Studying Robot Social Cognition Within A Developmental Psychology Framework

Kerstin Dautenhahn

Department of Cybernetics
University of Reading
United Kingdom
K.Dautenhahn@rdg.ac.uk

Aude Billard

Department of AI
University of Edinburgh
United Kingdom
A.Billard@epfl.ch

Abstract

This paper discusses two prominent theories of cognitive development and relates them to experiments in social robotics. The main difference between these theories lies in the different views on the relationship between a child and its social environment: a) the child as a solitary thinker (Piaget) and b) the child in society (Vygotsky). We discuss the implications this has on the design of socially intelligent agents, focusing on robotic agents. We argue that the framework proposed by Vygotsky provides a promising research direction in autonomous agents. We give examples of implementations in the area of social robotics which support our theoretical considerations. More specifically, we demonstrate how a teacher-learner setup can be used to teach a robot a proto-language. The same control architecture is also used for a humanoid doll robot which can interact with a human by imitation. Another experiment addresses dynamic coupling of movements between a human and a mobile robot. Here, emergent robot-human interaction dynamics are influenced by the temporal coordination between the robot's and the human's movements.

1 Introduction

Grounding artificial life 'as-it-could-be' in knowledge on biological life 'as-we-know-it' is a core issue in the field of artificial life. Researchers are investigating how in hardware or software one can construct artifacts that behave like living organisms, in their own, 'characteristic' ways and by carefully exploiting the 'natural' dynamics of artifact-environment interactions ([Pfe95]). Robots are increasingly used as research tools as part of a 'comparative study'. On the one hand artificial systems serve as models of natural systems in order to investigate open questions in biology ([TJ94]), on the other hand natural systems can serve as models for the construction of artificial systems. Examples are the study of cricket phonotaxis ([Web95]), collective behaviour of social insects [DGF+91]), walking machines imitating stick-insects [CBC+94], fly-like robot vision systems [FPB91]), and many more.

Recently, in studies on robotic social behaviour (in interactions with other robots and/or humans) theories from developmental psychology ([HJ96]) came into focus. This paper discusses that a particular theory developed by Vygotsky which emphasises the importance of teaching and social interactions in the development of children's cognitive skills, can offer a framework to investigate the development of 'robots in society'. This framework is based on the assumption that in the same way as intelligence and embodiment make only sense with respect to an agent's environment, social behaviour can only be evaluated with respect to other agents which are part of the system, including the human experimenter.

To give an example of a socially coupled system: If a human is playing with a dog then although the dog can be well trained and the human can have the strong intention to pursue a particular goal, the interaction dynamics between the two agents which are both influencing each other's behaviour cannot be completely predicted, the dynamics emerge from interactions between the two socially situated agents.

The rest of the paper is divided as follows: Section 2 provides the theoretical background from developmental psychology. The next two sections (3, 4) summarise some of our work on social robotics, including robot-robot and robot-human interaction. Section 5 discusses how the robotics experiments support our theoretical framework. The paper concludes with section 6.

2 Theories in Developmental Psychology

In the following we contrast two influential 20th-century theories on children's cognitive development ([Le95]).

- Piaget and his view of *the child as a solitary thinker*: The cognitive development from child to adult is based on qualitatively different universal stages. The social context might assist development, but the child's own activity plays the essential role in the progression of cognitive stages. At the centre of Piaget's theory is the isolated child as a 'little scientist', namely exploring and testing the world on its own. The child goes through different stages or periods according to an invariant sequence which holds across cultures: sensorimotor stage (up to 2 years), preoperational stage (2-7 years), concrete operational stage (7-12 years), formal operational stage (12 years and older). Transitions involve the processes of assimilation, accomodation and equilibration. Piaget believed that cognitive development is a spontaneous process and does not depend on teaching by adults.
- Vygotsky and his view of *the child in society*: Interactions with adults and peers and teaching are essential for cognitive development, the social and cultural context matters. Human cognitive capacities change as a result of historical development and new cultural tools (technological and psychological tools). His work addresses how the child acquires cultural tools through development and in interaction with other persons. He argues that concepts, language, voluntary attention and memory originate in culture, i.e. are interpersonal processes before they become internalised by the child as intrapersonal processes. Vygotsky stresses the role of teaching, in particular for guiding the development of abstract modes of thought. Intrinsic (within individual) and external (cultural) play an equal part in Vygotsky's theory, which like Piaget's assumes development in stages. According to Vygotsky social interactions determine structure and pattern of internal cognition. He claims that the mechanism underlying higher mental functions are grounded in social interaction, and that all higher mental functions are internalised social relationships.

