Experiments on human-robot communication with **Robota,** an imitative learning and communicating doll robot.

Aude Billard*,
Kerstin Dautenhahn†,
Gillian Hayes*
*Dept. of Artificial Intelligence, University of Edinburgh, 5 Forrest Hill, Edinburgh EH1 2QL
†Department of Cybernetics at the University of Reading
PO Box 225 Reading, RG6 6AY United Kingdom
Contact:
audeb@dai.ed.ac.uk
Tel: +44-131-650-4491
Fax: +44-131-650-6899

Abstract

Imitation and communication behaviours are important means of interaction between humans and robots. In experiments on robot teaching by demonstration, imitation and communication behaviours can be used by the demonstrator to drive the robot's attention to the demonstrated task. In a children game, they play an important role to engage the interaction between the child and the robot and to stimulate the child's interest. In this work, we study how imitation skills can be used for teaching a robot a symbolic communication system to describe its actions and perceptions.

We report on experiments in which we study human-robot interactions using a doll robot. **Robota** is a robot, whose shape is similar to that of a doll, and which has the capacity to learn, imitate and communicate. Through simple phototaxis behaviour, the robot can imitate (mirror) the arms and head's movements of a demonstrator. The robot is controlled by a Dynamical Recurrent Associative memory Architecture (DRAMA), which allows learning of time series. We carry out experiments, where the robot is taught to perform different sequences of actions and to label these action sequences with different 'names'. In further experiments, the robot is taught combination of words, which form grammatically correct sentences, to describe its actions and perceptions of touch on different parts of its body. Finally, we carry out tests with 5-years old children, who teach the robot words to label different parts of its body and simple action sequences.

Results demonstrated the validity of the imitative strategy for teaching a robot complex sequences and combinations of sensor and actuator inputs. In particular, tests with children suggest that the imitative and communicative behaviours of Robota makes it an interesting toy for children. Moreover, if Robota was provided with more complex actuator capabilities, it would also make an interesting robotic platform for research on human-robot interaction and especially on robot learning from demonstration and by imitation.

1 Introduction

Social intelligence relates to behavioural skills which animals acquire and use through their social interactions with conspecific agents. Recently, people began investigating the design of social skills, such as imitation and communication, for a robot as a means for interacting with and, in particular, teaching the robot [11, 10, 12]. *Learning by imitation or observation* is thought to play a key role in the development of social skills in primates and humans [15, 17]. Similarly, imitation is an interesting means for guiding the robot attention. It can be used to teach the robot new motor skills by having it observing and then imitating the actions of that of a demonstrator. E.g. in [11, 8], a robot learns to perform context dependent action patterns, by following and thus replicating the movement of a second robot. In [7, 16, 13], a robot learns to perform sequences of actions by observing and then reproducing those of a human demonstrator. Social interactions, such as imitation, which create a coordination between the interactive agents, are fundamental to the development of linguistic competences in children [9, 14]. In previous and present work, we study how imitation skills can be used for teaching a robot a symbolic communication system to describe its actions and perceptions.

Inspired by psychological and ethological studies of the development of communication skills in children and primates, we determined a number of key features for the cognitive and behavioural mechanisms of our agent [2]. We require for the agent to possess associative capabilities for spatio-temporal association across multiple sensor-actuator modalities and behaviour skills for coordination of actions with that of a second agent. We developed a control architecture, DRAMA (Dynamical Recurrent Associative Memory Architecture), made of a connectionist model, which provides the necessary cognitive mechanisms for imitative learning and communication in robots [5]. We implemented the model in a number of teacher-learner experiments with physical and simulated mobile robots being taught either by another robot or by a human [2, 4, 5, 6]. The experiments demonstrated the robustness of the architecture for extracting spatio-temporal regularities in an very noisy (highly variable) data set. In a first set of experiments, we studied transmission of a vocabulary to label actions and perception of a robot [4, 6]. Further, we carried out simulation studies, in which we investigated transmission of the vocabulary among a group of robotic agents [2]. In a third set of experiments, we investigated learning of a sequence of the robot's perceptions, while wandering in a physically constrained environment (series of corridors) [5].

