

Research article

DRAMA, a connectionist architecture for online learning and control of autonomous robots: experiments on learning of a synthetic proto-language with a doll robot

Aude Billard

The author

Aude Billard is at LAMI, EPFL, IN-F Ecublens, Lausanne, Switzerland.

Keywords

Language, Neural networks, Robots

Abstract

Currently, academic as well as industrial research investigates the design of social skills for a robot, which could facilitate the robot's interaction with humans. Work in this area pointed out three important abilities required by the robot, namely the capacity for communicating, for learning and for imitating. The work presented in this paper integrates these three aspects by building robots capable of learning a language through the observation and imitation of a teacher agent. This paper presents DRAMA, an artificial neural network architecture which allows online learning of spatio-temporal regularities and time series in a robot's sequence of perceptions. The model has been implemented in a number of experiments, using wheeled robots, in which the robot learned the topography of its environment relative to fixed landmarks. This paper presents the latest implementation of the model in Robota, a doll robot, which is taught a synthetic proto-language.

Electronic access

The research register for this journal is available at

<http://www2.mcb.co.uk/mcbrr/ir.asp>

The current issue and full text archive of this journal is available at

<http://www.emerald-library.com>

Introduction

There is an high potential in terms of both industrial applications and academic research behind research on robot teaching by demonstration. The last 20 years have seen a breakthrough in the development of optimisation algorithms used to provide robots with learning capabilities. Learning methods have been applied so far essentially to solving problems related to dealing with a static environment, e.g. learning the topography of an office environment, recognising faces from sets of static images. Recently, a new direction on robot learning has emerged where people investigate robot learning from interacting with animated agents, either humans or other robots. If one could teach the robot to perform complex skills by simply demonstrating them rather than programming them explicitly, one would reduce importantly the programming cost involved in the robot's construction. The demonstrator might not even need to be the programmer. Thus, in applications where flexibility, adaptability, and social skills of robots are required (e.g. in the area of service robotics) machines could be delivered with little predefined domain knowledge, but a rich learning potential, and knowledge about tasks and actions could be trained during interactions with human operators and users.

In this line of research, recent work investigates the design of imitative skills for the robot, to be used as part of experiments on robot teaching by demonstration. There, the robot's imitative behaviour is used to teach the robot complex motor skills, by the robot first observing and then repeating the movements of the demonstrator. In Kuniyoshi and Inove (1994); Mataric (1994) and Schaal (1997), the robot replicates three-dimensional movements of a human demonstrator's arm movements, while in Demiris *et al.* (1997) and Scassellati (1998), the robot replicates the movements of the demonstrator's head. This

The author thanks Kerstin Dautenhahn for comments on this paper and to her and Gillian Hayes for their guidance with the author's PhD research; also Clive Loughlin for comments on an earlier version of this paper; the technicians of the Department of Artificial Intelligence at the University of Edinburgh, for their help in constructing the electronic parts of the robots; and Auke Jan Ijspeert for his help in constructing the doll robot's body. This research was supported by a grant from the Swiss National Science Foundation.

paper presents experiments, in which the robot's imitative skills are used for teaching the robot a language. When replicating the teacher's actions, the robot makes similar perceptions to that of the teacher. The robot learns the language by attaching the teacher's words to the observations it makes, which, because of the imitation game, are similar to that of the teacher.

Communication skills are fundamental to human society. They play an essential role for the transmission of knowledge across generations and are at the basis of social relationships between individuals. Humans do not communicate only with each other. They also address verbally or through gesture their pets and other familiar animals, which they expect to show some degree of understanding of their talking, such as obedience to orders, responding to their name. Similarly, a minimum of communicative capabilities could be expected from artificial agents, such as robots, which are expected to share several aspects of human daily life (e.g. vacuum cleaner, food/money distributors, autonomous toys). For instance, human-robot interaction might be facilitated if the robot is able to learn and remember its different users' preferences. Such ideas are currently implemented in the Internet domain, by defining user preference sensitive search engines (e.g. "softbots" or "situated" computer interfaces (Lueg and Pfeifer, 1997)), and discussed in research on robot entertainment (Blumberg *et al.*, 1996; Maes, 1995).