Piaget's theory of cognitive development has been challenged from various directions, e.g. questioning the universal nature of the sequences of the stages, or arguing for a more continuous notion of development using a dynamical systems account (e.g. [ST93]). Recent experiments have confirmed Vygotsky's belief in the essential role of social interaction and teaching as a scaffolding mechanism which is important for the child in order to reach higher levels of competence and control based on current skills. Hereby concepts are not taught directly but through social interaction, the child's experiences are re-arranged, a shared understanding develops between the child and its interaction partner.

Especially interesting to us are theories regarding the child's language development. Language is a complex cognitive skill, and for many application areas a language-like system can be important for a robot to possess. A Piagetean viewpoint sees language as a product of the cognitive development of mental representations, while Vygotsky believes that the sole primary function of language is communication with peers and adults, and that language develops exactly in this context. In section 3 we summarise briefly the experiments we carried out in teaching a robot a language-like system. We followed Vygotsky's approach in these experiments by providing the robot with the primary ability to socially interact with a teacher, by imitating the teacher's movements. Imitation is used as a means to synchronise the movements of teacher and learner robot which implicitly directs the learner's attention towards the teacher's movements.

As discussed in [HJ97], [Dau97] the synchronisation and coordination of movements between humans and their environment seems to play a crucial role in the development of children's social skills. In particular, getting the interaction dynamics right between infant and caretaker seems to be a central step in the development of social skills, since they require coordinating one's external and/or internal states with another agent, to become engaged in the situation. The states need not be exactly the same, as dancers in a group can dance different movement patterns, but their states are temporally coordinated. Moreover, dancing in a group is more than the sum of its parts, a dance is an emergent pattern in which different individual dancers take part and synchronise their movement towards each other and the group as a whole. In section 4 we summarise a simple robotic experiment which studies how socially interesting phenomena can arise from simple interaction dynamics. Here, temporal synchronisation of two agents, a robot and a human, can be exploited for teaching purposes, so that the robot's actions are influenced by way the human is responding to its actions.

In [HJ96] Horst Hendriks-Jansen recognised the link between recent context-oriented findings in developmental psychology on the one hand and behaviour-oriented autonomous agent research on the other hand. Hendriks-Jansen argues that we can learn from robotic experiments about the importance of agent-environment interactions, thus robots might be appropriate tools for a synthetic psychology ([Bra84]). Following this line of thought, we view the relationship between developmental psychology and research into autonomous robots as follows: 1) Robots can serve as tools to test theories and mechanisms suggested by developmental psychology, 2) Robots can be built on the basis of findings in developmental psychology in order to develop cognitively plausible designs and control architectures. The authors' work on socially intelligent robots is along the second line of investigation, i.e. we are building cognitively plausible robots which are however not exactly copying natural systems, but which are based on findings in psychology. Generally, our work is related to research studying robot-human interaction (e.g. [KDK97], [Kun97]). The most closely related work in the field of social robotics is currently done at the MIT AI-Lab ([Bre98], [Sca99]) where humanoid robotic platforms are studied which can socially interact with humans, e.g. by imitating humans and responding to a human's behaviour with changes in facial expressions. This work is different from our work which studies teaching a robot a 'language', and processes of movement shaping between humans and machines. However, the general approach of putting social skills and the dynamics of interaction at the centre of robotic intelligence, and to pursue a direction grounded in developmental psychology theory is a common ground.

The next section summarises robotic experiments based on our previous theoretical arguments.

3 Robotic Experiments on Teaching a Proto-Language

We carried out a number of experiments [Bil98], in which an autonomous mobile robot was taught a synthetic language-like system (proto-language). The proto-language consists of a lexicon and of combinations of words of the lexicon, which form English proto-sentences. The robots' learning and behavioural capabilities are provided by a connectionist model. It is a Dynamic Recurrent Associative Memory Architecture (DRAMA) [BH99], which allows learning of spatio-temporal regularities and time series. Learning of the proto-language, i.e. signal-perception association and combination, results from multiple spatio-temporal association across the robot's sensor-actuator state space. We sketch briefly some of these experiments in the following.

