

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/2388403>

# Embodied Cognitive Science

ARTICLE · APRIL 2001

Source: CiteSeer

---

CITATIONS

131

---

READS

53

## 1 AUTHOR:



[Rens Kortmann](#)

Delft University of Technology

23 PUBLICATIONS 196 CITATIONS

SEE PROFILE

# Embodied cognitive science

Rens Kortmann

Universiteit Maastricht, IKAT  
Neural networks and adaptive behaviour group  
P.O. Box 616  
NL-6200 MD Maastricht  
The Netherlands  
`kortmann@cs.unimaas.nl`

## Abstract

The paper provides an introduction to the field of embodied cognitive science from a biological and behavioural perspective. We show how the field of neuro-ethology can help transform cognitive science from a representational to an embodied perspective. The transformation is necessary to introduce a bottom-up approach to understanding cognition in order to resolve some fundamental problems with classical cognitive science. We give examples of current research by which we characterise the key idea of embodied cognitive science: study cognition by starting with the low-level behaviour of simple animals.

## 1 Introduction

Embodied cognitive science studies how *complete agents* cope with the challenges of their environment (Clark, 1996; Pfeifer and Scheier, 1999). Complete agents are natural or artificial systems (animals or machines) that possess a body (Varela, Thompson, and Rosch, 1991) and are situated (Clancey, 1997) in their environment, i.e., they perceive the environment only through their own sensors. Moreover, complete agents possess characteristics such as autonomy, self-sufficiency, and adaptivity (Pfeifer, 1996). The ultimate aim of embodied cognitive science is to understand how high-level cognitive processes such as reasoning and language arose from low-level interactions with the environment such as object manipulation and corrective eye movements. The main motivation for this approach is the argument that during natural evolution, high-level cognitive capacities and the associated brain areas (neo-cortex in humans) developed from low-level sensori-motor couplings in the deep parts of the brain (e.g., the limbic system in humans) (Kalat, 2001).

Embodied cognitive science, therefore, contrasts traditional cognitive science (e.g., Stillings *et al.*, 1995) and classical artificial intelligence (e.g., Winston, 1977). The latter fields aim at understanding and synthesising high-level cognition using a computational theory of mind (Pylyshyn, 1984; Sterelny, 1990), i.e., interpreting human mental processes as computations on symbolic representations.

Embodied cognitive science builds upon the large scientific heritage of fields such as neuro-ethology (Guthrie, 1980; Camhi, 1984; Hoyle, 1984), cybernetics (Wiener, 1948), and ecological psychology (Gibson, 1979). These areas of research share two features in their scientific method: the bottom-up approach and comparative study.

In this paper, we identify the relation of embodied cognitive science and neuro-ethology. The following sections first treat the essence of traditional cognitive science and classical artificial intelligence (section 2) followed by the pith of neuro-ethology (section 3). We indicate that traditional cognitive science ignores particular aspects of intelligent behaviour, which leads to fundamental problems with the field. Furthermore, neuro-ethology focusses exactly on the aspects that are disregarded by traditional cognitive science (Beer, 1990). In section 4, we show how the research questions of neuro-ethology help transform traditional cognitive science into embodied cognitive science, which *does* address the issues disregarded by traditional cognitive scientists. In addition, we elucidate how embodied cognitive science relieves some of the problems faced by neuro-ethologists. In section 5 we discuss the scope and limitations of embodied cognitive science. Finally, we draw conclusions in section 6.

## **2 Traditional cognitive science and classical artificial intelligence**

Cognitive science is traditionally a science of the intelligent mind (Stillings *et al.*, 1995). Some cognitive scientists restrict their focus to human intelligence, whereas others also include the abstract theory of intelligent processes and computer intelligence (Simon and Kaplan, 1989). In the 1960s, psychologists, philosophers, computer scientists, linguists, and neuroscientists joined forces in order to understand complex mental tasks such as thinking, remembering, and language (Gardner, 1985). For example, the interpretation of speech can be studied from a neuroscientific perspective (sound reception and neural processing in the brain), a linguistic perspective (parsing, word morphology, etc.), a computer science perspective (automatic speech recognition, natural language processing), a philosophical perspective (e.g., semantics), and a psychological perspective (intention of a speech act, etc.).