Communicating with a robot has at least two advantages. The robot could explain its behaviour, which would make the task of its programmer easier in case of failure and would make the robot more trustable for its user. The robot could be commanded verbally, which makes it a more natural way (requiring no training) of interacting with the machine than by using a remote control device. It could also be taught by its user or programmer, e.g. teaching a new vocabulary to be attached to new observations. While the robot could be provided with a built-in communication protocol, it would, however, show more adaptability towards the task and environment in which it is to be used, if it was able to learn the language. If the robot could learn the meaning of each word of the language and the language's specific syntax, the robot would not be restricted to be used only in specific domains or a specific country. This paper

reports on the development of a control architecture, which allows an autonomous robot to learn a language, which it is taught either by a human instructor or by another robot.

There are mainly three types of algorithms used to enable learning in autonomous mobile agents: reinforcement learning (RL), evolutionary algorithms (EAs) and artificial neural networks (ANNs) architectures. Because we aimed for a general robot control mechanism, as independent as possible from the particular environment and hardware used for the implementation, ANNs were more interesting to us than RL and EA techniques. In contrast to these other techniques, ANNs require little knowledge of the task by not relying on the design of a good evaluation function for the robot's performance. However, a current disadvantage of ANNs over RL is that they are often computationally too heavy to enable online and on-board computation. In addition, they often require numerous examples, i.e. a long training phase, before performing well. Thus, training the robot and then testing its performances are often two separate phases. Online computation means that learning of the robot is done simultaneously to the robot's control and is used immediately by the robot's controller. Online learning is a fundamental requirement to achieve robust and adaptable robots. In our case, this was a requirement quite difficult to satisfy as we use robots with poor computational capacities. Our robots are provided with a micro-controller with 512k byte EPROM space and 128k byte Static RAM. The central processing unit (CPU) is a Phillips 93C100 series 68000 compatible running at 30MHz. It has no facilities for floating points and all calculations have to be done with integers.

Driven by these considerations, we develop a control architecture which allows on-line learning and control of an autonomous robot. The model is made of a connectionist architecture, called DRAMA, providing online learning of spatio-temporal regularities and of time series in the robot's perceptions. We implemented the model in a number of experiments with wheeled robots, where the robot learned spatial regularities of its environment by recognising landmarks and temporal regularities by recording the time delay for travelling from one landmark to the other. Further, experiments were done in

which a learner robot was taught by a human instructor and by another robot a vocabulary to label the landmarks. The vocabulary consisted of radio bit-strings, emitted by the teacher robot, which the robot learned to associate (by the same associative process as for learning the landmarks) to its observation of the landmarks.

This paper presents a new implementation of DRAMA in an autonomous doll-shaped robot, which is taught a synthetic proto-language. The robot is provided with the ability of imitating/mirroring the arm and head movements of a human demonstrator. The robot's imitative behaviour is used by the instructor to lead the robot to perform sequences of actions, during which she teaches the robot different words to label the different sequences. The robot is also provided with a number of body sensors to detect touch on different parts of its body and to measure its body inclination. The teacher exploits the robot's ability at learning time series (property of the DRAMA architecture), by teaching the robot combinations of words, which form English proto-sentences, to describe its actions and its different perceptions.

The robot's controller

The robot's controller consists of two parts (see Figure 1), a set of event detector modules and a central associative module, namely the DRAMA architecture. Information from the sensors and actuators is encoded in binary (0/1) bit-strings and processed in two stages. It is first segmented in the event detector module before being forwarded to the associative module, so that only information which results from a change between the previous and current state of the sensor or actuator is transferred. Figure 1 shows a schematic representation of the robot's controller. All the

robot's sensors, i.e. tilt sensor (inclination), keyboard and body switches, and actuators (arms and head motors) are inter-connected through the associative module.