Teaching a proto-language to a robot. In a first set of experiments, we studied transmission of a vocabulary from a teacher robot to a learner robot. While the learner robot follows, and thus, implicitly imitates/replicates the movements of the teacher robot, it is taught a vocabulary to label its perception of inclination [BH98] (see figures 1, 2), orientation, its actions and observations of objects [BD97]. The learner robot associates the teacher's radio signals (which the teacher emits to describe its own perceptions) with its (the learner's) own perception, which, because of its spatial closeness to the teacher agent, is very similar to that of the teacher. Further, we carried out simulation studies, in which we investigated transmission and use of the vocabulary among a group of nine robotic agents [BD99]. We showed that the teacher-learner scenario used in the first set of experiments to transmit a vocabulary from one teacher robot to one learner robot, scales up successfully to transmit a vocabulary among a group of robots.

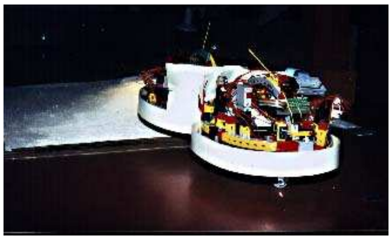

Figure 1: Teacher and learner robot in the physical set-up.

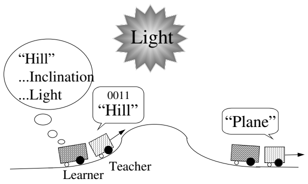

Figure 2: The learner robot follows the teacher robot which climbs a hill. It learns the words 'hill' and 'light' emitted by the teacher which describe its current perceptions.

Robota, an Interactive Robotic Doll. Experiments were done using Robota, a doll-like robot, which was able to mirror the arm and head movements of a human instructor (see figure 3) [BDH98]. Through

Figure 3: The doll robot mirrors the head and arm movements of the demonstrator. It does so by detecting variations in the infrared emissions of the sensors on the demontrator's glasses and in her hands.

this imitative scenario, Robota was taught to perform and label dance patterns (the robot is controlled by the DRAMA architecture). The robot was taught a synthetic proto-language, composed of a lexicon and combinations of words of the lexicon which formed proto-sentences. These described the robot's interaction with the teacher, e.g. "I move left arm", "You touch right foot", "Feed me". These experiments demonstrated that the robot directed by the DRAMA architecture could learn a basic 'language' or 'proto-language', which shares some properties with natural language: 1) each word carried a specific meaning (e.g. *arm, head, right, left*, etc.); 2) words can be combined and the combination can be given a different meaning, the meaning of each combination being determined by the order of appearance of each word in the combination (e.g. *You touch right arm, I move head right*; 3) the conceptual meaning of each word is learned implicitly by only presenting each word as part of complete sentences, which can then be used in new word combinations; 4) precedence between words' appearance in the combination is learned and can be used to infer the correct order when constructing a new word combination.

Simple sentence examples were used in these experiments, in which the words could easily be tied to the taught concepts. Moreover, the 'language' the robot was to learn was regular; that is, the robot's learning task was to recognise temporal regularities in the words' ordering across the taught sentences and to correlate the words' usage with its sensors and actuators' activity (the physical context of the teaching gives meaning to the demonstrator's utterance). As such, these experiments were a first step towards demonstrating the validity of the system (the learning architecture and the imitative strategy) for teaching a robot a symbolic communication system, in this case a regular language. However, it remains to be shown how the system could scale up to learning a complete language with grammar structure and irregularities. The fact that the DRAMA architecture is comparable in function to a Hidden Markov model or other recurrent neural networks, which are models currently used in techniques of Natural Language Processing, suggests that the model could scale up successfully to learning any regular language. It is however unclear whether the model could allow learning of irregular languages, such as human languages, which requires complex syntax, the understanding of ambiguity and the integration of exceptions to grammatical rules.

Synthesis. Little work has yet been done in teaching a physical robot a synthetic form of communication (e.g. [YS93], [SV97]). Our work differs from those studies in several aspects: 1) The learning and behavioural capacities of the robots result from a single cognitive architecture; it is a connectionist model which has general ability for extracting spatio-temporal regularities in a dynamic environment [BH99]. 2) The proto-language which the robot is taught is not restricted only to a lexicon, where each word of the lexicon relates to a single specific perception, as in the two compared studies. 3) We consider both the cognitive and behavioural skills behind a robot's learning of a proto-language. In particular, we point out the importance of a movement coordination between the communicative agents, for the agents to share a common perceptual context, onto which they can ground a common understanding of the proto-language. Movement imitation is in our experiments the means by which we create a coupling between the two agents. In this aspect, our work brings a novel contribution to current research in those areas, addressing the *symbol grounding problem* [Har90] from a behavioural by opposition to a pure cognitivist point of view, for which the notions of embodiment and situatedness are key issues for symbolic cognition.

