

# Cognitive Developmental Robotics As a New Paradigm for the Design of Humanoid Robots

Minoru Asada<sup>1</sup> and Karl F. MacDorman<sup>2</sup> Hiroshi Ishiguro<sup>3</sup> Yasuo Kuniyoshi<sup>4</sup>

<sup>1</sup> Adaptive Machine Systems, Osaka University, Suita, Osaka 565-0871 Japan

<sup>2</sup> Department of Systems and Human Science, Osaka University, Toyonaka, Osaka 560-8531 Japan

<sup>3</sup> Computer and Communication Sciences, Wakayama University, Wakayama 640-8441 Japan

<sup>4</sup> Humanoid Interaction Laboratory, ETL, Tsukuba 305-8568, Japan

**Abstract.** This paper proposes *cognitive developmental robotics* as a new principle for the design of humanoid robots. This principle may provide ways of understanding human beings that go beyond the current level of explanation found in the natural and social sciences. Furthermore, a methodological emphasis on humanoid robots in the design of artificial creatures holds promise because they have many degrees of freedom and sense modalities and, thus, must face the challenges of scalability that are often side stepped in simpler domains. We examine the potential of this new principle as well as issues that are likely to be important to CDR in the future.

## 1 Introduction

Robot heroes and heroines in science fiction movies and cartoons like Star Wars in the US and Astor Boy in Japan have attracted us so much, which as a result has motivated many robotics researchers. These robots, unlike special purpose machines, are able to communicate with us and perform a variety of complex tasks in the real world. What do present-day robots lack that prevents them from realizing these abilities? We advocate a need for *cognitive developmental robotics* (CDR), which aims to understand the cognitive developmental processes that an intelligent robot would require and how to realize them in a physical entity. However, cognitive developmental robotics has just started, and its definition, design principle, and methodology have not yet been established. In this paper, we examine the potential for cognitive developmental robotics to further our understanding of both humans and machines. We hope this stimulates many researchers — not simply in robotics but also in other disciplines — to discuss and tackle this controversial new paradigm.

The key aspect of CDR is its design principle. Existing approaches often explicitly implement a control structure in the robot's 'brain' that was derived from a designer's understanding of the robot's physics. According to CDR, the structure should reflect the robot's own process of understanding through interactions with the environment. Since both CDR and the traditional approach may lead to similar results, CDR may seem unnecessary if we evaluate it merely in terms of task performance. However, we believe CDR holds promise in the long-run both in terms of producing humanlike behavior and because it can serve as a testbed for cognitive theories. Furthermore, traditional approaches in AI and engineering tend to break down in natural settings, where the robot's body and environment are difficult to model and can change unpredictably [1, ?].

Brooks et al. [2] proposed the methodology for alternative essences of intelligence as a humanoid design principle, which consists of parallel themes: development, social interaction, embodiment, and integration. Any of these themes seems essential to CDR, and we share very similar concepts. But in CDR we emphasize more fundamental issues of cognitive development and propose a more constructivist approach.

Cognition and development have been key issues for human intelligence, and recent progress in these disciplines promoted a new area called *developmental cognitive neuroscience* (DCN) [3], which emerged at the interface between two of the most fundamental questions that challenge mankind. The first one concerns the relation between mind and body, and especially between the physical substance of the brain and the mental processes it supports (cognitive neuroscience). The second concerns the origin of organized biological structure, such as the highly complex structure of the adult human brain (development). Johnson claimed that we can cast light on these two questions by focusing on the relation between the postnatal development of the human brain and the cognitive processes it supports.

The basic idea seems applicable to the approach of CDR since it has to deal with cognitive processes during the development of a robot's brain. However, the difference between CDR and DCN is that CDR

is a synthetic or constructivist approach with the potential to test its models by implementing them in humanoid robots. The cycle of fault diagnosis and reimplementing may iterate many times in order to refine the model [4]. The idea is that this process of refinement might result in a useful model of human interaction.