In this paper, we report on a new set of experiments, using a different robotic platform, namely a doll-like robot called *Robota*, for studying issues related to human-robot interaction. This new implementation of the DRAMA architecture aims 1) to verify the computational capacity of the model at learning highly redundant combinations and sequences of inputs and 2) to test whether using a robot which has a familiar and somewhat 'cute' appearance might prove appealing (rather than frightening) for people to interact with.

Robota is a robot, whose shape is similar to that of a doll, and which has the capacity to learn, imitate and communicate. Using a simple phototaxis behaviour, the robot can imitate (mirror) the arms and head's movements of a demonstrator. The robot is controlled by a Dynamical Recurrent Associative memory Architecture (DRAMA), which allows learning of time series. We carry out experiments, in which the robot is taught to perform different sequences of actions and to label these action sequences with different 'names'. In further experiments, the robot is taught combination of words, which form grammatically correct sentences, to describe its actions and perceptions of touch on different parts of its body. Finally, we carry out tests with 5-years old children, who teach the robot words to label different parts of its body and simple actions sequences. Note that the aim of these tests was not to provide a psychological study of the children's behaviour or social intelligence, but rather to give a qualitative evaluation of the children' understanding and interest of the game we proposed.

This paper is divided as follows. We first describe the doll robot's hardware and controller. The control architecture is in great detail described in previous publications, therefore we only give a short summary in this paper. We then report on the experiments we carried out both with adults and children. We conclude this paper by discussing the experiments' results.

2 The doll robot

The hardware of the doll robot is made partly of plastic parts (arms, legs and head) which were taken out of a commercial doll and of LEGO pieces, which form the central part of the body. The robot controller is made of a micro-controller with 512k byte EPROM space and 128kS byte Static RAM. The CPU (central Processing Unit) is a Phillips 93C100 series 68000 compatible running at 30 Mhz. Figure 1 show pictures of the doll robot.

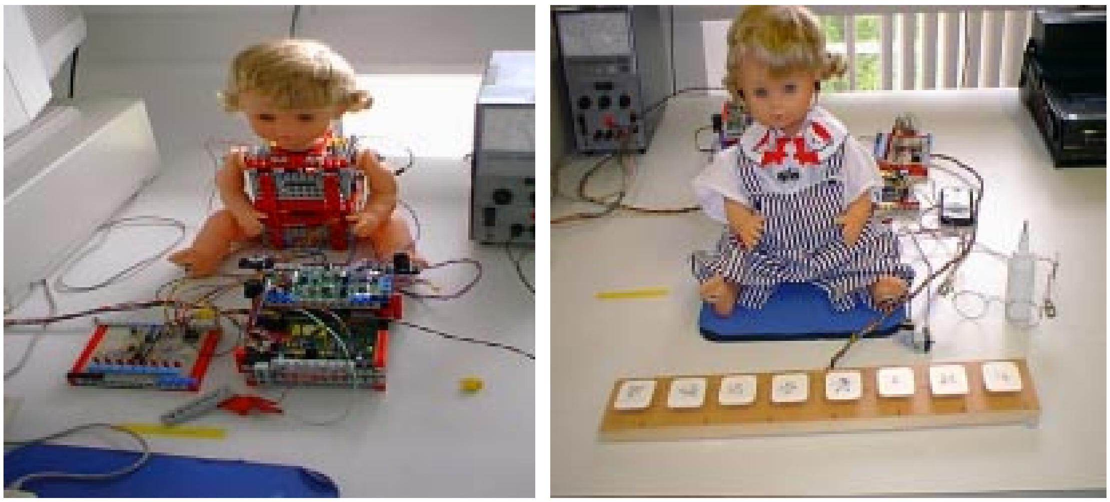

Figure 1: *Left:* The LEGO structure of the robot doll with the mini keyboard; *Right:* The doll dressed up, with, on the right side, the feeding bottle, the glasses and the infra-red hand sensors.