4 Engagement in Human-Robot Interaction

In [Dau99] we describe experiments which address a dynamics approach towards robot-human interaction, based on ideas on movement synchronisation in social interaction which we discussed previously ([Dau97]). The experiments are based on the concept of *temporal coordination* as a 'social feedback' signal for reinforcement learning in robot-human interaction. This section briefly summarises the experiments.

The experiments consists of one mobile robot (a Lego robot, in another trial a *fischertechnik* robot) and a human with a stationary video camera pointing at her. The camera image is used to classify hand movements of a human in front of the camera into six categories: a) moving horizontally from right to left or left to right, b) moving vertically up or down the screen, c) moving the hand in circles either clockwise or anti-clockwise. Information about the classification of the movements is sent to the robot via radio-link. The robot is controlled using a PDL ([Ste94]) architecture which makes the system relatively robust to noisy input, namely incorrect 'classification' of the input signals.

The control program which runs on the mobile robot can run in two modes: in the *autonomous mode* it performs a sequence of movements autonomously, and depending on the feedback from the human certain movements can be selected. Functionally, this controller behaves like an action-selection and reinforcement mechanism. In the *slave mode* the six possible inputs (movements by the human) are associated to four possible outputs (movements of the robot): turning left, turning right, moving forward, moving backwards. Functionally, this controller behaves like a Hebbian learning controller. Figure 4 shows the basic setup of the experiments and the association matrix.

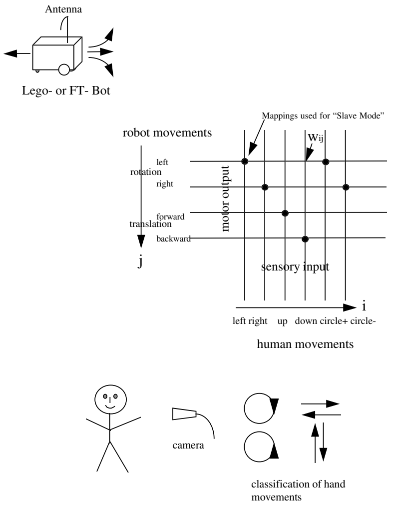

Figure 4: The basic experimental set-up and the association matrix.

Generally, changes in the association matrix were due to 'temporal coordination', i.e. the temporal relationship between the movements of human and robot. Temporal coordination means that mappings between movements of the robot and the human can change dynamically in a coupled way, the mappings need not be static.

Synthesis. In the experiment outlined above we studied the temporal coordination between a human and a mobile robot which changed, depending on the reactions or the feedback by a human, its movement repertoire. A very simple association matrix was used for training purposes, however, it turned out during demonstrations of this system that it was rather the human than the robot which was the learner is these experiments. In the slave mode humans very quickly realised that the robot's movement were correlated to their own movements and that the robot could be operated like a passive puppet-on-a-string toy. However, the 'puppet' was sensitive to how long humans interacted with it and how 'attentive' they were (e.g. adapting the speed of their own movements to the robot's speed). A cooperative human paid attention to the robot's movement and kept it moving, 'neglect' made the robot slow down and finally stop. The robot could also be operated so that it finally only performed those movement(s) to which the human gave longest response and attention. The robot therefore adapted to the human and 'personalised', i.e. after a while only reacting to the human's 'favourite' movement. This also occurred in the autonomous mode, however then the human could only select from a given repertoire of movements, i.e. the human could shape the robot's autonomous behaviour. A cooperative human learnt quickly to give the appropriate feedback in order to keep the robot moving. Depending on the human's preference the robot then (in the autonomous mode) ended up performing only one or a few different movements. Thus, the behaviour of the robot finally was typical of the human who interacted with it. Despite the fact that we used a deterministic robot controller, robot-human interaction patterns were emergent since they could not be predicted from the robot controller alone: the robot's movements changed according to the way humans reacted to it, and vice versa, and these coupled dynamics could not be predicted and could only be studied in real world experiments. Using the notion of emergence as defined in Artificial Life ([Lan89]), we can say that simple local rules (the 'genotype' of the system: simple control rules for the robot) executed in an environment could give rise to emergent patterns (the 'phenotype') of robot-human interactions.