Pivotal in the approach of traditional cognitive science is the role of

symbolic representations. Pylyshyn (1984) states that the distinctive feature of *cognisers* – the subjects of study for cognitive scientists – is that their behaviour is based on representations. The idea is based on the representational theory of mind (Fodor, 1981; Sterelny, 1990) that states that all mental processes can be described in terms of computations on symbolic representations. In other words, mental activity is equivalent to the execution of an algorithm. For example, traditional cognitive science describes a person’s reaction to a visual stimulus as follows. First, the mind transforms the incoming image into some symbolic input representation. The form of this representation could be a logical proposition, a feature vector, or some other description. Second, the input representation is compared to representations stored in memory. Third, from the outcomes of the comparison an output representation is created that contains the appropriate motor actions. Notice that in the traditional approach to cognitive science, the questions of how an input representation is built and of how an output representation is executed by the motor system, is often ignored. The main focus lies on the algorithmic processing in the middle.

In the 1930s, the conception of mental activity as computations on symbolic representations was supported when Alan Turing presented the universal Turing machine – a hypothetical machine that can compute any possible mathematical function (Haugeland, 1985). From Turing’s idea, elaborated by von Neuman’s architecture for general-purpose computers (Stillings *et al.*, 1995), the field of artificial intelligence (AI) was originated in 1956 (Gardner, 1985). Together with the traditional cognitive scientists, early researchers in artificial intelligence envisioned the human mind as a machine performing computations on symbolic representations to solve problems (Newell and Simon, 1981). Moreover, given the fact that the Turing machine could compute any mathematical function, the researchers felt they were able to build a computer that emulated human-level intelligence.

The achievements of classical AI lie mainly in the domain of mathematical problem solving and reasoning with explicit knowledge. For instance, in 1997 the computer chess programme Deep Blue beat the world champion Kasparov. Moreover, knowledge-based decision support systems are now wide-spread in companies and other organisations. The applications of classical AI research outperform humans in particular when a computer’s processing speed and memory directly compete with that of a human brain. For instance, Deep Blue’s success is mainly due to the fact that it could compute the consequences of its moves quicker than Kasparov could. Also, logical decision making is better performed by computers that can handle vast numbers of rules as compared to humans experts that rely on intuition much more than on logical inference – e.g., legal decision making by computers is more consistent than by humans (Van den Herik, 1997).

In other domains, though, the performance of artificial systems does not compare to that of natural systems (animals) by far. These domains

typically are related to the real world and an artificial system's interaction with that real world (Brooks, 1990; Wilson, 1991; Clark, 1996; Clancey, 1997; Arkin, 1998; Pfeifer and Scheier, 1999). For instance, the mobile robot Shakey was built to navigate in a laboratory environment using a camera and a symbolic planning system (Nilsson, 1984). The robot's behaviour is brittle and sensitive to noise in the environment (Brooks, 1986; Arkin, 1998). The problems mentioned here are captured by two important theoretical issues regarding the symbolic representation of the real world: the frame problem (McCarthy and Hayes, 1969; Dennett, 1984) and the symbol grounding problem (Harnad, 1990). A thorough treatment of the problems reaches beyond the scope of this paper. We therefore restrict ourselves to giving a brief description in order to indicate the problems faced by traditional cognitive science and classical AI when applying the representational theory of mind to the real-world domain. For a full discussion of the problems we refer to the literature cited above.

The frame problem addresses the difficulties with representing change. If an intelligent system represents its environment symbolically, then how does it update the representation over time? This might not seem a problem for an abstract domain such as chess: there is a finite number of board positions and chess pieces. In contrast, changes in the real world are so fast and abundant that it is hardly plausible that intelligent systems (animals or machines) could efficiently use symbolic representations of the world for all of their behavioural patterns (Brooks, 1991; Kirsh, 1991). The symbol grounding problem discusses how symbolic representations relate to the real world. Again, in highly abstract domains this problem is less significant as such a system needs not know the meaning of the symbols it uses: a decision-support system need not know the meaning of the rules it operates on in order to deduce a valid conclusion from the input it is supplied with. However, a system operating autonomously in a real-world needs to know the relation between real-world objects (food, predators) and the system's symbolic representations of those objects: the meaning of symbols must be grounded in the system's own interaction with the real world (Pfeifer and Scheier, 1999). Gibson (1979) discusses the grounding of classification tasks in terms of affordances: food can be eaten, predators can be escaped from. Traditional cognitive science focusses on high-level cognitive capacities, not on low-level interactions with the real world. Therefore, for Shakey a goal position is a meaningless goal position, supplied by a human experimenter.