The DRAMA network

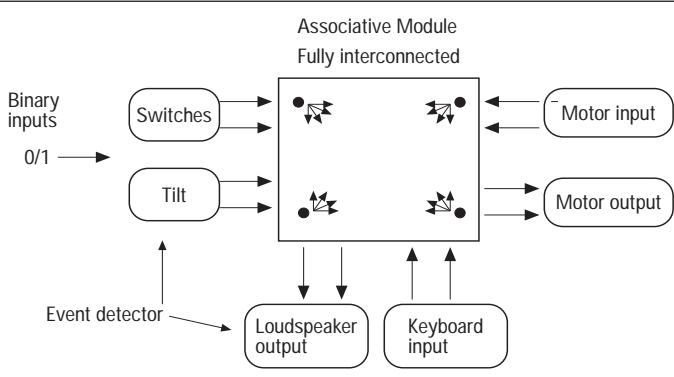
The associative module of the robot's controller is made of an artificial neural network, where each unit is connected to one unit of the event detector modules and represents one bit of sensor or actuator information. It is a dynamical recurrent associative memory (DRAMA), whose mathematical description is given in (Billard and Hayes, 1999). We mention here only the characteristics which are of interest for the reader's understanding of the experiments. DRAMA is a fully connected network with self-connections on each unit. Each connection is associated with two parameters, a "time parameter", which records the time delay between activation of each unit linked by the connection, and a "confidence factor" which records the frequency of activation of the connection. Connection parameters are updated following Hebbian rules (see equations (1) and (2)). That is, whenever two new inputs are presented to the associative module, which causes the activation of two or more units in the network, they are associated (i.e. their connections parameters are updated). Because of the recurrent structure of the network, sensor and actuator information (i.e. their corresponding network units' activation) is memorised for a fixed duration (the memory duration is determined by the decay rate of unit activation along the self-connection on the unit). This allows association to be made between time delayed presentations of two inputs to the network. Association results in the record (in the time parameter) of the time delay between the two input presentations (i.e. the time delay between a robot's action and the consecutive robot's perception of a change in its sensor measurement, e.g. compass measurement).

$$tp_{ij}(t) = \frac{tp_{ij}(t-1) \cdot \frac{cf_{ij}}{a} + \frac{y_i(t)}{y_j(t)}}{\frac{cf_{ij}}{a} + 1} \quad (1)$$

where cf is confidence factor; tp is time parameter; x is input and i, j are units nm.

$$cf_{ij}(t) = cf_{ij}(t-1) + a \quad (2)$$

Figure 1 Schematic representation of the robot's controller



$$y_i(t) = F(x_i(t) + tp_{ii} \cdot y_i(t-1) + \sum_{j \neq i} G(tp_{ij}, cf_{ij}, y_j(t-1))) \quad (3)$$

$$y_i(t) = \begin{cases} 1 & \text{if } x_i(t) > 0 \\ tp_{ii} \cdot y_i(t) & \text{otherwise} \end{cases}$$

or

$$\sum_{j \neq i} G(tp_{ij}, cf_{ij}, y_j(t-1)) > 0$$

$$G(tp_{ij}, cf_{ij}, y_j(t-1)) = \begin{cases} 1 & \text{if } |tp_{ij} - y_j(t-1)| < e \\ & \text{and } cf_{ij} > \frac{\max_{\{y_j > 0\}}(cf_{ij})}{T} \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

$$F(x) = x \text{ if } x \leq 1$$

Learning is based on a statistical process of elimination among all associations made across all the sensor and actuator state space of the robot. That is, the most often associated input pairs (the rate of association is measured by the confidence factor parameter) are considered as correctly correlated. In the connectionist terminology this means that retrieval of the units' activity in the network results from a winner-take-all mechanism on the connection parameters. Equation (3) gives the unit activation function. The factor G encapsulates the winner-take-all conditions, which can be paraphrased as follows: a unit is activated by other active units, with which it has a non-zero connection:

- (1) if the time delay since the other unit's activation is equal within some error margin to the memorised temporal correlation (time parameter); and
- (2) if the confidence factor of the connection is sufficiently important (over a fixed threshold).

The model allows learning of spatio-temporal regularities in the robot's perception, by making spatio-temporal associations across all the robot's sensor and actuator modalities. Spatial regularities of unit activation (which represent correlations between different sensor-actuator states) are recorded through the confidence parameters. The temporal regularity of these units' correlation (time delay between each module state change) is recorded in the time parameters. Because learning, i.e. update of the network connection parameter, and control, i.e. retrieval of the units' activity, requires only one time step of computation, the model is fast and therefore allows on-line learning and control of autonomous robots.

Predefined knowledge can easily be implemented from the start or added in the course of the robot's normal operation, by predefining or changing the network's connectivity. This knowledge of (spatial and/or temporal) correlations between the robot's specific sensor-actuator modalities can be given by fixing the connection parameters between the network's units corresponding to these sensor-actuator systems. This transmitted information is then directly incorporated into the robot's current knowledge, which is represented by the network's connectivity state. A priori information can easily be transmitted to the robot via downloading of parameter values during run time, or by adding these values to the robot's controller as a supplementary software. Similarly, one could imagine wiring together two robots in order to directly transfer one robot's experience, i.e. its network's structure, to another robot. However, this could only be possible when the two robots are sufficiently similar (in their internal functioning and sensor capabilities) for the experience of one robot to be meaningful to the other. In the opposite case (as in human-robot communication), such a data transfer would not be possible and then external symbolic communication as described by this work would be necessary.

