . The three conditions are (1) no feedback (solid lines), (2) feedback during occlusion (dotted lines), and (3) feedback at all times (dashed lines).

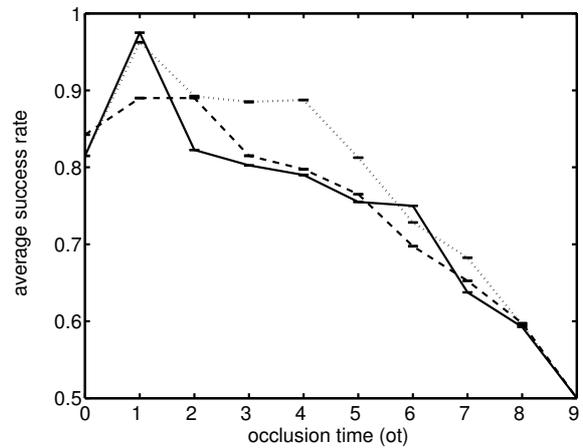


Figure B.3: Average success rate of MLP-controlled agents as a function of occlusion time ot . The three conditions are (1) no feedback (solid lines), (2) feedback during occlusion (dotted lines), and (3) feedback at all times (dashed lines).

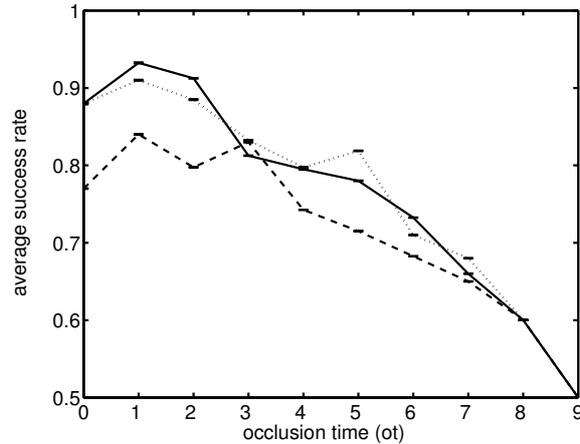


Figure B.4: Average success rate of RNN-controlled agents as a function of occlusion time ot . The three conditions are (1) no feedback (solid lines), (2) feedback during occlusion (dotted lines), and (3) feedback at all times (dashed lines).

(see figure B.4). It implies that feed forward-controlled agents benefit from output-input feedback to cope with occlusion from external input, while recurrent-controlled agents do not. Apparently, the advantage of internal simulation of perception and behaviour in ACP* is restricted to agents with feed forward neurocontrollers.

B.3 Discussion and conclusions

Our results show that, for this simple task, feed forward-controlled agents benefit from the ability to simulate perception and behaviour internally. This benefit was also pointed out by Ziemke, Jirnhed, and Hesslow (Ziemke *et al.*, 2005), who trained agents to follow ‘blindly’ corridors without collision, using predictions of sensory input instead of real sensory input. Ziemke *et al.* found agents with recurrent neurocontrollers to perform worse than agents controlled by feed forward neurocontrollers. However, they did not compare these results with the performance of agents without the ability to simulate perception and behaviour. Therefore, they failed to note that there may be no benefit from internal simulation of perception and behaviour for recurrent-controlled agents.

RNN neurocontrollers without an output-input feedback mechanism are still able to generate internal input, because they have recurrent connections on their hidden units. However, the feedback produced by the output-input feedback mechanism differs from RNN-type feedback in that the latter maps hidden activity onto the hidden layer, rather than mapping output activity onto the input layer. Apparently, this difference is immaterial to the success rate while coping with occlusion in the categorisation task.

Our results agree with those of Ziemke et al., but show in addition that the benefit from internal simulation of perception and behaviour is restricted to agents with feed forward neurocontrollers. Our results suggest that any recurrency, be it by means of an internal input or by means of RNN-type feedback, suffices to deal with occluded external input. On the basis of this finding, we may draw two closely related conclusions: (1) the ‘simulation hypothesis’ may be too specific, and (2) predicting future perception depends on neural recurrency in general, rather than on the ability to simulate perception by feeding back actions. Further studies are needed to elucidate the types of recurrence required to deal with events or objects that are temporarily out of view.