From the previous treatment of classical AI it becomes clear that explanations of intelligent behaviour in terms of the processing of symbolic representations might be appropriate in abstract domains, such as mathematics and logic – in more dynamic domains found in the natural world, the frame problem and symbol grounding problem point out to the need of increased understanding of an intelligent system's interaction with the real world. In other words: “understanding the nature of cognition requires

considering more than the complex problem solving and learning of human experts and their tutees.” (Clancey, 1997, p. 6). In the next section we provide an introduction to a scientific field in which the interaction of animals with the environment is cardinal. It will be shown that this field in fact helps transform cognitive science from a representational to an embodied perspective.

### 3 Neuro-ethology

Neuro-ethology is the study of natural behaviour and the neural mechanisms mediating it (Guthrie, 1980; Camhi, 1984; Hoyle, 1984). The field originated from two other biological disciplines: neuro-biology and ethology that, by themselves, differ significantly in research interest and methods. Neuro-biology investigates the workings of neural structures in animals under controlled circumstances (Kandel, 1976). The experimental set-up allows neuro-biologist to obtain strict stimulus-response characteristics of a brain area. As a drawback, the laboratory conditions might cause the test animals to behave in a non-natural way which in turn places a bias on the experimental results (Huber, Franz, and Bühlhoff, 1998). In contrast, ethologists study the behaviour of animals in their natural habitat (field experiments) (Slater, 1985). The approach usually guarantees the display of natural behaviour by the animals, but makes brain recordings virtually impossible. As a result, ethological research usually yields no results beyond the level of descriptions of natural behaviour (Camhi, 1984).

Neuro-ethology aims at combining the explanatory power of neuro-biology with the ecological validity of ethological research. In order to reach the aim, neuro-ethologists first observe the behaviour of an animal in its natural habitat and derive problems faced by the animal’s neural system to produce the behaviour. These problems are subsequently studied in neuro-biological experiments under as natural circumstances as possible. Performing the experiments usually involves returning to the first phase (observing natural behaviour) when the neuro-biological results raise questions on the behavioural studies. Therefore, an iterative process (in which behavioural and neuro-biological experiments alternate) leads to the resolution of research questions that include the following topics: signal detection and recognition (e.g., calling songs or olfactory trails); co-ordination (e.g., in flight or when walking); localisation (sound sources, landmarks, etc.); and orientation (e.g., in navigation).

The approach is usually centred around a particular common behaviour faced by many animals, such as finding a mate, or avoiding predation. Subsequently, a target animal is chosen to study the behaviour and its neural underpinnings (Camhi, 1984). Although the particular details of neural mechanisms for similar behaviours can vary significantly amongst different

animal species, the nature of neural signalling is common in most animals and many species are constrained by the same physical limitations of their bodies and of their environment. For instance, different insect species have compound eyes and share the same habitat in which they display similar behaviours, such as feeding or target pursuit. Studying the neural mechanism mediating behaviour in one animal species, therefore is likely to reveal insight into the mechanisms of other species. In other words, neuro-ethology relies on comparative study.