The robot has three motors, for moving each arm and the head separately. It has also a LEGO sound emitter, which is used to simulate the robot crying. 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 detectors consists of an emitter and a receptor part. 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 arm. That is, when the demonstrator moves the left arm in front of the robot, the right detector on the robot is activated, which triggers the robot right arm waving movements (going up and down with a fixed interval). 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 (and vice-versa for the right). As a result, the robot appears to mirror the demonstrator arms and head movements. Figure 2 shows a schematic representation of the sensor disposition on Robota's body.

The robot is provided with a simple communication system, which consists of a keyboard and a loudspeaker (pocket recorder). We use, in fact, two keyboards containing 8 keys, a mini one, as shown on figure 1 left, with a set of eight light bulbs on top of it and a big one as shown on figure

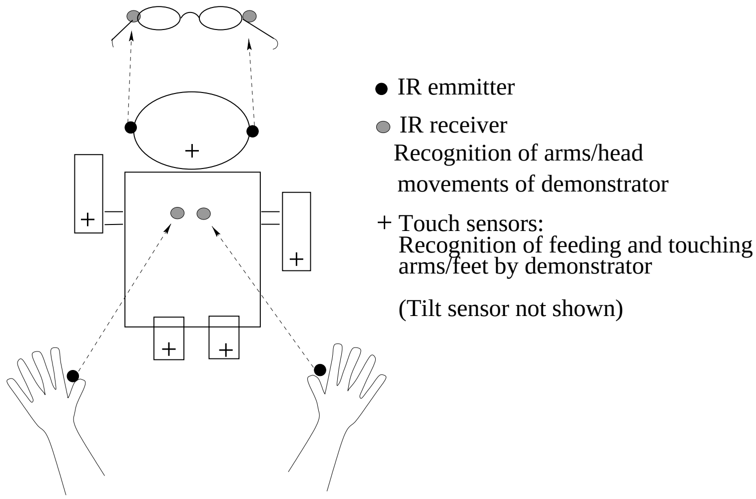

Figure 2: Schematic representation of the sensors disposition on Robota's body.

1 right. Big and small keyboards have the same functionality, the big one was used with the children, as it is easier to manipulate. The pocket recorder is used by the demonstrator to record sounds, spoken words or sentences (a sound slot last for 2 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 keyboard, each key representing a different word; the robot answers to the demonstrator by reading back the sound slots of the corresponding words of the keyboard.

3 The robot's controller

The robot's controller consists of a Dynamical Recurrent Associative Memory (DRAMA), a complete description of which can be found in [5]. 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-recurrent 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. Inputs to the network consist of the robot's sensor and actuator information, which is encoded in binary (0/1) bit-strings. In the experiments, the network is updated continuously whenever a variation (one-bit activation or deactivation) in a sensor or actuator input is measured. Retrieval of the network output units' activity proceeds from a Winner-take-all mechanism. A unit is activated by other active units, with which it has a non-zero connection, (i) if the time delay since the other unit's activation is equal to the memorised temporal correlation (time parameter), (ii) if the confidence factor of the connection is sufficiently important (over a fixed threshold) and (iii) if the two first conditions are satisfied by all active units, with which the unit has a non-zero connection. In the experiments, retrieval of the network outputs is used to determine the robot's behaviour at each time step, that is activation of the robot's motors for moving arms and head and the robot's speaking (activation of the lights on the mini-keyboard and reading of the recorder).

As an associative memory, the network has a maximal capacity of order 2 of its size (number of units). Using the full capacity of the network, retrieval performance is perfect up to a proportion of 30% of noisy data. The model allows learning of time series. In particular, it allows learning of overlapping sequences, that is sequences composed of one or more common patterns. This combinatorial property of the network is used in the experiment to teach the robot different sequences of actions, which contain the same four basic actions, and different word combinations, formed with the same basic eight words, to describe its actions with grammatically correct sentences. The robot's imitative behaviour, which results from phototaxis on the infra-red sensor measurements, was implemented by predefining the network's connections between infra-red sensors and corresponding actuators. The random movements of the robot's head and arms in response to the child's touch and the crying behaviour, in the experiments with children (see section 5), were defined as separate processes.

4 Experiments on learning dance patterns and word combinations

The potential of the robot's learning controller (DRAMA architecture) at learning complex times series of sensor and actuator patterns was evaluated through two experiments, which we report in the following. Each test was performed by 5 different persons.