List of abbreviations and symbols

α	angle
β	orientation of the nearest food element
μ	descriptor of the shape of the step-length distribution
a	rounded output (i.e., action)
AI	Artificial Intelligence
ACP	Active Categorical Perception (original model)
ACP*	Active Categorical Perception (revised model)
b	number of blind sensors
C	number of balls positioned correctly
CA	sum of correctly avoided objects
CC	sum of correctly caught objects
d	direction of movement
DC97	sToL model reported by Dehaene and Changeux (1997)
F	fitness
FA	sum of avoided objects that should have been caught
FC	sum of caught objects that should have been avoided
G	a two-dimensional grid
h	number of hidden nodes
i	index
I	sensor activation (i.e., input)
I_{norm}	normalised input
j	object element
j_{max}	width of an object
L	vertical and horizontal size of the foraging environment
l	step length
M	number of moves made
M_{min}	minimal number of moves
MAPP	computational model of expert chess memory designed by Simon and Gilmarin (1973)
MLP	multilayer perceptron

n	number of food elements
O	output pattern
ot	occlusion time
P	perceptron
PSLF02	sToL model reported by Polk <i>et al.</i> (2002)
q	index
r	radius
RNN	recurrent neural network
S	number of solved problems
s	number of functional sensors
S_i	i^{th} sensor
S_{max}	maximum number of problems
sd	standard deviation
\overline{sr}	average success rate
sToL	situated Tower of London
SSA	sensory state-action
SST	sensory state-transition
t	time
ToL	Tower of London
TVSS	tactile visual substitution system
u	units in degrees
x	horizontal position
x_{max}	width of grid G
xs	horizontal position of the first sensor
y	vertical position
y_{max}	height of grid G
Z	normalising constant

Summary

The notion of representation is well-defined within the traditional computational theory of mind. However, in the relatively novel theory of situated cognition this is not the case. The focus of this thesis is on the nature of representation in situated systems, i.e., situated representation.

In the first chapter the problem concerning situated representation is outlined. The chapter indicates that the cognitive sciences are in need of an operationalisation of the notion of situated representation. To investigate a possible realisation of such an operationalisation the following problem statement is formulated: what is the nature of representation in situated systems? Subsequently, two research questions are formulated to investigate the problem statement: (i) to what extent can we identify where the knowledge resides that is used by a situated system to solve a certain task? and (ii) how is this knowledge accessed and used by a situated system when solving a certain task? Furthermore, in this chapter, the methodology of our investigation is described in terms of five conditions for models of cognition (situatedness, embodiment, cognition, parsimony, and transparency) and a short outline of the three models applied in the thesis (a robot model of active categorical perception, a robot model of the Tower of London task, and a robot model of foraging) is given.

The second chapter provides background information. It elaborates on representation issues in situated systems and on the two types of representation that such systems use, internal representation and external representation. This elaboration guides the empirical study of situated representation in artificial systems in the third, fourth, and fifth chapter.

In the third chapter, the behaviour of robots in a simple model of active categorical perception is examined. The effective sensorimotor mapping of optimised situated robots clarifies the notion of situated representation. The model provides a unique opportunity to study situated representations, because the perceptual ambiguity in the model forces successful situated robots to represent adequately the information they need to perform a given task. The findings obtained by employing the active categorical perception model and the subsequent analyses lead us to four conclusions: (i) reactive robots can cope with perceptual ambiguity in the context of active categorical perception, (ii) reactive robots can organise their behaviour according to sensory stimuli that are no longer present using the environment as

an external memory, (iii) reactive robots incorporating a non-linear sensorimotor mapping are better capable of dealing with perceptual ambiguity in an active categorical perception task than those incorporating a linear mapping, and (iv) sensory state-transition diagrams provide insight into the behavioural strategies employed by reactive robots to deal with perceptual ambiguity and their use of the environment as an external memory. Moreover, the findings obtained by employing the active categorical perception model and the subsequent analysis demonstrate that representation in situated systems can be internal and/or external. This indicates that the operationalisation of the notion of situated representation should allow internal representation and external representation.