Next to the term ‘comparative study’ also the phrase ‘bottom-up approach’ is central to the field of neuro-ethology. The phrase indicates that through studying relatively simple behaviours, such as orientation or localisation, neuro-ethologists eventually aim at explaining high-level behaviour such as decision making and the orchestration of highly complex behavioural patterns. We notice this approach contrasts the ‘top-down’ approach adopted by traditional cognitive scientists and researchers of classical AI. Whereas neuro-ethologists aim at understanding complex behaviours in terms of low-level interactions of an animal with its environment, the representations-focused cognitive science and AI researchers study high-level behaviour directly, while disregarding the underlying low-level behaviours that support the cognitive processes. As an example we mention the research on fear by LeDoux (1996) and co-workers. In the human brain he identified a cortical and sub-cortical pathway mediating fear-responses. When presented with a fearsome stimulus, a subject’s cortical pathway elicits the awareness of fear and is therefore associated with high-level cognition. At the same time, the much faster sub-cortical pathway triggers a non-conscious fight-or-flight reaction which is associated with a low-level reactive response. The low-level behaviour can be overruled by the high-level, cognitive behaviour, but can never be replaced by it: the cognitive high road is too slow to trigger a timely response to a fearsome stimulus. In summary, a low-level, non-cognitive brain area is fundamental to a fear-response, even in humans. This observation supports a bottom-up approach to intelligent behaviour, as adopted in neuro-ethology.

**Limitations of the neuro-ethological paradigm** As was mentioned before, neuro-biological and thus neuro-ethological experiments need a controlled laboratory set-up in order to allow for meaningful measurements in the animal’s brain. Using the existing techniques it is still impossible to measure neural activity in freely moving animals. Instead, to obtain valid stimulus-response characteristics, the animals that are investigated are usually placed in open-loop conditions, i.e., they are fixated and thus do not receive feedback from their actions. For instance, the activity of movement-sensitive neurons in flies is measured while the animal is fixated in a hollow tube with wax (Mastebroek, Zaagman, and Lenting, 1980). While neuro-

ethologists put a lot of effort in emulating as natural experimental conditions as possible, still there exists a fundamental problem: it is unclear whether the behaviour of animals and its neural activity is similar in open-loop and closed-loop conditions, i.e., when freely moving or not.

Moreover, neuro-ethological research can not study multiple behaviours in an animal simultaneously: again, in order to obtain valid stimulus-response characteristics, complex behaviours are decomposed into simpler behavioural patterns such as orientation or flight course stabilisation that are studied separately (Camhi, 1984). The synthesis of neural activity for the complex behaviour from the components is not a straightforward procedure, i.e., neural mechanisms cannot simply be added to account for the explanation of complex behaviour (Huber *et al.*, 1998). Another potential risk in neuro-ethology is assuming the existence of different neural mechanisms for different behavioural patterns, whereas one neural mechanism can mediate several behaviours. For instance, the phonotactic behaviour (sound-seeking) of crickets consists of a sound recognition and localisation component. Webb (1995) found that both behaviours are likely to be mediated by the same neural structure (see also section 4).

In summary, neuro-ethological research faces two methodological problems: lack of ecological validity due to closed-loop experiments and the decomposition of natural behaviour. We propose that embodied cognitive science can relieve the problems by employing robots to model aspects of animal behaviour. Robots can be used in closed-loop conditions and the modelling aspect allows for the synthesis of behaviours.

## 4 Embodied cognitive science

From the previous sections we conclude that neuro-ethology treats some of the shortcomings of traditional cognitive science. Section 2 showed the focus of traditional cognitive science and classical AI which was to describe and synthesise high-level cognitive processes through computations on symbolic representations. Their major shortcoming is the disregard of low-level behaviour that does not require symbolic representations in its explanation. Section 3 gave an introduction to the field of neuro-ethology that fixates on the neural mechanisms underlying low-level behaviour in animals. Drawbacks of the latter field include artificial experimental set-ups and the decomposition of behaviour.

We argue that embodied cognitive science can combine the best of both worlds. First, employing a neuro-ethological focus on low-level behaviour provides a better understanding of how cognisers (Pylyshyn, 1984) interact with the real-world and how they use their low-level behaviour as a foundation for high-level cognitive processes. Second, the modelling techniques from AI can be used to solve parts of the problems faced by neuro-ethology.



In the following, we shall elaborate on the arguments mentioned and by giving examples of current work in embodied cognitive science and the achievements.