Potentially this method can be used to adapt the behaviour of robot to a human's individual needs and preferences, in particular if the 'movements' which we used become complex behaviours and can be shaped individually. This is done in a non-symbolic way, without any reasoning involved except for defining an association matrix and detecting temporal coordination. More sophisticated learning architectures could be based on such a system, e.g. for the study of imitation (see section 3). This is particularly attractive if the robot has more degrees of freedom than the simple system we used in this robot-human interaction experiments. Future work in this area has to show whether and how the approach can scale up to more complex robotic platforms, and to areas where humans have long periods of interaction with a robot, e.g. in service robotics (e.g. [WAP+98]). The concepts might also be tested with non-robotic socially intelligent agents ([Dau98]).

We believe that synchronisation of movements can contribute to life-like behaviour as well as appearance can. However, in robot-human interaction so far the analysis of the human's behaviour resulting in a symbolic description which can then be used to control a robot's behaviour has been a predominant approach. Often, body movements are used by computationally expensive vision routines which extract information on position or gestures, rather then using the dynamic nature of the movements itself. Temporal coordination might be a means to link the human's and the robot's dynamics in a way which appears 'natural' to humans, an approach which is also applicable to interactive agent systems like ALIVE (Artificial Life Interactive Video Environment, [MBD+95]). The 'dancing' experiments described in this section were strongly inspired by Simon Penny's PETIT MAL, an interesting example of a non-humanoid but socially successful mobile robot [Pen97]. A double pendulum structure gives the robot an 'interesting' (very smooth behaviour transitions) and at the same time unpredictable movement repertoire, pyro-electric and ultra-sonic sensors enable the robot to react to humans like approaching or avoiding. The system has been running at numerous exhibitions and attracted much attention despite of its technological simplicity. The robot is a purely reactive system without any learning or memory functionality, the complexity lies in the balanced design of this system, and not in its hardware and software components. Robot-human interactions with PETIT MAL generate interesting dynamics which cannot be explained or predicted from the behaviour of the human or the robot alone. This implementation at the intersection of interactive art and robotics demonstrates the power of dynamics in human-robot social interactions. Combining learning and movement shaping techniques which the authors investigate with interesting designs like PETIT MAL outline a direction for building minimal designs for socially competent robots.

5 Discussion: Socially Intelligent Robots

Our research is based on the assumption that in order to study the cognitive development of robots we have to consider the 'robot in society', i.e. using Vygotsky's approach to see social interactions as fundamental, and as a context which can scaffold the development of cognitively richer functionalities. Our work on socially intelligent robots (see sections 3 and 4) shows the following correspondences to Vygotsky's approach to children's cognitive development:

1. Communication for the sake of communication. We study robotic language for the primary purpose of communication (see section 3). In a later step we investigate the way how the agents might use their communication skills, but the experiments are not based on the learner's fitness advantage. The experiments do therefore complement studies on the evolution or learning of communication where fitness or task achievement usually play a central role ([MdRM96], [Pao97], [Nob98]).

2. Learning in a social context. We investigate learning by teaching. In our experiments the learner using an imitative following strategy is literally guided by the teacher through the environment. This is more related to Vygotsky's 'arranging experiences ' ideas rather than didactic teaching. The teacher robot in our experiment does not specifiy what should be learnt, the symbols which are communicated are from the learner's point of view treated like any other perceptual input.

3. Internalisation. The learner learns on the basis of its own sensorimotor experiences; the learner's own activity is (as both Piaget and Vygotsky argue) at the centre of the learning process. As symbols are learnt that describe certain situations (characterised by sensorimotor experience), they form a basis to be used in sequences, e.g. to produce a sequence of actions, or a sentence. In section 4 we studied the interaction dynamics which form a basis for learning and the engagement of two agents in a social dialogue.

4. Shared understanding. The learner robot's aquisition of a proto-language is the result of repeated social interaction, until the learner agent has sucessfully 'understood' the word, which means that it was able to use its associative memory in order to reproduce the word in a similar context. A learner can become a teacher so that a shared understanding can develop in a group of agents.