The fourth chapter studies the nature of internal representation. Internal representation is often associated with planning in symbol manipulation tasks. In order to study the nature of internal representation, in this chapter we investigate representation in a situated robot model of a typical planning task that requires symbol manipulation, the Tower of London task. The results obtained with the situated Tower of London model and the subsequent analyses lead us to conclude that the ability to perform (situated) symbol manipulation by internal simulation of perception and behaviour allows the robot to plan ahead in time. Our second conclusion is that representation of both the current and future states of the environment occurs through the mapping of sensor-array activations to actions. For the current state the activation is received from the environment and for the future state the activation is received from the internally generated expected state. The two conclusions indicate that in order to operationalise situated internal representation, the operationalisation should allow internal simulation of perception and behaviour.

The fifth chapter studies the nature of external representation of a situated robot performing a foraging task in a stochastic environment. In order to investigate how the externally represented knowledge is accessed and used by the situated robot, we analyse the robot-environment interaction in the situated model of foraging by two different types of analysis (microscopic and macroscopic). The analyses of the results obtained with the situated model of foraging lead us to three conclusions: (i) macroscopic analysis may predict a universal property that can be explained at the microscopic level by microscopic analysis, (ii) macroscopic analysis may complement microscopic analysis in the study of adaptive behaviour, and (iii) macroscopic analysis may be preferred over microscopic analysis, owing to its power to reveal universal properties. Moreover, the experiment with the situated model of foraging and both the analyses show that external representation may reside in the average properties of the interaction with the environment. Robots in the model of foraging do not represent individual food elements (or their physical locations) by element-specific interaction, but represent the uniform distribution of those food elements in their average interaction with the environment. These findings indicate that in order to incorporate external representation into the operationalisation of the notion of situated representation, the operationalisation should allow representation by the average interaction with the environment.

The sixth chapter combines the results of the three investigations reported in the

preceding chapters. On the basis of these investigations we formulate a new operationalisation of the notion of representation. The new operationalisation holds that for an entity to be adequately represented by a system, it is implied that the system is able to perform and/or simulate internally the entity-specific interaction with the environment. Four advantages of the new operationalisation over its non-situated counterpart are discussed, these advantages concern: (i) external representation and internal representation, (ii) the representation debate, (iii) situated accounts of cognition and awareness, and (iv) the symbol grounding problem. Thereafter, the two operationalisations are related to each other, from which we arrive at the belief that the operationalisation of non-situated representation should be replaced by the new operationalisation. Furthermore, in this chapter, a discussion on the possible implications of the new operationalisation indicates that the new operationalisation may have implications for the fields of artificial intelligence, cognitive neuroscience, and cognitive psychology.

In the seventh chapter we answer the research questions formulated in the first chapter by stating that: (1) we can identify where the knowledge resides that is used by a situated system to solve a certain task to the extent that we can reveal the coordination between the sensory and motor system(s) of a system and relate it to the environmental dynamics, and (2) the knowledge which a situated system uses to solve a certain task is accessed and used by:

- (i) exploiting the attractors in the interaction with the environment (chapter 3),
- (ii) simulating interaction with the environment internally (chapter 4), and
- (iii) exploiting the average properties of the interaction with the environment (chapter 5).

Furthermore, in the seventh chapter, we answer the problem statement formulated in the first chapter by stating that in emphasising the role of interaction for cognition in the theory of situated cognition the operationalisation of situated representation is essential. We conclude by stating that, in a situated system, representation is as strongly rooted in the environment as the system itself, i.e., *representation is situated in nature*.

Samenvatting

De notie van representatie kent een heldere definitie in de traditionele ‘computational theory of mind’. De relatief nieuwe ‘theory of situated cognition’ ontbeert een heldere definitie van representatie. Dit proefschrift richt zich op de aard van representatie in gesitueerde systemen.

In het eerste hoofdstuk wordt het probleem van gesitueerde representatie aan de orde gesteld. Het hoofdstuk geeft aan dat de cognitiewetenschappen een operationalisatie van de notie van gesitueerde representatie behoeft. Om een dergelijke operationalisatie op zijn mogelijkheden te onderzoeken wordt de volgende probleemstelling geformuleerd: wat is de aard van representatie in gesitueerde systemen? Vervolgens worden er twee onderzoeksvragen geformuleerd om de probleemstelling te bestuderen: (i) in hoeverre kan worden geïdentificeerd waar de kennis wordt gerepresenteerd die gesitueerde systemen gebruiken om een bepaalde taak uit te voeren? en (ii) hoe wordt deze kennis benaderd en gebruikt door een gesitueerd systeem bij het uitvoeren van een bepaalde taak? Tevens wordt in dit hoofdstuk de methodologie van het onderzoek beschreven in termen van vijf condities voor modellen van cognitie (‘situatedness’, ‘embodiment’, ‘cognition’, ‘parsimony’, and ‘transparency’) en wordt een korte schets gegeven van de drie modellen die in dit proefschrift worden toegepast (een robotmodel van actieve categorische perceptie, een robotmodel van de ‘Tower of London’ taak en een robotmodel van foerageergedrag).