Pfeifer and Scheier (1997) investigated a classification task in terms of sensori-motor co-ordination. Traditionally, classification is viewed as the mapping of an instance to a symbolic class representation (Estes, 1994). For instance, when seeing a particular animal, the viewer creates a symbolic representation of the animal and maps it to the class ‘dogs’. Alternatively, Pfeifer and Scheier adopt a more ‘Gibsonian’ perspective on the task. They interpret classification of objects in terms of the possible actions an agent can perform with them: graspable objects vs non-graspable objects, pushable objects vs non-pushable objects, etc. In an experiment a mobile robot was placed in an arena with cylindrical objects from two classes: objects with either a large or a small diameter. When the robot met an object, it started to circle around the object and measured the angular velocity. Moreover, the robot possessed a ‘tail’ by which it could grasp the objects with small diameter, but not the large ones. Using Hebbian learning, a neural network was trained that decided for the robot to grasp an object. The decision was based on the angular velocity that was measured while circling the object. In this approach, sensori-motor patterns of activity (time-series of sensor and motor values), rather than symbols, constitute class representations. The approach fits better to the way children and animals learn to classify objects than the classical approach does (Smith, 1994). Moreover, the approach describes how an agent (animal or robot) grounds the meaning of an object class in its behaviour.

As a second example, we mention the work by Webb (1995; 1998) on cricket phonotaxis. The work proposes a neural mechanism of female crickets’ classification and localisation of male crickets’ calling songs. The mechanism was implemented on a robot model and yielded robust behaviour comparable to that of real crickets. Again, the classification task (distinguishing between different calling songs, belonging to males of different cricket species) was based on a behavioural response (approaching the conspecific male), not on a symbolic classification. The work gives another hint that symbolic representations are not always necessary to explain high-level cognitive tasks (categorisation). Moreover, Webb’s approach to studying cricket phonotaxis adds an important aspect to cognitive science and neuro-ethology: robotic modelling. The approach relieves two problems of neuro-ethology: first, robots can be used in closed-loop conditions, i.e., it is possible to record the activity of the artificial neural mechanism that controls the robot, while the robot can move freely in a desirable environment; second, robots can be employed to model the orchestration of behaviours. Again, activity in the robot’s artificial neural system can be recorded while the robot performs various behaviours in parallel. For a more complete treatment on using robots to model animals we refer to the work of Webb

(1995; 1998; 1999).

Next we shall mention a few examples of robot models of animal behaviour and the neural mechanisms underlying it. Some of the models were implemented in a physical robot, whereas others were built as computer simulations. Already in 1982 Arbib modelled visuo-motor co-ordination in frogs and toads in a system called *Rana computatrix*. Some ten years later Cliff (1991) modelled visually-guided behaviour of the hoverfly *Syrirtta pipiens* in a computer simulator called SYCO (*Syrirtta computatrix*). Other animals that were investigated include cockroaches (Beer and Chiel, 1993), lampreys (Ijspeert, Hallam, and Willshaw, 1999), salamanders (Ijspeert and Arbib, 2000), more flies (Franceschini, Pichon, and Blanes, 1992; Huber *et al.*, 1998), ants (Lambrinos *et al.*, 2000), and bats (Peremans *et al.*, 2000).

**Applications** Besides using robots in order to investigate animals, the field of embodied cognitive science has produced a wide range of applications by taking inspiration from biology. For instance, the field of behaviour-based robotics (Arkin, 1998) employs principles of animal behaviour to the design of autonomous systems (cars, wheelchairs, space rovers). The field of evolutionary robotics (Harvey *et al.*, 1997; Nolfi and Floreano, 2000) applies the theory of evolution to the development of robot controllers. Finally we mention the use of optic flow in visually-guided behaviour in robots (Weber, Venkatesh, and Srinivasan, 1997; Srinivasan *et al.*, 1997).

## 5 Discussion

In this paper we have discussed important aspects of traditional cognitive science, classical artificial intelligence and neuro-ethology. We pointed out to a few problems relating to the fields and showed that embodied cognitive science can relieve some of the problems through employing robot modelling. In fact, Clancey (1997) views the researchers of embodied cognitive science as a new generation of cyberneticists (Wiener, 1948) to stress the importance of the perception-action loop in embodied cognitive science research. Cybernetics studied the control of low-level behaviour in animals and machines already in the 1940s and 1950s. However, the artificial systems built by early cyberneticists (e.g., Walter, 1950) used simple controllers based on analogue hardware. In contrast, with present-day computing power, the use of neural network models in robots is facilitated.