Since brain science is primarily concerned with structural details of human brains at the microscopic level, it may not be well suited to providing a comprehensive model of human activity or explaining how brains support it. Sometimes, however, the social sciences have attempted to understand human activities at a purely macroscopic level — without concern for the biological structure of individuals (for example, humans are sometimes treated as black boxes). By providing a means of scrutinizing and testing models and finding alternatives, CDR can help bridge macroscopic and microscopic approaches. We expect that, through the process of designing and implementing humanoid robots, a new way of understanding human beings will develop that differs significantly from the ways in which humans are understood in the natural and social sciences. In addition, we believe that robots that can emerge symbols through social interaction will have the best chance of one day approaching human capacity.

We already mentioned one side of the design principle of CDR: the design of self-developing structures inside the robot’s brain. But if we consider human beings and other intelligent species, we find another side to this story. Individuals cannot reach their full potential without nurturing relationships. Parents, teachers, and other adults adapt themselves to the needs of children according to each child’s level of maturity and the particular relationship they have developed with that child [5]. So the other side of CDR’s design principle concerns environmental design: how to set up the environment so that the robots embedded therein can gradually adapt themselves to more complex tasks in more dynamic situations. It may include instruction from a human or robot. Fig.1 shows a typical method of designing the embedded structure and the environment.

The rest of this article is organized as follows: First, we briefly examine the view that embodiment is the least requirement for cognitive development. Then we explain the design principles and the approaches of CDR. Finally, we discuss future directions for CDR.

## 2 Physical Embodiment and Interactions

Owing in part to the influence of a series of papers by Brooks (cf., [6, 7]), artificial intelligence researchers now consider physical embodiment to be necessary for designing the structure of intelligent systems. A physical body enables an agent to interact with its environment, which we may expect could lead to the emergence of intelligent behavior and internal organization. Robotics researchers have never really disputed the need for a physical body because it is essential to their research. Therefore, few have entered into a critical dialogue concerning the relationship between having a body and the emergence of intelligence. Here, we review the significance of embodiment [8].

1. Perception and action are not separable but tightly coupled.
2. Under resource-bounded conditions (memory, processing power, controller, etc.), an agent is able to learn a sensorimotor mapping from experience (interactions with the environment).
3. As the complexity of its task or the environment increases, the agent is able to adapt itself to these changes by learning from the consequences of its actions and adapting this knowledge to new situations.

No one seems to have any objection to 1 and maybe 2. Pfeifer and Scheier explained “embodiment” in a variety of contexts in their book [9] with reference to Brooks’s definition [6]. However, they seem to put more emphasis on physical coupling than cognitive and physical developmental processes. A typical example is passive dynamic walking [10] by exploiting the system’s dynamics. Here, we focus on *cognitive* (as opposed to physical) development by adding 2 and 3 since current technology does not really support the realization of a growing, changing body.

## 3 Cognitive Science, Developmental Psychology, Neuroscience and Cognitive Developmental Robotics

*Developmental Cognitive Neuroscience* [3] has emerged from cognitive science, developmental psychology, and neuroscience partly owing to the recent progress of imaging technology in brain science. A

fundamental controversy in cognitive science concerns the relative importance of nature and nurture in determining the structure and behavior of individuals. One extreme is that gene coding has all kinds of information necessary for development. The other extreme is that much of the information involved in the formation of a human mind comes from the environment. Both viewpoints are lacking. Neither the nature nor the nurture side address how new information emerges, as Johnson pointed out. In the last decade new evidence has revealed that complicated interactions between genes, developmental processes, and the environment lead to the emergence of structural organization and behavior at many levels [11].

The view that conventional robotics methods correspond to the nature because they embedded robot behaviors explicitly while CDR corresponds to the nurture because of its dependence on learning from interactions is misguided because neither the nature nor the nurture view explains how new forms of social interaction emerge. CDR aims at a constructivist approach to realizing a mechanism that can adapt to complicated and dynamic changes in the environment based on its capacity for interaction.