6 Conclusion

This paper gave a brief introduction to Vygotsky's theory of cognitive development, contrasted to Piaget's approach which is still very influential in artificial intelligence and autonomous agents research. Experiences and learning of the isolated agent are widely seen as the first step towards agent development, social matters are generally considered as a subsequent step, building on and adding to the first stage. We presented in this paper experiments on socially intelligent robots based on the strict assumption that sociality lies at the heart of cognitive development. A general control architecture was described which allows control of agents which can learn a synthetic proto-language and action sequences taught by a teacher. Moreover, learners can become teachers to others, potentially resulting in culturally shared knowledge and experience. We gave an example of robot-human interaction where exploiting temporal coordination in the interaction dynamics let to interesting emergent phenomena based on the movement dynamics of two agents (robot and human) which are socially coupled in an interaction situation.

Our experiments demonstrated that focusing on the social context and interaction dynamics leads to interesting experiments which can contribute to socially intelligent robotic agents. As we showed, this approach challenges the view that 'intelligence' has to be solely part of the robot control system. Instead, global interaction patterns emerge from the interaction dynamics of the agents engaged in an embodied and socially situated 'dialogue'.

References

[BD97] A. Billard and K. Dautenhahn. Grounding communication in situated, social robots. In *Proc. TIMR, Manchester, Towards Intelligent Mobile Robots TIMR UK 97, Technical Report Series of the Department of Computer Science, Manchester University*, 1997.

[BD99] A. Billard and K. Dautenhahn. Experiments in social robotics: grounding and use of communication in autonomous agents. *Adaptive Behaviour Journal, Special Issue on Simulation Models of Social Agents*, 7(3-4), 1999.

[BDH98] A. Billard, K. Dautenhahn, and G. Hayes. Experiments on human-robot communication with robota, an imitative learning and communication doll robot. Technical Report CPM-98-38, Centre for Policy Modelling, Manchester Metropolitan University, UK, 1998.

[BH98] A. Billard and G. Hayes. Transmitting communication skills through imitation in autonomous robots. In A. Birk and J. Demiris, editors, *Learning Robots: A Multi-Perspective Exploration*. LNAI Series, Springer-Verlag, 1998.

[BH99] Aude Billard and Gillian Hayes. Drama, a dynamical recurrent associative memory architecture for on-line learning of temporal correlated events in autonomous mobile agents. Adaptive Behavior journal, 7(1), January 1999, 1999.

[Bil98] Aude Billard. Drama, a connectionist model for robot learning: Experiments on grounding communication through imitation in autonomous robots. PhD Thesis, Depart. of Artificial Intelligence, University of Edinburgh, 1998.

[Bra84] Valentin Braitenberg. *Vehicles: Experiments in Synthetic Psychology*. MIT Press, Cambridge, 1984.

[Bre98] C. Breazeal. A motivational system for regulating human-robot interaction. in *Proceedings of AAAI98*, Madison, WI, 1998.

[CBC+94] H. Cruse, Ch. Bartling, G. Cymbalyuk, J. Dean, and M. Dreifert. A neural net controller for a six-legged walking system. In P. Gaussier and J.-D. Nicoud, editors, *Proc. From Perception to Action Conference, Lausanne, Switzerland*, pages 55–65. IEEE Computer Society Press, 1994.

[Dau97] Kerstin Dautenhahn. I could be you – the phenomenological dimension of social understanding. *Cybernetics and Systems*, 25(8):417–453, 1997.

[Dau98] Kerstin Dautenhahn. The art of designing socially intelligent agents: science, fiction and the human in the loop. *Applied Artificial Intelligence Journal, Special Issue on Socially Intelligent Agents*, 12(7-8):573–617, 1998.

[Dau99] Kerstin Dautenhahn. Embodiment and interaction in socially intelligent life-like agents. In C. L. Nehaniv, editor, *Computation for Metaphors, Analogy and Agents*, pages 102–142. Springer Lecture Notes in Artificial Intelligence, Volume 1562, 1999.

[DGF+91] J. L. Deneubourg, S. Goss, N. Franks, A. Sendova-Franks, C. Detrain, and

L. Chrétien. The dynamics of collective sorting: robot-like ants and ant-like robots. In J. A. Meyer and S. W. Wilson, editors, *From Animals to Animats, Proc. of the First International Conference on simulation of adaptive behavior*, pages 356–363, 1991.

[FPB91] N. Franceschini, J. M. Pichon, and C. Blanes. Real time visuomotor control: from flies to robots. In *Proc. of IEEE Fifth International Conference on Advanced Robotics*, 1991.