Converting from a representational to an embodied perspective on intelligence requires some adaptive power. Modesty is called for when adopting a bottom-up approach to cognition, but people will be astounded by the ingenuity of the neural mechanisms underlying low-level behaviour in humans and other animals.

**Limitations** Using artificial systems (robots or computer simulations) to study real animals has some limitations of which we mention three. First, the role of noise is different in natural and artificial systems. For instance, photon absorption noise is Poisson distributed (Land, 1981), which does not necessarily have to be the case in artificial visual sensors. Second, the motor system of robots is usually extremely simplified which results in a difference between the sensory feedback experienced by robots and animals moving in the same environment. If, however, work focusses on the motor behaviour instead, the sensory processing is usually limited (Beer and Chiel, 1993; Ijspeert and Arbib, 2000). Third, an animal might use sensory input or environmental clues that are not modelled in the artificial system used. For instance, the role of the ultra-violet component of sunlight plays an important role in insect navigation (Lambrinos *et al.*, 2000). In a similar way, other sensory channels that we are currently unaware of might be of importance for the behaviour. Not implementing these channels in a robot model might lead to wrong explanations of the behaviour. This problem is even larger when computer simulations are used instead of physical robots. Computer simulations might represent a too simple environment and miss important sensory information. Instead, physical robots interact with the same environment as real animals and therefore are theoretically able to receive all sensory information available to real animals. The term ‘comparative study’ therefore also applies to embodied cognitive science.

All of the limitations mentioned above are general problems of the scientific method of modelling – they do not apply just to embodied cognitive science. Modelling involves abstracting away from the real system under investigation. All simplifications in the design of the model should be accounted for. For instance, using two wheels to model the motor system of an animal can be justified when yaw behaviour (rotation in the horizontal plane) is studied. Moreover, modelling is usually performed in an incremental fashion (adding more and more components to the model). In fact, the limitations could serve as the explanatory power of embodied cognitive science. For instance, when reproducing behaviour in a robot fails with the neural mechanism assumed to mediate the behaviour in real animals, this might indicate that more sensory channels are involved, or that the role of the actuators is more important than expected.

## 6 Conclusions

Altogether, we showed that neuro-ethology helps transform the perspective on cognitive science from one of representations to one of embodiment and situatedness. The method of applying robots to model animal behaviour and the neural mechanisms underlying the behaviour relieves some central methodological problems in traditional cognitive science and neuro-ethology.

First we noticed that traditional cognitive science and classical artificial intelligence underestimate the role of low-level real-world interaction in producing intelligent behaviour, whereas embodied cognitive science adopts a bottom-up approach starting with the low-level aspects of behaviour. Second we stressed that the problems in neuro-ethology regarding recording neural activity from living animals can be prevailed by using robot models. The use of robot models allows for comparative study of animals species. Summarising, embodied cognitive science combines the best of two worlds: traditional cognitive science and neuro-ethology.

**Future work** In future, we think embodied cognitive science could contribute to the scientific community by focusing on the relation between high-level cognitive tasks and the underlying low-level behaviour supporting it. For instance, the use of eye movements (low-level behaviour) is pivotal in understanding how high-level tasks such as face recognition are accomplished (Ballard, 1991). The grounding of symbolic representations is another important topic. Embodied cognitive science requires a change of mind (bottom-up instead of top-down approach) but offers important new insights to the study of intelligence.

## Acknowledgements

I would like to thank Eric Postma for his advice and useful discussions.

## References

- Arbib, M. (1982). Rana Computatrix: An Evolving Model of Visuo-motor Coordination in Frog and Toad. *Machine Intelligence*, Vol. 10, pp. 501–517.
- Arkin, R. (1998). *Behavior-based robotics*. MIT Press, Cambridge.
- Ballard, D. (1991). Animate vision. *Artificial intelligence*, Vol. 48, pp. 57–86.
- Beer, R. and Chiel, H. (1993). Simulations of cockroach locomotion and escape. *Biological neural networks in invertebrate neuroethology and robotics* (eds. R. Beer, R. Ritzmann, and T. McKenna), pp. 267–285, Academic Press, Boston.
- Beer, R. (1990). *Intelligence as adaptive behavior, an experiment in computational neuroethology*. Academic Press, Boston.
- Brooks, R. (1986). A robust layered control system for a mobile robot. *IEEE journal of robotics and automation*, Vol. 2, pp. 14–23.