Het tweede hoofdstuk geeft achtergrondinformatie bij de in het eerste hoofdstuk geformuleerde probleemstelling. Het hoofdstuk beschrijft uitvoerig wat representatie in gesitueerde systemen is en vermeldt twee typen representatie die door dergelijke systemen worden gebruikt, interne en externe representatie. De beschrijving geeft aan waarop het empirisch onderzoek naar gesitueerde representatie in kunstmatige systemen zich zal richten in het derde, vierde en vijfde hoofdstuk.

In het derde hoofdstuk wordt het gedrag van robots in een eenvoudig model van actieve categorische perceptie bestudeerd. De perceptuele ambiguïteit vereist dat robots acties uitvoeren om de objecten te kunnen classificeren. De effectieve koppeling tussen de sensoren en de actuatoren van geoptimaliseerde robots verheldert de notie van gesitueerde representatie. De bevindingen verkregen uit het model van actieve categorische perceptie en de daaropvolgende analyses leiden tot de volgende vier conclusies: (i) reactieve robots kunnen omgaan met perceptuele ambiguïteit in een actieve categorische perceptietaak, (ii) reactieve robots zijn in staat om gedrag te

organiseren aan de hand van stimuli die niet langer aanwezig zijn door de omgeving te gebruiken als een extern geheugen, (iii) reactieve robots met niet-lineaire koppelingen tussen sensoren en actuatoren kunnen beter omgaan met perceptuele ambiguïteit in een actieve categorische perceptietaak dan robots met lineaire koppelingen tussen sensoren en actuatoren, en (iv) ‘sensor state-transition’ diagrammen geven inzicht in de gedragsstrategieën die worden toegepast door reactieve robots ten aanzien van perceptuele ambiguïteit en de omgeving als extern geheugen. Bovendien tonen de bevindingen die verkregen zijn met het model van actieve categorische perceptie, en de daarbij behorende analyses, aan dat representatie door gesitueerde systemen intern en/of extern kan plaatsvinden. Dit duidt erop dat de operationalisatie van de notie van gesitueerde representatie dient te voorzien in zowel interne als externe representatie.

Het vierde hoofdstuk bestudeert de aard van interne representatie. Interne representatie wordt doorgaans geassocieerd met planning in symboolmanipulatietaken. Om de aard van interne representatie te bestuderen onderzoeken we in dit hoofdstuk een gesitueerd robotmodel van de ‘Tower of London’ taak, een typische plannings-taak waarbij symboolmanipulatie een grote rol speelt. Op basis van de bevindingen verkregen uit het gesitueerde ‘Tower of London’ model en de daarbij behorende analyses mogen we concluderen dat (gesitueerde) symboolmanipulatie door interne simulatie van perceptie en gedrag aan de basis ligt van het planningsvermogen van de robot. Bovendien concluderen we dat zowel de huidige als de toekomstige toestand van de omgeving door de robot wordt gerepresenteerd door sensoractivatie aan acties te koppelen. De huidige toestand van de omgeving wordt gerepresenteerd door activatiepatronen die ontvangen worden vanuit de omgeving. De toekomstige toestand wordt gerepresenteerd door intern gegenereerde activatiepatronen. Onze conclusies geven aan dat de operationalisatie van de notie van gesitueerde representatie het mechanisme van interne simulatie van perceptie en gedrag dient te omvatten.