In the first experiment, the robot was taught four to eight 'dance' patterns, that is different sequences of head and arms movements. In figure 3, we show an example of teaching of eight dance patterns. In order to teach the robot, the demonstrator first performs herself the dance, moving sequentially her arms and head, which the robot immediately imitates in response to the infra-red emission of the sensors attached to the glasses of the demonstrator and the sensors which the demonstrator holds in her hands (see explanation in section 2). At the end of the dance, the demonstrator presses a key on the keyboard. The robot associates the pressing of the key (which activates the corresponding sensors unit) with the complete sequence of actions which it has memorised, that is, it updates the parameters of the connections linking the sensor unit corresponding to the key to the actuator units which have been activated by the different actions of the dance (recall that each unit keeps a memory of its activation for a fixed delay, see explanation in section 3). For each key, the demonstrator teaches a different dance pattern. The demonstrator verifies if the robot has correctly learned the dance patterns, by pressing again the key. This activates retrieval of the learned actions sequence by reading backwards the DRAMA connections from the keyboard switches inputs to the robot's actuator outputs. 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 (see the two criteria for the retrieval of unit activity described in section 3). Each dance pattern was learned after 1 to 3 trials, depending on the precision of the demonstrator's movements. The eight sequences were taught in about 10 minutes.

Learning is statistical, that is, correct association between the key and the corresponding action is learned once this has been recorded (in the confidence factor parameter, see section 3) more often than other combinations relative to a given threshold. Teaching errors, such as pressing the wrong key or showing the wrong actions, would not fail the learning but only delay it, as the correct sequence should be repeated a number of times until the incorrect associations are discarded. Note that the success of the learning in this experiment depends strongly on choosing correctly the value of the short term memory (memory of the unit activation), so that it is long enough to allow association of the complete sequence and short enough to distinguish between two sequences. In [3], we report on a quantitative evaluation of the influence of this parameter on the success of the experiment and in [5] we propose an algorithm for on-line tuning

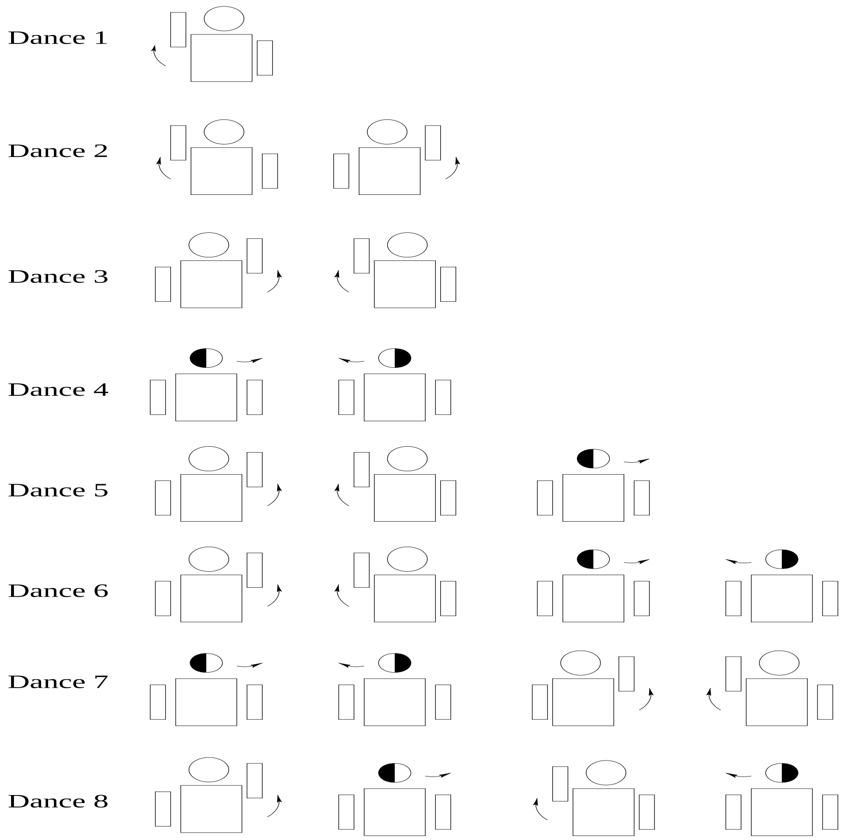

Figure 3: Examples of a teaching of eight dance patterns

of this parameter together with the learning of the connections parameters.