- Brooks, R. (1990). Elephants don't play chess. *Robotics and autonomous systems*, Vol. 6, pp. 3–15.
- Brooks, R. (1991). Intelligence without representation. *Artificial intelligence*, Vol. 47, pp. 139–159.
- Camhi, J. (1984). *Neuroethology*. Sinauer associates, Sunderland.
- Clancey, W. (1997). *Situated cognition: on human knowledge and computer representations*. Cambridge University Press, New York.
- Clark, A. (1996). *Being there: putting brain, body, and world together again*. MIT Press, Cambridge.
- Cliff, D. (1991). The Computational Hoverfly: A Study in Computational Neuroethology. *From Animals to Animats: Proceedings of the First International Conference on the Simulation of Adaptive Behaviour* (eds. J.-A. Meyer and S. Wilson), MIT Press, Cambridge.
- Dennett, D. (1984). Cognitive Wheels: The frame problem of AI. *Minds, Machines and Evolution* (ed. C. Hookway), pp. 129–152, Cambridge University Press, Cambridge.
- Estes, W. (1994). *Classification and cognition*. Oxford University Press, New York.
- Fodor, J. (1981). *Representations: philosophical essays on the foundations of cognitive science*. Harvester Press, Brighton.
- Franceschini, N., Pichon, J.-M., and Blanes, C. (1992). From insect vision to robot vision. *Philos. Trans. R. Soc. Lond. Biol.*, Vol. 337, No. 1281, pp. 283–294.
- Gardner, H. (1985). *The mind's new science*. Basic Books, New York.
- Gibson, J. (1979). *The ecological approach to perception*. Houghton Mifflin, Boston.
- Guthrie, D. (1980). *Neuroethology: an introduction*. Wiley, New York.
- Harnad, S. (1990). The symbol grounding problem. *Physica D*, Vol. 42, pp. 335–346.
- 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.
- Haugeland, J. (1985). *Artificial intelligence : the very idea*. MIT Press, Cambridge.

- Hoyle, G. (1984). The scope of neuroethology. *The Behavioural and Brain Sciences*, Vol. 7, pp. 367–412.
- Huber, S., Franz, M., and Bühlhoff, H. (1998). On robots and flies: modeling the visual orientation behavior of flies. Technical Report 56, Max-Planck-Institut für biologische Kybernetik, Tübingen, Germany.
- Ijspeert, A. and Arbib, M. (2000). Visual tracking in simulated salamander locomotion. *From Animals to Animats, Proceedings of the 6th International Conference on the Simulation of Adaptive Behavior* (eds. J. Meyer, A. Berthoz, D. Floreano, H. Roitblat, and S. Wilson), pp. 88–97, MIT Press, Cambridge.
- Ijspeert, A., Hallam, J., and Willshaw, D. (1999). Evolving swimming controllers for a simulated lamprey with inspiration from neurobiology. *Adaptive Behavior*, Vol. 7, No. 2, pp. 151–172.
- Kalat, J. (2001). *Biological psychology*. Brooks/Cole publishing company, Pacific Grove, 7th edition.
- Kandel, E. (1976). *Cellular basis of behavior*. W.H. Freeman and Co., San Francisco.
- Kirsh, D. (1991). Today the earwig, tomorrow man? *Artificial intelligence*, Vol. 47, pp. 161–184.
- Lambrinos, D., Möller, R., Labhart, T., Pfeifer, R., and Wehner, R. (2000). A mobile robot employing insect strategies for navigation. *Robotics and Autonomous Systems*, Vol. 30, pp. 39–64.
- Land, M. (1981). Optics and vision in invertebrates. *Handbook of sensory physiology* (ed. H. Autrum), Vol. VII/6B, pp. 471–592, Springer, Berlin.
- LeDoux, J. (1996). *The emotional brain : the mysterious underpinnings of emotional life*. Simon and Schuster, New York.
- Mastebroek, H., Zaagman, W., and Lenting, B. (1980). Movement detection: performance of a wide-field element in the visual system of the blowfly. *Vision research*, Vol. 20, pp. 467–474.
- McCarthy, J. and Hayes, P. (1969). Some philosophical problems from the standpoint of artificial intelligence. *Machine intelligence*, Vol. 4, pp. 463–502.
- Newell, A. and Simon, H. (1981). Computer science as empirical inquiry: symbols and search. *Mind design – philosophy, psychology, artificial intelligence* (ed. J. Haugeland), pp. 35–66, MIT Press, Cambridge.
- Nilsson, N. (1984). Shakey the robot. Technical Report 323, AI Center, SRI International, Menlo Park.