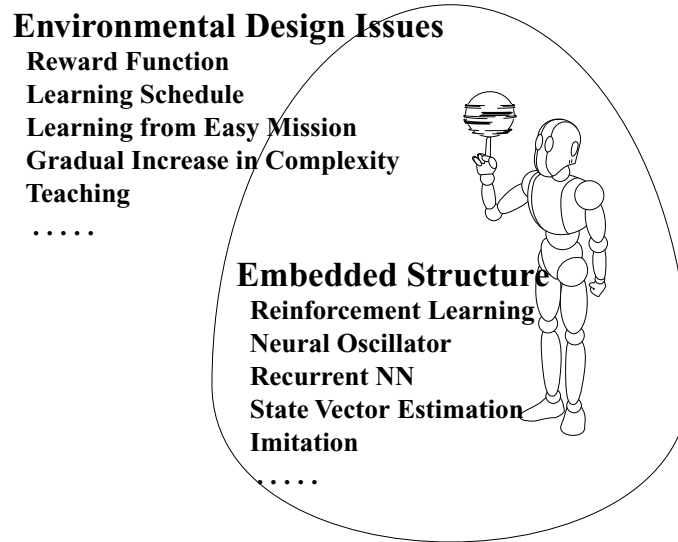


Fig. 1. The interaction between the embedded structure (e.g., robot brain) and the environment

## 4 The Design Principle of CDR

From the standpoint of engineering, there are two sides to the design principle of CDR: (1) how to design a robot brain whose embedded structure can learn and develop; and (2) how to create a social environment capable of supporting the development of cognitive processes.

### 4.1 Embedded structure

The embedded structure is a mechanism inside the robot that efficiently supports its interaction with an environment. The information obtained through interaction will differ qualitatively depending on the size and organization of the robot's functional modules, which may range from the neural level to larger units such as visual and motor subsystems. However, a common feature is that new information emerges inside the robot. Reinforcement learning, which maps from sensory information to actuator outputs, is a typical example of a functional module.

### 4.2 Environmental design issues

The conventional robot design principle has put much more emphasis on the embedded structure than on environmental issues, although the resulting behaviors seriously depend on both. Environmental design

issues are essential for a robot with embedded structure to learn and develop so that it can gradually adapt itself to more complicated environments. Environmental design issues include all kinds of factors that come from outside the robot. How other active agents respond is key to multi-agent learning, whether they be cooperative (e.g., rescue activities in a disaster situation), competitive (a prey surrounded by predator), or both (game situations such as RoboCup [12]). Furthermore, other agents can be coaches or teachers who can communicate with robots by various means. From the viewpoint of facilitating a robot's development, learning from easy missions (LEM) [13], a learning schedule [8], or a gradual increase in domain complexity [14] are typical approaches.

## 5 Approaches to CDR

Although a full scale implementation of a humanoid robot, built according to the principles of CDR, currently stands beyond our reach, for the time being we can focus on essential issues in CDR while keeping in mind the long-term goal.

### 5.1 Development

We have examined developmental issues within the reinforcement learning paradigm because reinforcement learning enables complex behavior to emerge through interaction without making many, often untenable assumptions about the structure and initial state of the internal mechanisms of cognition. Yet the flexibility of reinforcement learning has also been its weakness; it results in a huge space of possible states and actions to explore. Only recently have researchers begun to develop powerful nonlinear algorithms that may be able to generalize across that space efficiently.[15, 16].

**Guidance by starting with easy tasks.** Although human beings live long enough for the various stages of cognitive development to unfold gradually [17], robots have not yet attained that level of reliability. *Robot shaping* [18] or *learning from easy missions* (LEM) [13] provide typical and intuitive methods for accelerating learning. In LEM the essential problem is how to define easy missions. One solution is that the robot starts close to the goal state in the state space and is gradually moved further from the goal state as learning progresses. A distance measure is defined for the state space, and changes in the Q-values are used to determine when to shift to more difficult situations.

**Environmental complexity control.** Generally, the state space consists of multiple state axes, which leaves the question of how to define closeness to the goal state. This raises a more general issue. How do we define the complexity of the environment in terms of the developmental stage of the robot, and how do we adjust the environment to meet the robot's changing developmental needs? While at first it may seem that we are looking at the problem the wrong way around — adapting the task to the robot rather than the robot to the task — this is in fact what parents do naturally in finding stimulating, age-appropriate ways of interacting with their children. Asada et al.[8] defined the complexity of the environment in terms of the relationship between self-induced motor commands and changes in sensory input.