The transferred information can represent the results of the robot's learning during a previous run or the programmer's knowledge (when it is easily translatable to the robot). For instance, in the experiments reported here, predefined behaviours for obstacle avoidance in the vehicle robots and infra-red phototaxis, i.e. tracking of infra-red light (for the imitative behaviour of the doll robot), were predefined by fixing the connection parameter values between the robot's motors and the robot's bumpers and infra-red sensors so as to perform the desired behaviours. Similarly, in the experiments in which teaching of the vocabulary was given by a teacher robot (Billard and Hayes, 1998; 1999) (see below), the teacher robot's a priori knowledge of the vocabulary was determined by fixing the connection parameter values between the robot's radio sensor (words are radio bit-strings) and motors (for words describing the robot's movements).

In addition, learning, i.e. update of the network's connectivity for representing the measured correlations can easily be biased by erasing the network's connections between

particular sensor-actuator modalities, when one does not expect relevant correlations to be found between these systems. Downloading of a priori information or biasing the learning process might be useful in order to speed up the learning process, as they can reduce the amount of irrelevant information to discard before achieving learning (as learning results from a statistical process of elimination between correct and incorrect associations). For example, a spot welding robot may come pre taught with knowledge of how to spot weld and how not to hit the object to be welded, leaving just the location of each weld to be learnt. This would immediately bypass about 90 per cent of the current teaching effort.

Generally, by using software packages or data download, a robot could inherit knowledge which need not be learned from scratch, e.g. a spot welding robot may be delivered with a software package with knowledge of how to spot weld and how not to hit the object to be welded, leaving just the location of each weld to be learnt. The amount of teaching necessary to adapt the robot to its specific task and environment could in this way be significantly reduced.

Experiments on landmarks recognition and labelling

We implemented the DRAMA architecture in a number of experiments with wheeled LEGO or FISCHERTECHNIK robots which are widely used tools for research on mobile robots. In these experiments the robot learned to recognise landmarks (boxes, light, aluminium covered ground areas). That is, the robot learned the locations of different objects, by associating each object's features with its polar co-ordinates, which the robot calculated from measurement with an odometric sensor and a compass. The robot learned also to label the landmarks following a vocabulary, i.e. set of radio signals, which it was taught by another robot in one case (Billard and Hayes, 1999) and by a human instructor in another case (Billard and Dautenhahn, 1998). Teaching was done as part of a following strategy, whereby the learner robot followed its instructor around the environment. The robot then learned to associate the radio signal, emitted by the instructor when facing one of the objects, to its own observation of the object (i.e. through its sensor measurements characterising the object's features and position). The following strategy, leading to the imitation of the teacher's robot movements by the learner robot, was

used to teach the learner robot a vocabulary to describe its actions (turn, move, stop), its orientation relative to the four cardinal points (Billard and Hayes, 1998) and its inclination, when moving up/down a hilly environment (Billard, 1998).

Experiments on teaching a doll robot

The DRAMA architecture was implemented in a doll-shaped robot. The hardware of the doll robot is made partly of plastic parts (arms, legs and head) which were taken from a commercial doll and of LEGO pieces, which form the central part of the body (see Figures 2a and 2b).

The robot has three motors, for moving each arm and the head separately. The robot is provided with five touch sensors (electrical switches), placed under the feet, inside the hands and the mouth, a tilt sensor which measures the vertical inclination of the body (it distinguishes between horizontal and vertical positions) and four infra-red (IR) detectors. Each infra-red detector consists of an emitter and a receptor. Two of the IR receptors are placed on the robot's chest and measure the signals of the corresponding IR emitters which the demonstrator holds, one in each hand. The signal of each sensor is used to control each of the robot arms. That is, when the demonstrator moves his/her left arm in front of the robot, the right detector on the robot is activated, which triggers the robot's right arm waving movements, i.e. lifting the arm up and then down with a fixed time interval between the two movements. The two other IR emitters are placed on the robot's ears, while the two corresponding IR receptors are mounted onto a pair of glasses which the demonstrator wears. Phototaxis performed on the two IR signals is used to direct the robot's head. That is, when the demonstrator looks, e.g. to the left, the left detector on the glasses receives full activation while the right one receives none, which triggers the robot's head movement to the left, i.e. the robot turns the head to the left (and vice versa for the right). After a fixed time delay (about half a second), the robot moves the head back to facing the demonstrator. As a result, the robot appears to mirror the demonstrator's arm and head movements.

Figure 3 shows a schematic representation of the sensor positions on Robota's body.