Het vijfde hoofdstuk bestudeert de aard van externe representatie in een model van een gesitueerde robot die een foerageertaak uitvoert in een stochastische omgeving. Om te begrijpen hoe extern gerepresenteerde kennis wordt gebruikt door de gesitueerde robot, analyseren we de interactie tussen de robot en zijn omgeving door middel van twee typen analyse (microscopische en macroscopische). De analyses leiden tot de volgende drie conclusies: (i) macroscopische analyse kan een universele eigenschap voorspellen die op het microscopische niveau kan worden verklaard, (ii) macroscopische analyse kan microscopische analyse ondersteunen in het bestuderen van adaptief gedrag en (iii) macroscopische analyse kan worden verkozen boven microscopische analyse omdat het universele eigenschappen kan onthullen. Bovendien geven de experimenten met het gesitueerde model van foerageergedrag en de bijbehorende analyses aan dat de gemiddelde eigenschappen van de interactie met de omgeving de externe representatie vormt. Robots in het gesitueerde model van foerageergedrag representeren geen individuele voedsелеlementen (of hun fysieke lokatie) door element-specifieke interactie, maar representeren de uniforme distributie van de voedsелеlementen in de gemiddelde interactie met de omgeving. Deze bevindingen geven aan dat de nieuwe operationalisatie van representatie moet

voorzien in de gemiddelde interactie met de omgeving als basis voor externe representatie.

Het zesde hoofdstuk combineert de resultaten van de onderzoeken uit de voorgaande drie hoofdstukken. Op basis van deze resultaten formuleren we een nieuwe operationalisatie van de notie van representatie. De nieuwe operationalisatie houdt in dat een systeem, om een entiteit adequaat te representeren, in staat moet zijn om entiteit-specifieke interactie met de omgeving uit te voeren en/of intern te simuleren. Vervolgens worden in dit hoofdstuk vier voordelen van de nieuwe operationalisatie ten opzichte van zijn niet-gesitueerde tegenhanger behandeld die betrekking hebben op: (i) externe en interne representatie, (ii) het representatiedebat, (iii) gesitueerde benaderingen van cognitie en bewustzijn, en (iv) het ‘symbol grounding’ probleem. Hierna worden de twee operationalisaties met elkaar vergeleken. Deze vergelijking leidt tot de opvatting dat de operationalisatie van niet-gesitueerde representatie moet worden vervangen door de nieuwe operationalisatie. Bovendien worden in dit hoofdstuk de mogelijke implicaties van de nieuwe operationalisatie voor de onderzoeksgebieden kunstmatige intelligentie, cognitieve neurowetenschap en cognitieve psychologie besproken.

In het zevende hoofdstuk beantwoorden we de onderzoeksvragen die werden geformuleerd in het eerste hoofdstuk als volgt: (i) waar de door gesitueerde systemen gebruikte kennis wordt gerepresenteerd kan worden geïdentificeerd voor zover we de koppelingen tussen sensoren en actuatoren kunnen begrijpen en deze kunnen relateren aan de dynamiek van de omgeving, en (ii) deze kennis wordt benaderd en gebruikt door een gesitueerd systeem bij het oplossen van een bepaalde taak door:

- (i) het exploiteren van attractoren in de interactie met de omgeving (hoofdstuk 3),
- (ii) het intern simuleren van interactie met de omgeving (hoofdstuk 4) en
- (iii) het exploiteren van de gemiddelde eigenschappen van interactie met de omgeving (hoofdstuk 5).

Tenslotte wordt de probleemstelling uit het eerste hoofdstuk beantwoord door de claim dat de operationalisatie van gesitueerde representatie essentieel is bij het benadrukken van de rol van de omgeving voor cognitie binnen de ‘theory of situated cognition’. We concluderen tenslotte dat, in een gesitueerd systeem, representatie net zo sterk gekoppeld is aan de omgeving als het systeem zelf, in andere woorden: *representatie is gesitueerd van aard.*

Curriculum vitae

Michel van Dartel was born in Eindhoven, on May 26, 1976. He attended secondary school at the Bisschop Bekkers College in Eindhoven, where he received his HAVO diploma in 1993. In 1995 he received a degree in administrative economics from the Economisch Lyceum Eindhoven and, subsequently, in 1996, he received a *propedeuse* in Personnel and Organisation from the Hogeschool Eindhoven. Thereafter, he studied at the Universiteit Maastricht, where he received his M.Sc. in Cognitive Psychology in 2001. Throughout these years Michel van Dartel also concentrated on a skateboarding career, achieving sponsorship and coverage in various media. However, in 2001 he decided to focus on his scientific career by accepting a position as a Ph.D. researcher at the Institute of Knowledge and Agent Technology (IKAT) of the Universiteit Maastricht. Besides computer science, psychology, and skateboarding, his main interests are music and art. With his current music project 'The Rose Frustrates' he released several recordings. His artwork is strongly influenced by his scientific work, and has been exposed in two group exhibitions. His homepage is www.michelvandartel.nl.