1. **Self-induced movements in a static environment:** The agent can directly correlate its motor signals with changes in its sensory input (e.g., observing its own hand movements, exploring the environment by alternating foveations with eye saccades).
2. **Passive agents:** Depending on the actions of an agent or other agents, passive agents can be moving or still. A ball is a typical example. As long as passive agents are stationary, they can be treated as part of the static environment. But when they are in motion, there is no simple correlation between an agent's motor signals and sensory projections from the passive agent.
3. **Other active agents:** The relationship between sensory projections from active agents and self-induced movements is not at all a straightforward deterministic mapping. In the early stage, other agents may be treated as noise or disturbance because they lack direct visual correlation with self-induced motor commands. Later, they can be found from more complicated and higher order correlations (coordination, competition, etc.). The complexity is drastically increased.

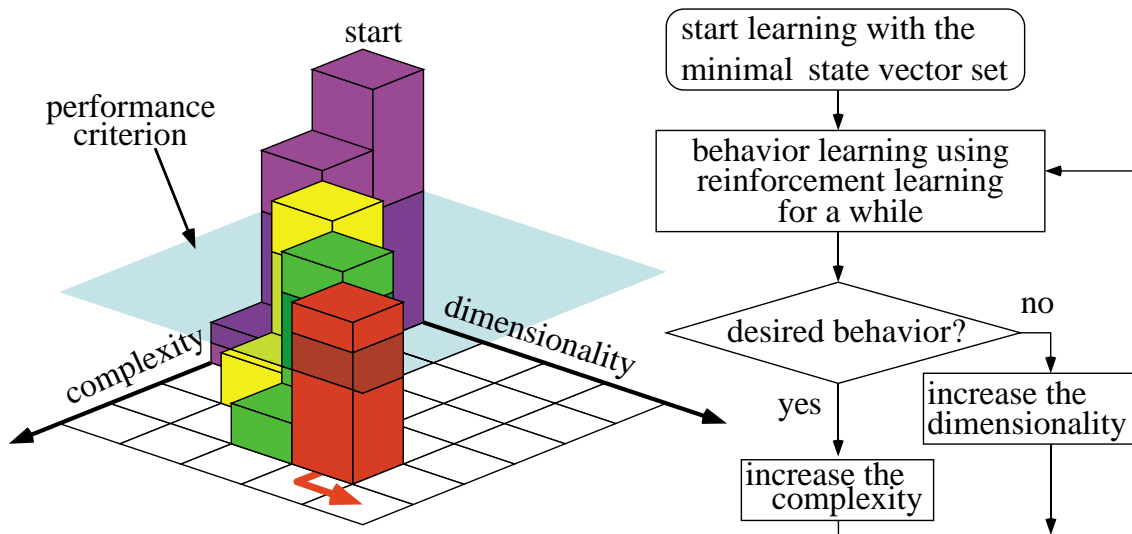


Fig. 2. Interaction between embedded structure and environment

To enable a robot to behave intelligently, the complexity of its internal representation should mirror that of the environment. The problem is how to map the environment’s relevant structure. Uchibe et al. [19] invented an algorithm to estimate which dimensions of the state vector best capture the environment’s complexity, and they applied it to improving shooting behavior in a simplified defender-versus-shooter soccer game [14]. Fig. 2 shows how the dimensionality of the state vector grows with the environment’s complexity — in this example, speed of the defender. As long as the performance (success rate) exceeds the threshold, we increase the complexity by increasing the speed. If the success rate drops below the performance criterion, the robot increases the dimensionality of the state vector.

The environmental complexity control method resulted in performance that slightly exceeded that of learning with a full state vector, and required only one third the time. The fact that both robot and environment are physical systems limits the time a robot has to reach task competency through experimentation. Therefore, *starting small* [11] can be important for bootstrapping learning mechanisms by exposing the learner to the environment’s most prominent features first. To some extent, we see this in child language learning [20].

If a robot develops expectations concerning how self-induced movements transform sensory projections from static objects, passive agents can be detected and modeled from correlations among violated expectations [21]. By diagnosing why expectations were violated, the robot develops higher and higher-order expectations that may scaffold ever more abstract learning.