The robot is provided with a simple communication system, which consists of a keyboard and a loudspeaker (which is a

Figure 2 (a) The LEGO structure of the robot doll with the mini keyboard; (b) The doll dressed up, with, on the right side, the feeding bottle, the glasses and the infra-red hand




(a)



(b)

commercial pocket recorder with 20 seconds recording time). The pocket recorder is used by the demonstrator to record sounds, spoken words or sentences (a sound slot lasts for two seconds), which correspond to the conceptual meaning the demonstrator attaches to each key of the keyboard. The demonstrator communicates with the robot by pressing the keys on the

Figure 3 Schematic representation of the sensor positions on Robota's body. Correspondences between movements of the demonstrator's arm and head and Robota's recognition of these movements are indicated



keyboard, each key representing a different word; the robot answers the demonstrator by reading back the sound slots of the corresponding words of the keyboard.

Experiments were carried out in which the robot was taught to perform and label sequences of actions, that is sequence of the robot's head and arm movements. In order to teach the robot, the demonstrator first performs herself the sequence of actions, moving sequentially her arms and head, which the robot immediately imitates in response to the infra-red reception of the sensors attached to the glasses of the demonstrator and the sensors which the demonstrator holds in her hands. At the end of the sequence, the demonstrator presses a key on the keyboard. The robot associates the pressing of the key with the complete sequence of actions which it has memorised (by property of association of the DRAMA architecture, as explained earlier). The demonstrator repeats the same process for each key, teaching a different action sequence each time. During the training, the demonstrator verifies if the robot has correctly learned the action sequences, by pressing again the keys. This activates retrieval of the associated action sequence in the DRAMA network, which leads to the sequential activation of the robot's arms and head motors. Note that because the time delay between each action is recorded in the network's time parameters, each action of the sequence is retrieved after the same time delay as observed during the demonstration. Results showed complete success for all teachings. That is, the robot learned correctly to perform the different

sequences taught by each of the instructors. Each sequence pattern was learned after one to three trials, depending on the precision of the demonstrator's movements. It took between ten to 15 minutes to teach the robot eight sequences. An example of teaching of the eight sequences is given in Figure 4.

Further, experiments were carried out, in which the robot was taught English proto-sentences, i.e. sequential typing of the keys of the keyboard (where each key represented a different word), for describing its actions and perceptions. Examples of taught sentences are: "You touch foot left" (four keys), which correspond to pressing the switch on the robot's left foot; "I move head right" for the robot turning the head to the right. Proto means that the sentences are incomplete; the determinants are missing and the plural conjugations are not respected. However, the grammatical ordering of the words is respected.

For teaching the robot, the demonstrator first presses the relevant switch or activates the robot's arm or head by moving the corresponding sensors, and then presses the corresponding keys on the keyboard, one after the other one, such as to conserve the grammatical order. The demonstrator checks if the robot has correctly learned by pressing the switch again or activating the robot's arm or

head. This results in retrieving the keys sequence, resulting in reading sequentially the loudspeaker's recorded slots, corresponding to each key, by reading backwards the DRAMA connections from the sensor-actuator inputs to the keyboard switches inputs. Note that as for the action sequences of the first experiment, the order of the keys sequence is conserved and so the grammatical order of the sentence it represents.

Note that the architecture is not limited to learning only eight action sequences or combinations of eight words. Only the hardware, that is the keyboard, was limited to contain only eight keys. There is no a priori limitation on the number of words or patterns it could learn; this depends on the number of inputs to the network, which fixes the maximal capacity of the network (Billard and Hayes, 1999). The reason for using a "mini" humanoid robot (a doll) was economical. It would now be interesting to implement the architecture in a more complex humanoid robot with higher degrees of freedom and better sensor capabilities.


Conclusion

This paper presented a novel connectionist architecture, DRAMA, used for online learning and control of an autonomous robot. The model allows learning of spatial regularities and time series in a robot's perceptions. It was successfully implemented in experiments for landmark recognition and labelling of an autonomous robot.

The DRAMA architecture shows characteristics which make it very relevant for all types of autonomous robotic agents. These characteristics are:

- (1) fast computation for the model to react in real-time, which was shown by the model's implementation for online and on-board learning of a computationally limited robot;
- (2) robustness and adaptability in the face of varying environmental constraints, demonstrated by successful learning of a robot when faced with a proportion of 30 per cent of noisy/incorrect data (Billard and Hayes, 1999);
- (3) as little built-in knowledge as possible to keep the system unspecific to a particular type of implementation (task, agent or environment), which was verified by the implementation of the model in different

Figure 4 Examples of a teaching of eight action sequences



types of robotic agents (simulated and real vehicles (Billard and Dautenhahn, 1998; Billard and Hayes, 1998) and a doll-shaped robot).