In the second experiment, the robot is taught combinations of keys to describe its actions and its perceptions of touch on its body parts (hands, feet and mouth). The demonstrator was asked to choose eight words among eleven, which were *I, You, move, turn, touch, arm, foot, head, left, right, up/down* and to stick each word on top of each key on the keyboard. The demonstrator was then asked to teach the robot to describe its actions (moving arms and head) and its perceptions of touch (pressing of the switches) on its feet, arms and mouth, by associating them with a combination of keys. These combination of keys represent grammatically correct sentences. An example of teaching was *I move right arm, I turn head left, You touch mouth, You touch left/right foot.* In order to teach 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 preserve 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 (which it shown by activating sequentially the light bulbs placed above each of the keys on the mini-keyboard, see picture 1), by reading backwards the DRAMA connections from the sensor-actuator inputs to the keyboard switches inputs. Note that as for the actions sequences of the first experiment, the order of the keys sequence is conserved and so the grammatical order of the sentence it represents.

Learning in this experiment results from the same principle as in the previous experiment on teaching dance patterns. The difference lies on the direction of association. In the dance experiment, we associated activation of key units to sequential activation of actuator units, while, in the second experiment, we associated activation of actuator or switch units to sequential activation of key units. Because actuator, keyboard and body switches information is processed similarly in the associative memory (training and retrieval processes is the same for all units, independently on the original type of the information they represent), correct learning and retrieving of the taught sequence, in the second experiment, results from the same mechanism as described for the first experiment.

An interesting result of this experiment was that, although the robot has been taught only complete sentences, it has learned implicitly the conceptual meaning of each of the words of the sentence. For instance, after having taught the two sentences *you touch left arm* and *I move right arm* for pressing the switch on the robot's left arm and activating the robot's movement of the right arm, both pressing the robot's left arm and moving its right arm results in the robot lighting up only the key corresponding to the word *arm*. This results from the Winner-take-all retrieval mechanism, which requires that all activated units agree in their vote for another unit activation. That is, a unit is activated only if all activated units satisfy the two criteria for activation retrieval of this particular unit (see section 3). In the previous example, the arm unit is the only unit for which both the switch and actuator units agree. Similarly the concepts of left or right can be retrieved if one moves the robot's left arm together with touching the switch on the left foot, after having taught the two sentences *I move left arm*, *you touch left foot*. From the eight words combinations which the demonstrators were given, they were able to retrieve the concepts for *I move* associated with activation of the actuators, *you touch* pressing of a switch, *arm, foot, left* and *right*. In addition, the robot had learned implicitly the sequential order of words occurrence, as those coming after a short or long time delay. From teaching *I move left arm* and *I move head right*, the robot inferred *I move right arm*; that is, it retrieved the correct word combination, in the correct order, although it had not been taught this sequence previously. However, the words *right* and *arm* where retrieved almost simultaneously, as they both appeared at the end of the sentences on the example. Note that the architecture is not limited to learning only eight dance patterns 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 (see [5]).

5 Tests with children

In the following section, we give a narrative account of the tests we carried out with children. It is by no means meant to be a psychological analysis of the children behaviour, but rather it aims to give a qualitative assessment of the children appreciation of the game. Through these tests, we wished to evaluate the potential of a robotic doll as a toy for children and in particular to evaluate the children' understanding and interest in the particular game we proposed.

The robot doll was tested with nine children of five years old (six boys and three girls[1]). The scenario of the game the children were asked to play consisted of the following:

1. The child is presented the doll as a young baby, that does not know how to speak yet. The task of the child is then to teach the doll to speak. The child is shown the keyboard; the meaning of each key is described by an icon on top of it. A word is prerecorded in the pocket recorder for each key: (from left to right) 'food', 'rocking', 'hello' (right arm up),

[1]This unequal proportion of boys to girls was involuntary.