**Learning schedules in multi-agent learning.** In the last example [14], there was only one learner, and the behavior of the other agent (defender) was controlled according to the skill level of the learner. What kind of control is possible if both agents are learners? In general, the simultaneous learning of action policies from rewards in a multi-agent environment is difficult for a number of reasons: (1) Initially, the agents’ learning focuses more on exploring the state space than exploiting a policy; since exploration depends on trying many unfruitful actions, the learning of both robots may diverge as they try to adapt to each other’s largely random behavior. (2) Even when each robot converges to its locally optimal policy, the robots may easily miss finding policies that — taken together — would be globally optimal.

We introduced a learning schedule [8] to address the first problem: only one agent is allowed to learn. The other agents have fixed policies until the learner’s skill reaches a pre-specified level. Then the next agent begins learning. We applied a cooperative task in the context of RoboCup [12], namely, passing and shooting. The learning schedule successfully resulted in mutual skill development, which is generally difficult to reach in the case of co-evolution [22].

## 5.2 Social interaction and communication

One way that CDR could contribute to our understanding of human beings is by providing models of the cognitive and social processes underlying the development of communication and by testing those models using humanoid robots. The transition from nonverbal to verbal communication is an active area of research that CDR can address. There is a large gap between primates and *Homo sapiens* that needs to be filled in (the missing link, cf. [23]). Since a survey of existing views in linguistics, philosophy, and sociology would be too broad in scope for the purposes of this paper, we focus instead on the issue of symbol emergence and language acquisition from the viewpoint of engineering design.

Over the last few decades many language researchers have reached a consensus that language is an innate ability, and human babies are born with a kind of “language faculty” or device [24]. Broca and Wernicke areas seem to be part of such device, but things are not so simple. All language abilities cannot be reduced to the activities of those areas; in fact many areas are related to language use implicitly. Also, given the biological continuity from primate to *Homo sapiens*, the claim that only human beings have a language device seems difficult to accept. Without question, human brains come into the world especially equipped for language. So the problems facing CDR are:

1. What kind of structure should be embedded inside a robot’s brain? Whether the structure be explicitly or implicitly specified, it should involve a new explanation of how evolution and development bridge the gap between nonverbal and verbal communication.
2. What analytical approaches should we use in CDR (e.g., structural/anatomical, behavioral, evolutionary)? and at what level of detail (e.g., gene and neuron or auditory subsystem)? These are important questions since it is impractical to try to reproduce in robots millions of years of biological evolution in all its detail.

Although current speech recognition and generation technology is useful in some contexts, it is doubtful that most AI systems based on this technology really understand (see Searle’s Chinese Room argument [25]) or can apply language in context as a tool for communication [4]. This is a problem generally for representations that cannot be related to sensorimotor activity [26]. In their book [11], Elman et al. showed that simple grammars could be learned by artificial neural networks. However, the inputs and outputs were just symbols without any semantics. We intend to start exploring how symbols emerge through social interaction. Steels and Voget [27] implemented adaptive language games using robotic agents, and their approach seems closer to ours, but they assumed a protocol for robots to communicate. Since CDR aims to offer at least a partial explanation of the transition from nonverbal to verbal communication, we should focus on how such a protocol could evolve or be learned by humanoid robots.

Schaal [28] surveyed imitative learning methods, and emphasized the importance of imitation as a route to developing humanoid robots. He focuses on efficient motor learning, the connection between action and perception, and modular motor control in the form of movement primitives. He pointed out open problems such as learning perceptual representations and movement primitives, movement recognition through movement generation, and understanding task goals. Schaal then discusses the recent finding that some neurons called “mirror neuron” were active both when the monkey grasps or manipulates objects and when it observes the experimenter making similar actions. Rizzolatti and Arbib [29] speculate that the ability of imitate actions and to understand them could have subserved the development of communication skills based on the fact that a similar system includes the Broca area, which is known to be related to speech generation in the human brain.

From the viewpoint of CDR as a design principle for humanoid robots, such a system probably should be included because capabilities of both motion generation by imitation and motion understanding (e.g., by comparison to one’s own motion repertoire) may be necessary. There seems to be two main kinds of imitation pathways: *visual imitation* (imitative learning by observation) and *auditory imitation* (imitative learning by listening).