SIKS Dissertation Series

1998

- 1 Johan van den Akker (CWI²) *DEGAS - An Active, Temporal Database of Autonomous Objects*
- 2 Floris Wiesman (UM) *Information Retrieval by Graphically Browsing Meta-Information*
- 3 Ans Steuten (TUD) *A Contribution to the Linguistic Analysis of Business Conversations within the Language/Action Perspective*
- 4 Dennis Breuker (UM) *Memory versus Search in Games*
- 5 Eduard W. Oskamp (RUL) *Computerondersteuning bij Straftoemeting*

1999

- 1 Mark Sloof (VU) *Physiology of Quality Change Modelling; Automated Modelling of Quality Change of Agricultural Products*
- 2 Rob Potharst (EUR) *Classification using Decision Trees and Neural Nets*
- 3 Don Beal (UM) *The Nature of Minimax Search*
- 4 Jacques Penders (UM) *The Practical Art of Moving Physical Objects*
- 5 Aldo de Moor (KUB) *Empowering Communities: A Method for the Legitimate User-Driven Specification of Network Information Systems*
- 6 Niek J.E. Wijngaards (VU) *Re-Design of Compositional Systems*
- 7 David Spelt (UT) *Verification Support for Object Database Design*
- 8 Jacques H.J. Lenting (UM) *Informed Gambling: Conception and Analysis of a Multi-Agent Mechanism for Discrete Reallocation*

2000

- 1 Frank Niessink (VU) *Perspectives on Improving Software Maintenance*
- 2 Koen Holtman (TU/e) *Prototyping of CMS Storage Management*
- 3 Carolien M.T. Metselaar (UvA) *Sociaal-organisatorische Gevolgen van Kennistechnologie; een Procesbenadering en Actorperspectief*

²Abbreviations: SIKS – Dutch Research School for Information and Knowledge Systems; CWI – Centrum voor Wiskunde en Informatica, Amsterdam; EUR – Erasmus Universiteit, Rotterdam; KUB – Katholieke Universiteit Brabant, Tilburg; KUN – Katholieke Universiteit Nijmegen; RUL – Rijksuniversiteit Leiden; TUD – Technische Universiteit Delft; TU/e – Technische Universiteit Eindhoven; UL – Universiteit Leiden; UM – Universiteit Maastricht; UT – Universiteit Twente, Enschede; UU – Universiteit Utrecht; UvA – Universiteit van Amsterdam; UvT – Universiteit van Tilburg; VU – Vrije Universiteit, Amsterdam.

- 4 Geert de Haan (VU) *ETAG, A Formal Model of Competence Knowledge for User Interface Design*
- 5 Ruud van der Pol (UM) *Knowledge-Based Query Formulation in Information Retrieval*
- 6 Rogier van Eijk (UU) *Programming Languages for Agent Communication*
- 7 Niels Peek (UU) *Decision-Theoretic Planning of Clinical Patient Management*
- 8 Veerle Coupé (EUR) *Sensitivity Analysis of Decision-Theoretic Networks*
- 9 Florian Waas (CWI) *Principles of Probabilistic Query Optimization*
- 10 Niels Nes (CWI) *Image Database Management System Design Considerations, Algorithms and Architecture*
- 11 Jonas Karlsson (CWI) *Scalable Distributed Data Structures for Database Management*

2001

- 1 Silja Renooij (UU) *Qualitative Approaches to Quantifying Probabilistic Networks*
- 2 Koen Hindriks (UU) *Agent Programming Languages: Programming with Mental Models*
- 3 Maarten van Someren (UvA) *Learning as Problem Solving*
- 4 Evgueni Smirnov (UM) *Conjunctive and Disjunctive Version Spaces with Instance-Based Boundary Sets*
- 5 Jacco van Ossenbruggen (VU) *Processing Structured Hypermedia: A Matter of Style*
- 6 Martijn van Welie (VU) *Task-Based User Interface Design*
- 7 Bastiaan Schonhage (VU) *Diva: Architectural Perspectives on Information Visualization*
- 8 Pascal van Eck (VU) *A Compositional Semantic Structure for Multi-Agent Systems Dynamics*
- 9 Pieter Jan 't Hoen (RUL) *Towards Distributed Development of Large Object-Oriented Models, Views of Packages as Classes*
- 10 Maarten Sierhuis (UvA) *Modeling and Simulating Work Practice BRAHMS: a Multiagent Modeling and Simulation Language for Work Practice Analysis and Design*
- 11 Tom M. van Engers (VU) *Knowledge Management: The Role of Mental Models in Business Systems Design*