'Yeah' (two arms up), 'No' (side movements of the head), 'hand', 'left foot', 'right foot'. The child is explained that (s)he has to teach the doll the words on the keyboard.

2. The child is shown how (s)he can direct the movements of the doll's arms and head by moving his/her own arms and head, while holding the two infra-red emitters in each hand and wearing the glasses (the children liked very much wearing glasses!).

3. After five minutes of playing, the doll begins to cry (emitting a sharp continuous sound). The child is told that the sound meant that the doll is crying and that it cries because it wants something. "Has the child any idea what the baby-doll could want?" (All children responded immediately that 'it wants to eat'). The child is then shown the baby feeding bottle and asked to feed the doll (which they did with great pleasure). After some seconds of feeding, the baby stops crying and the child is congratulated to have found out what the baby-doll wanted.

4. The child is then asked to teach the doll how to say that it wants to eat. The child is shown how to teach the doll using the keys on the keyboard (first feed, then press the first key on the left). The child is asked to feed the doll again, which prompts the doll to say 'food' (retrieving of the correct association between the mouth switch and the key, which activates reading the corresponding word on the recorder). The child is then congratulated because (s)he has managed to teach the doll.

5. The child is then asked to try to teach the doll another word of the keyboard. (Children spontaneously tried to teach the words for hand or left/right foot.) The child is asked to teach the doll to say 'hello'. For this, (s)he must first wave the right arm, which prompts the doll to wave its left arm, and then press the corresponding key on the keyboard. The child is then shown that if he/she presses the key 'hello' again, the doll waves its arm and then say the word 'hello', showing that it has well learned the correspondence between the key, waving the left arm and saying hello. The child teaches the robot further to wave both arms and say 'yeah' and to shake the head and say 'no'. (The experience had to be repeated two or three times before the child would understand that the robot was actually learning the relationship between actions and saying the words).

6. The doll begins to cry again and says 'food'. The child feeds the doll and the doll stops crying.

7. The child carry on teaching other words on the keyboard.

8. The doll begins to cry again. (All children immediately said that the doll wanted to eat and then tried to feed it. However, the doll did not stop crying.) The child is then told that the doll probably wants to be rocked. The child is shown how to gently rock the doll. After a couple of rocking movements, the doll stops crying. The child is then asked to teach the doll to say the word 'rocking', by first rocking the doll and then pressing the corresponding key.

9. The doll begins to cry again and says 'food'. The child feeds the doll and the doll stops crying. End of the game.

At the end of the game, all children were asked the following questions: Did you like the game? To which they all, apart from one, answered a definitive 'yes'. The child, who did not like it, explained to his father that he did not like it because he had difficulty to understand the words the doll was saying and to press the switches on the doll's body (critics related to the hardware rather than the game itself), he however very much liked feeding the doll and having it moving when pressing the keys, as it reminded him his 'Tamagoshi' (Japanese children game; the child was Japanese). Do you have a doll at home? Only three of them did not have one. Is this doll different from another doll? They all said no (probably thinking of the shape of the doll, see answer to next question). Do you like this doll? They all said 'yes'. Is this doll better than other dolls? All minus the one who did not like the game said 'yes'. Why? Answers were 'Because it can do things', 'Because it moves/speaks'.

Figure 4: Kyo Takiguchi 'feeding' Robota

The children took a real pleasure in the fact that the doll was responding to their touching her. Little behaviours as small random movements of the head and arms has been implemented as a reaction to the child touching the switches on the doll's hands, feet and mouth. The robot would then appear to sometimes wave the head from left to right when the child tried to feed it, as if it did not wanted to be fed. The children liked also the fact of being able to command the robot's arms and head movements, although they took less pleasure in this than in touching the robot. One reason is probably that they do not control well enough their own movements yet and had thus problems to direct correctly the robot's movements. For the robot to respond correctly to head and arms' movements, you had to make very precise movements (lift the arms straight and in a narrow region around the middle of the robot's body, and small slow side movements of the head). The children had problems to achieve such precision. It was not clear whether all children had really understood the learning process behind the robot's speaking, that is the relationship between their touching the robot, pressing the keys and the robot's consequent speaking, especially as some children would press the wrong key, which would result in the robot speaking a different word than the expected one. These observations suggest that the game reached the limits of competence of so young children, as it required from them to concentrate for a long period of time on a repetitive task. However, we should mention that most of the children were keen on continuing playing with the doll after they had been told that the game was finished (each game lasted for about 20 minutes; this limit was chosen in order not to exceed the concentration capacity of young children). None of the children had had time to teach the entire vocabulary during the game, nor to understand fully all the different capabilities of the robot (especially to understand all the correlations between their acting on the robot and the robot's movements and speaking). The complexity of the robot seemed to intrigue them rather than to intimidate them.