The existing methods (e.g., [30, 31, 32]) for the former often assume the global coordinate transformation from a god’s eye viewpoint; however, CDR should focus on how such a transformation emerges through interactions between humanoid robots and/or humans (e.g., in a manner similar to how a baby learns from its parents). Asada et al. [33] proposed an imitation system which recovers the other agent’s view without any knowledge of a global coordinate transformation, but it assumes that the other agent has the same body structure (see Fig 3). They expect to offer a route to motion generation, understanding, and “mind reading” (the theory of mind [34], [35]).




Fig. 3. An imitation system

Imitation must extend beyond mere mimicry to the ability to generate something new. Since the mechanical structure of the human speech generation system is quite complex and a sophisticated integration of voluntary and involuntary muscle controls is necessary to generate sounds [23], auditory imitation has an essential problem: At which level should the robot start to imitate?

The term *social dynamics* can be used to characterize largely unconscious relationships among people that may facilitate our sense of “shared mind.” Without talking to each other, people can share much information since they are embedded in the dynamics, but it is still not well understood how this dynamics developed. Aspects of the dynamics are so-called “shared attention,” “subjective distance,” “group behavior,” and so on. One goal of CDR is to make its structure clear and build robots that function according to it. Concepts like “everyday cognition” or “distributed cognition” may catch its essence [36].

## 6 Conclusion

We have discussed a variety of issues concerning CDR, most of which are far from resolution since CDR has just started. Among them, two issues seem essential for future arguments. The first one concerns the definition of environmental complexity, the adjustment of which can be expected to aid development. However, the definition itself involves a contradiction because before encountering a new environment the robot cannot define the complexity. In the method reviewed [14], the complexity corresponds to the dimension of the estimated state vector, which is obtained in an off-line process. But it seems difficult to estimate the full dimensionality of the state vector accurately before learning begins. One alternative is to develop an online method of state vector estimation.

The second issue is imitation in social interaction between learner and teacher. There are several levels of interactions, each of which has its own issues. If the teacher knows everything about the learner like a god, the teacher can guide the learning process optimally. However, both learner and teacher must realistically have only limited, perspective-dependent knowledge. A further issue concerns the kinds of explicit or implicit means of communication available (e.g., visual, auditory) and the extent to which the teacher knows the learner’s state. Imitative learning seems essential to developing cognitive processes for both motion generation and language acquisition.

The ultimate aim of CDR is both for us to know ourselves by building robots and to build robots capable of functioning in society. Specifically, we want to build robots that can relate to people and each other as individuals by developing specific relationships, just as people do. We believe that in the process of building robotics that can function in society, we cannot help but learn about ourselves too.