[Har90] S. Harnad. The symbol grounding problem. *Physica D 42*, pages 335–346, 1990.

[HJ96] Horst Hendriks-Jansen. *Catching Ourselves in the Act: Situated Activity, Interactive Emergence, Evolution, and Human Thought.* MIT Press, Cambridge, Mass., 1996.

[HJ97] Horst Hendriks-Jansen. The epistomology of autism: making a case for an embodied, dynamic, and historical explanation. *Cybernetics and Systems*, 25(8):359–415, 1997.

[KDK97] Volker Klingspor, John Demiris, and Michael Kaiser. Human-robot-communication and machine learning. *Applied Artificial Intelligence Journal*, 11:719–746, 1997.

[Kun97] Yasuo Kuniyoshi. Fusing autonomy and sociability in robots. In W. Lewis Johnson, editor, *Proc. of the First International Conference on Autonomous Agents*, pages 470–471, 1997.

[Lan89] Christopher G. Langton. Artificial life. In C. G. Langton, editor, *Proc. of an Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems, Los Alamos, New Mexico, September 1987*, pages 1–47, 1989.

[Le95] Victor Lee and Prajna Das Gupta (eds.). *Children's Cognitive and Language Development.* Blackwell, 1995.

[MBD+95] P. Maes, B. Blumberg, T. Darrell, A. Pentland, and A. Wexelblat. Modeling interactive agents in alive. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI95-2). Montreal, Quebec, Canada, 1995*, pages 2073–2074, 1995.

[MdRM96] Antonio Murciano and Jose del R. Millan. Learning signaling behaviors and specialization in cooperative agents. *Adaptive Behavior*, 5(1):5–28, 1996.

[Nob98] Jason Noble. Tough guys don't dance: Intention movements and the evolution of signalling in animal contests. In R Pfeifer, B. Blumberg, J.-A. Meyer, and S. W. Wilson, editors, *Proceedings of the Fifth International Conference on Simulation of Adaptive Behavior: From Animals to Animats 5*, pages 471–477. The MIT Press, 1998.

[Pao97] Ezequiel Di Paolo. An investigation into the evolution of communication. *Adaptive Behavior*, 6(2):285–324, 1997.

[Pen97] Simon Penny. Embodied cultural agents: at the intersection of robotics, cognitive science and interactive art. In *Socially Intelligent Agents*, pages 103–105. AAAI Press, Technical report FS-97-02, 1997.

[Pfe95] Rolf Pfeifer. Cognition – perspectives from autonomous agents. *Robotics and Autonomous Systems*, 15:25–46, 1995.

[Sca99] Brian Scassellati. Imitation and mechanisms of joint attention: A developmental structure for building social skills. In C. L. Nehaniv, editor, *Computation for Metaphors, Analogy and Agents*, pages 176–195. Springer Lecture Notes in Artificial Intelligence, Volume 1562, 1999.

[ST93] L. B. Smith and E. Thelen, editors. *A Dynamic Systems Approach to Development: Applications.* A Bradford Book, The MIT Press, 1993.

[Ste94] Luc Steels. Mathematical analysis of behavior systems. In P. Gaussier and J.-D. Nicoud, editors, *Proc. From Perception to Action Conference, Lausanne, Switzerland*, pages 88–95. IEEE Computer Society Press, 1994.

[SV97] L. Steels and P. Vogt. Grounding adaptive language games in robotic agents. In P. Husbands and I. Harvey, editors, *Proceedings of the Fourth European Conference on Artificial Life, ECAL97*, pages 473–484. MIT Press, 1997.

[TJ94] C. Taylor and D. Jefferson. Artificial life as a tool for biological inquiry. *Artificial Life*, 1:1–13, 1994.

[WAP+98] D. M. Wilkes, A. Alford, R. T. Pack, T. Rogers, R. A. Peters II, and K. Kawamura. Toward socially intelligent service robots. *Applied Artificial Intelligence*, 12(7-8):729–766, 1998.

[Web95] Barbara Webb. Using robots to model animals: a cricket test. *Robotics and Autonomous Systems*, 16:117–134, 1995.

[YS93] H. Yanco and L. A. Stein. An adaptive communication protocol for cooperating mobile robots. In H.L. Roitblat J.-A. Meyer and S.W. Wilson, editors, *From Animals to Animats 2: Proceedings of the Second International Conference on the Simulation of Adaptive Behavior*, pages 478–485. The MIT Press, Bradford Books, 1993.