In this paper, a novel implementation of the DRAMA model into a doll-shaped robot was presented, for teaching the robot a synthetic proto-language. The experiments investigated the use of an imitative scenario, whereby the robot mirrors the arm and head movements of a human instructor, for teaching the robot to perform different sequences of actions. Imitation, communication and learning are important skills to possess by a robot expected to interact with humans through daily tasks. The doll robot is provided with these three competencies, which are completely controlled by the DRAMA architecture.

Further reading

In order to allow easy readability of the article by a large audience, the description of the robot's control architecture and of the experiments was kept qualitative. A mathematical description and formal analysis of the DRAMA architecture can be found in Billard and Hayes (1999). A complete technical report of the experiments with *Robota* is reported in Billard (1998). Experiments on landmark recognition and labelling are reported in Billard and Dautenhahn (1998) (teaching done by a human instructor) and in Billard and Hayes (1999) (teaching done by a second robot). Experiments on learning a vocabulary for its actions and orientation are given in Billard and Hayes (1998). A complete technical report of all experiments and of the development of the DRAMA architecture is given in the author's PhD thesis (Billard, 1998).

References

- Billard A. (1998), "DRAMA, a connectionist model for robot learning: experiments on grounding communication through imitation in autonomous robots", PhD thesis, Department of Artificial Intelligence, University of Edinburgh, submitted on 30 September.
- Billard, A. and Dautenhahn, K. (1998), "Grounding communication in autonomous robots: an experimental study", in Recce, M. and Nehmzow, U. (Eds), *Robotics and Autonomous Systems Journal*, Special Issue on *Scientific Methods in Mobile Robotics*, Vol. 24 Nos 1-2.
- Billard, A. and Hayes, G. (1998), "Transmitting communication skills through imitation in autonomous robots", in Birk, A. and Demiris, J. (Eds), *Learning Robots: A Multi-Perspective Exploration*, LNAI Series, Springer-Verlag.
- Billard, A. and Hayes, G. (1999), "DRAMA, a connectionist architecture for control and learning in autonomous robots", *Adaptive Behaviour Journal*, Vol. 7 No. 1, Society of Adaptive Behaviour.
- Blumberg, B. and Todd, P. and Maes, P. (1996), "No bad dogs: ethological lessons for learning", *From Animals to Animats 4, Proceedings of the 3rd International Conference on the Simulation of Adaptive Behaviour*, The MIT Press/Bradford Books.
- Demiris, J., Rougeaux, S., Hayes, G.M., Berthouze, L. and Kuniyoshi, Y. (1997), "Deferred imitation of human head movements by an active stereo vision head", *Proceedings of the 6th IEEE International Workshop on Robot Human Communication*, IEEE Press, Sendai, Japan, pp. 88-93.
- Kuniyoshi, M.I. and Inoue, I. (1994), "Learning by watching: extracting reusable task knowledge from visual observation of human performance", *IEEE Transactions on Robotics and Automation*, Vol. 10 No. 6, pp. 799-822.
- Lueg, C. and Pfeifer, R. (1997), "Cognition, situatedness, and situated design", *Proceedings of the Second International Conference on Cognitive Technology (CT 97)*, Aizu, Japan, 25-28 August.
- Maes, P. (1995), "Artificial life meets entertainment: interacting with lifelike autonomous agents", *Special Issue on New Horizons of Commercial and Industrial AI*, Vol. 38 No. 11, November, *Communications of the ACM*, ACM Press, pp. 108-114.
- Mataric, M.J. (1994), "Learning motor skills by imitation", in Stanford University (Ed.), *AAAI Spring Symposium toward Physical Interaction and Manipulation*, 21-23 March.
- Scassellati, B. (1998), "Imitation and mechanisms of joint attention: a developmental structure for building social skills on a humanoid robot", in Nehaniv C. (Ed.), *Proceedings of Computation for Metaphor, Analogy, and Agents* conference, Springer-Verlag series.
- Schaal, S. (1997), "Learning from demonstration", in Mozer, M.C. and Jordan, M. and Petsche, T. (Eds), *Advances in Neural Information Processing Systems*, Vol. 9, MIT Press, Cambridge, MA, pp. 1040-6.