- Nolfi, S. and Floreano, D. (2000). *Evolutionary robotics – the biology, intelligence, and technology of self-organizing machines*. MIT Press, Cambridge.
- Peremans, H., Müller, R., Carmenta, J., and Hallam, J. (2000). A biomimetic platform to study perception in bats. *Proceedings of the SPIE international symposia on sensor fusion and decentralized control in robotic systems III*. In press.
- Pfeifer, R. and Scheier, C. (1997). Sensory-motor coordination: the metaphor and beyond. *Robotics and autonomous systems*, Vol. 20, pp. 157–178.
- Pfeifer, R. and Scheier, C. (1999). *Understanding intelligence*. MIT-press, Cambridge.
- Pfeifer, R. (1996). Building “Fungus Eaters”: Design principles of Autonomous Agents”. *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behaviour*, MIT Press, Cambridge.
- Pylyshyn, Z. (1984). *Computation and cognition: toward a foundation for cognitive science*. MIT Press, Cambridge.
- Simon, H. and Kaplan, C. (1989). Foundations of cognitive science. *Foundations of cognitive science* (ed. M. Posner), pp. 1–47, MIT Press, Cambridge.
- Slater, P. (1985). *An introduction to ethology*. Cambridge University Press, Cambridge.
- Smith, E. T. L. (1994). *A dynamic systems approach to the development of cognition and action*. MIT Press, Cambridge.
- Srinivasan, M., Chahl, J., Nagle, M., and Zhang, S. (1997). Embodying natural vision into machines. *From living eyes to seeing machines* (eds. M. Srinivasan and S. Venkatesh), pp. 249–265, Oxford University Press, Oxford.
- Sterelny, K. (1990). *The representational theory of mind: an introduction*. Blackwell, Oxford.
- Stillings, N., Weisler, S., Chase, C., Feinstein, M., Garfield, J., and Rissland, E. (1995). *Cognitive science – an introduction*. MIT Press, Cambridge, 2nd edition.
- Van den Herik, J. (1997). From chess moves to legal decisions: a position statement. *Proceedings of JURIX '97: the tenth conference on legal knowledge based systems* (ed. A. Oskamp), pp. 107–109, Gerard Noodt Instituut, Nijmegen.

- Varela, F., Thompson, E., and Rosch, E. (1991). *The embodied mind*. MIT Press, Cambridge.
- Walter, W. (1950). An imitation of life. *Scientific American*, Vol. 182, No. 5, pp. 42–45.
- Webb, B. (1995). Using robots to model animals: a cricket test. *Robotics and autonomous systems*, Vol. 16, pp. 117–134.
- Webb, B. (1998). Robots, crickets and ants: models of neural control of chemotaxis and phonotaxis. *Neural Networks*, Vol. 11, pp. 1479–1496.
- Webb, B. (1999). A framework for models of biological behaviour. *International journal of neural systems*, Vol. 9, No. 5, pp. 375–381.
- Weber, K., Venkatesh, S., and Srinivasan, M. (1997). Insect inspired behaviours for the autonomous control of mobile robots. *From living eyes to seeing machines* (eds. M. Srinivasan and S. Venkatesh), pp. 227–249, Oxford University Press, Oxford.
- Wiener, N. (1948). *Cybernetics, or control and communications the animal and the machine*. Wiley, New York.
- Wilson, S. (1991). The animat path to AI. *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior* (eds. J.-A. Meyer and S. Wilson), pp. 15–21, MIT Press, Cambridge.
- Winston, P. (1977). *Artificial intelligence*. Addison-Wesley, Reading.