2002

- 1 Nico Lassing (VU) *Architecture-Level Modifiability Analysis*
- 2 Roelof van Zwol (UT) *Modelling and Searching Web-based Document Collections*
- 3 Henk Ernst Blok (UT) *Database Optimization Aspects for Information Retrieval*
- 4 Juan Roberto Castelo Valdueza (UU) *The Discrete Acyclic Digraph Markov Model in Data Mining*
- 5 Radu Serban (VU) *The Private Cyberspace Modeling Electronic Environments Inhabited by Privacy-Concerned Agents*
- 6 Laurens Mommers (UL) *Applied Legal Epistemology; Building a Knowledge-based Ontology of the Legal Domain*
- 7 Peter Boncz (CWI) *Monet: A Next-Generation DBMS Kernel For Query-Intensive Applications*

- 8 Jaap Gordijn (VU) *Value Based Requirements Engineering: Exploring Innovative E-Commerce Ideas*
- 9 Willem-Jan van den Heuvel (KUB) *Integrating Modern Business Applications with Objectified Legacy Systems*
- 10 Brian Sheppard (UM) *Towards Perfect Play of Scrabble*
- 11 Wouter C.A. Wijngaards (VU) *Agent Based Modelling of Dynamics: Biological and Organisational Applications*
- 12 Albrecht Schmidt (UvA) *Processing XML in Database Systems*
- 13 Hongjing Wu (TU/e) *A Reference Architecture for Adaptive Hypermedia Applications*
- 14 Wieke de Vries (UU) *Agent Interaction: Abstract Approaches to Modelling, Programming and Verifying Multi-Agent Systems*
- 15 Rik Eshuis (UT) *Semantics and Verification of UML Activity Diagrams for Workflow Modelling*
- 16 Pieter van Langen (VU) *The Anatomy of Design: Foundations, Models and Applications*
- 17 Stefan Manegold (UvA) *Understanding, Modeling, and Improving Main-Memory Database Performance*

2003

- 1 Heiner Stuckenschmidt (VU) *Ontology-Based Information Sharing in Weakly Structured Environments*
- 2 Jan Broersen (VU) *Modal Action Logics for Reasoning About Reactive Systems*
- 3 Martijn Schuemie (TUD) *Human-Computer Interaction and Presence in Virtual Reality Exposure Therapy*
- 4 Petkovic (UT) *Content-Based Video Retrieval Supported by Database Technology*
- 5 Jos Lehmann (UvA) *Causation in Artificial Intelligence and Law – A Modelling Approach*
- 6 Boris van Schooten (UT) *Development and Specification of Virtual Environments*
- 7 Machiel Jansen (UvA) *Formal Explorations of Knowledge Intensive Tasks*
- 8 Yong-Ping Ran (UM) *Repair-Based Scheduling*
- 9 Rens Kortmann (UM) *The Resolution of Visually Guided Behaviour*
- 10 Andreas Lincke (UT) *Electronic Business Negotiation: Some Experimental Studies on the Interaction between Medium, Innovation Context and Cult*
- 11 Simon Keizer (UT) *Reasoning under Uncertainty in Natural Language Dialogue using Bayesian Networks*
- 12 Roeland Ordelman (UT) *Dutch Speech Recognition in Multimedia Information Retrieval*
- 13 Jeroen Donkers (UM) *Nosce Hostem – Searching with Opponent Models*
- 14 Stijn Hoppenbrouwers (KUN) *Freezing Language: Conceptualisation Processes across ICT-Supported Organisations*
- 15 Mathijs de Weerdt (TUD) *Plan Merging in Multi-Agent Systems*
- 16 Menzo Windhouwer (CWI) *Feature Grammar Systems - Incremental Maintenance of Indexes to Digital Media Warehouse*