## References

1. K. F. MacDorman. Grounding symbols through sensorimotor integration. *Journal of the Robotics Society of Japan*, 17(1):20–24, 1999.
2. C. Breazeal (Ferrell) R. Irie C. Kemp M. Marjanovic B. Scassellat Brooks, R.A. and M. Williamson. Alternate essences of intelligence. In *AAAI-98*, 1999.
3. Mark H. Johnson. *Developmental Cognitive Neuroscience*. Blackwell Publisher Inc., Cambridge, Massachusetts 02142, USA, 1996.
4. S. J. Cowley and K. F. MacDorman. Simulating conversations: The communion game. *AI & Society*, 9(3):116–137, 1995.
5. R. A. Hinde. *Individuals, Relationships and Culture: Links between Ethology and the Social Sciences*. Cambridge University Press, Cambridge, UK, 1987.
6. R. A. Brooks. “A robust layered control system for a mobile robot”. *IEEE J. Robotics and Automation*, RA-2:14–23, 1986.
7. R. A. Brooks. “Elephants don’t play chess”. In P. Maes, editor, *Designing Autonomous Agents*, pages 3–15. MIT/Elsevier, 1991.
8. Minoru Asada, Eiji Uchibe, and Koh Hosoda. Cooperative behavior acquisition for mobile robots in dynamically changing real worlds via vision-based reinforcement learning and development. *Artificial Intelligence*, 110:275–292, 1999.
9. Rolf Pfeifer and Christian Scheier99. *Understanding Intelligence*. The MIT Press, Cambridge, Massachusetts 02142, USA, 1999.
10. T. McGeer. “passive walking with knees”. In *Proc. of 1990 IEEE Int. Conf. on Robotics and Automation*, 1990.
11. J. Elman, E. A. Bates, M. Johnson, A. Karmiloff-Smith, D. Parisi, and K. Plunkett. *Rethinking Innateness: A Connectionist Perspective on Development*. The MIT Press, Cambridge, Massachusetts 02142, USA, 1996.
12. Minoru Asada and Hiroaki Kitano, editors. *RoboCup-98: Robot Soccer World Cup II*. Springer, Lecture Note in Artificial Intelligence 1604, 1999.
13. M. Asada, S. Noda, S. Tawaratumida, and K. Hosoda. Purposive behavior acquisition for a real robot by vision-based reinforcement learning. *Machine Learning*, 23:279–303, 1996.
14. Eiji Uchibe, Minoru Asada, and Koh Hosoda. Environmental complexity control for vision-based learning mobile robot. In *Proc. of IEEE Int. Conf. on Robotics and Automation*, pages 1865–1870, 1998.
15. Leemon C. Baird and Andrew W. Moore. Gradient descent for general reinforcement learning. *Advances in Neural Information Processing Systems*, 1999.
16. Remi Munos and Andrew W. Moore. Barycentric interpolators for continuous space and time reinforcement learning. *Advances in Neural Information Processing Systems*, 1999.
17. Jean Piaget. *Child’s Conception of the World*. Littlefield Adams, 1990.
18. Marco Dorigo. *Robot shaping: An experiment in behavior engineering*. MIT Press, Cambridge, MA, 1997.
19. E. Uchibe, M. Asada, and K. Hosoda. “state space construction for behavior acquisition in multi agent environments with vision and action”. In *Proc. of ICCV 98*, pages 870–875, 1998.
20. E. L. Newport. Maturation constraints on language learning. *Cognitive Science*, pages 11–28, 1990.
21. K. F. MacDorman. Responding to affordances. In *Proc. of 2000 IEEE Int. Conf. on Robotics and Automation*, 2000.
22. D. Floreano, S. Nolfi, and F. Mondada. Competitive co-evolutionary robotics: From theory to practice. In *Proc. of the 5th Int. Conf. on Simulation and Adaptive Behaviors – From animals to animats 5 –*, pages 515–524, 1998.
23. Terrence W. Deacon. *The Symbolic Species: The co-evolution of language and the brain*. W. W. Norton & Company, New York, London, 1998.
24. Norm Chomsky. *Language and Mind*. New York: Harcourt Brace Jovanovich, 1972.
25. J. R. Searle. Minds, brains, and programs. *Behavioral and Brain Sciences*, pages 417–457, 1980.
26. S. Harnad. The symbol grounding problem. *Physica D*, 42:335–346, 1990.
27. L. Steels and P. Voget. Grounding adaptive language games in robotic agents. In *Proc. of the 4th European Conf. on Artificial Life*, pages 474–482, 1997.
28. Stefan Schaal. Is imitation learning the route to humanoid robots? *Trends in Cognitive Science*, 1999.
29. G. Rizzolatti and M. A. Arbib. Language within our grasp. *Trends Neuroscience*, 21:188–194, 1998.
30. K. Ikeuchi and T. Suehiro. Toward an assembly plan from observation. *IEEE Trans. on R&A*, 10:368–385, 1994.

31. Y. Kuniyoshi, M. Inaba, and H. Inoue. Learning by watching. *IEEE Trans. on R&A*, 10:799–822, 1994.
32. M. Kawato F. Gandolfo, H. Gomi, , and Y. Wada. Teaching by showing in kendama based on optimization principle. In *Proc. of the International Conf. on ANN*, pages 601–606, 1994.
33. M. Asada, Y. Yoshikawa, and K. Hosoda. Learning by observation without three-dimensional reconstruction. In *Proc. of the 6th International Conf. on Intelligent Autonomous Systems (IAS-6)*, page (to appear), 2000.
34. D. Premack and G. Woodruff. Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, pages 515–526, 1978.
35. R. Krebs J. and R. Dawkins. Animal signals: Mind reading and manipulation. In In J. R. Krebs and N. B. Davies, editors, *Behavioural Ecology: An Evolutionary Approach (2nd ed.)*. Blackwell, Oxford, 1984.
36. Gavriel Salomon (Ed.). *Distributed Cognitions: Psychological and Educational Considerations*. Cambridge University Press, Cambridge, UK, 1996.