- 17 David Jansen (UT) *Extensions of Statecharts with Probability, Time, and Stochastic Timing*
- 18 Levente Kocsis (UM) *Learning Search Decisions*

2004

- 1 Virginia Dignum (UU) *A Model for Organizational Interaction: Based on Agents, Founded in Logic*
- 2 Lai Xu (UvT) *Monitoring Multi-party Contracts for E-business*
- 3 Perry Groot (VU) *A Theoretical and Empirical Analysis of Approximation in Symbolic Problem Solving*
- 4 Chris van Aart (UvA) *Organizational Principles for Multi-Agent Architectures*
- 5 Viara Popova (EUR) *Knowledge Discovery and Monotonicity*
- 6 Bart-Jan Hommes (TUD) *The Evaluation of Business Process Modeling Techniques*
- 7 Elise Boltjes (UM) *Voorbeeld_{IG} Onderwijs; Voorbeeldgestuurd Onderwijs, een Opstap naar Abstract Denken, vooral voor Meisjes*
- 8 Joop Verbeek (UM) *Politie en de Nieuwe Internationale Informatiemarkt, Grensregionale Politie Gegevensuitwisseling en Digitale Expertise*
- 9 Martin Caminada (VU) *For the Sake of the Argument; Explorations into Argument-based Reasoning*
- 10 Suzanne Kabel (UvA) *Knowledge-rich Indexing of Learning-objects*
- 11 Michel Klein (VU) *Change Management for Distributed Ontologies*
- 12 The Duy Bui (UT) *Creating Emotions and Facial Expressions for Embodied Agents*
- 13 Wojciech Jamroga (UT) *Using Multiple Models of Reality: On Agents who Know how to Play*
- 14 Paul Harrenstein (UU) *Logic in Conflict. Logical Explorations in Strategic Equilibrium*
- 15 Arno Knobbe (UU) *Multi-Relational Data Mining*
- 16 Federico Divina (VU) *Hybrid Genetic Relational Search for Inductive Learning*
- 17 Mark Winands (UM) *Informed Search in Complex Games*
- 18 Vania Bessa Machado (UvA) *Supporting the Construction of Qualitative Knowledge Models*
- 19 Thijs Westerveld (UT) *Using generative probabilistic models for multimedia retrieval*
- 20 Madelon Evers (Nyenrode) *Learning from Design: facilitating multidisciplinary design teams*

2005

- 1 Floor Verdenius (UvA) *Methodological Aspects of Designing Induction-Based Applications*
- 2 Erik van der Werf (UM) *AI techniques for the game of Go*
- 3 Franc Grootjen (RUN) *A Pragmatic Approach to the Conceptualisation of Language*
- 4 Nirvana Meratnia (UT) *Towards Database Support for Moving Object data*
- 5 Gabriel Infante-Lopez (UVA) *Two-Level Probabilistic Grammars for Natural Language Parsing*
- 6 Pieter Spronck (UM) *Adaptive Game AI*

- 7 Flavius Frasinca (TUE) *Hypermedia Presentation Generation for Semantic Web Information Systems*
- 8 Richard Vdovjak (TUE) *A Model-driven Approach for Building Distributed Ontology-based Web Applications*
- 9 Jeen Broekstra (VU) *Storage, Querying and Inferencing for Semantic Web Languages*
- 10 Anders Bouwer (UVA) *Explaining Behaviour: Using Qualitative Simulation in Interactive Learning Environments*
- 11 Elth Ogston (VU) *Agent Based Matchmaking and Clustering - A Decentralized Approach to Search*
- 12 Csaba Boer (EUR) *Distributed Simulation in Industry*
- 13 Fred Hamburg (UL) *Een Computermodel voor het Ondersteunen van Euthanasiebeslissingen*
- 14 Borys Omelayenko (VU) *Web-Service configuration on the Semantic Web; Exploring how semantics meets pragmatics*
- 15 Tibor Bosse (VU) *Analysis of the Dynamics of Cognitive Processes*
- 16 Joris Graaumanns (UU) *Usability of XML Query Languages*
- 17 Boris Shishkov (TUD) *Software Specification Based on Re-usable Business Components*
- 18 Danielle Sent (UU) *Test-selection strategies for probabilistic networks*
- 19 Michel van Dartel (UM) *Situated Representation*