If time allows, further tests will be done with older children and again with young children but in using a simplified program (no keyboard, teaching only dance patterns). An interesting test would be to have the robot staying in the nursery for several hours or days and observe the progression of understanding and interest of the children over this particular game.

6 Discussion

Experiments carried out with adults showed that the system was able to learn complicated actions sequences and to distinguish between them by labelling them differently. Important was that the sequences could be made of the same set of actions, ordered differently, or of a subset of these actions, ordered similarly or differently than they are in the longer sequences. This combinatorial property of the system was exploited further to teach the robot to associate combinations of words, which formed grammatical sentences, to describe its actions and perceptions of touch on different parts of its body, e.g. *I move arm left, you touch right foot.* By exclusion property of the Winner-take-all retrieval mechanism, concepts of *left, right, I move, you touch, arm, foot* could be extracted from teaching full sentences only. Moreover, the robot would also learn the sequential ordering of words appearance and would reproduce the sentence with the same ordering.

This second experiment demonstrated that the learning architecture could allow learning of a basic 'language' or 'proto-language', which shares some properties with natural language: 1) each word (key in our experiment) can carry a specific meaning; 2) words can be combined and the combination can be given a different meaning while not losing the meaning of each word taken separately; 3) different combinations of the same words can be given different meanings, the meaning of each combination being determined by the order of appearance of each word in the combination; 4) conceptual meanings of each word can be learned implicitly by only presenting them as part of complete sentences, which can then be used to infer new word combinations; 5) precedence between words appearance in the combination is learned and can be used to infer the correct order when constructing a new word combination.

Experiments done with children of 5 years old showed the potential of the system as a game for children. The children enjoyed playing with the robot because they could interact with it in different ways. The robot would answer to the children touching specific parts of its body, by making small movements or little noise. It would mimic the child's head and arms' movements. It would speak words which are part of the child's every-day vocabulary (e.g. *food, hello, no*). Imitation is a game that young children like to play with each other and their parents, thus it was easy for them to understand that they could interact with the robot in this way.

Communication and imitation were the principal means of interaction between the child and the robot. They were easy to understand for the child as they are social interactions, which he has already mastered. More generally, communication and imitation would also be important means of interaction between humans and robots. In the introduction, we mentioned how verbal communication and imitation could be used by the demonstrator to drive the robot's attention to the task she is demonstrating. The experiments we reported here showed how, while imitating the movements of a human demonstrator, a robot could be taught to perform and then label sequences of movements. In our system, learning, communication and imitation behaviours result from the same process, namely training and retrieving of the DRAMA architecture. This results in a computationally fast and cheap system, which allowed its implementation in a computationally limited hardware system. From a robotic point of view, these are important characteristics which makes the system particularly relevant for further implementation for real time control and learning in more complex robotic platforms (that is robots with finer sensor capabilities and more complex actuators, which require more complex computation and are thus very sensitive to the speed of computation).

7 Conclusion

This paper reported on experiments, in which a mobile robot is taught by a human instructor to perform sequences of actions. Teaching is based on an imitation game, in which the robot mirrors the arms' and head's movements of the teacher. The instructor teaches the robot to distinguish between each action sequences by assigning them a different label. The learning capacity of the robot are provided by a connectionist model, DRAMA, which allows learning of time series. The capacity of the model to learning sequences and combinations of inputs was tested further in an experiment in which the robot was taught grammatically correct sentences to describe its actions and perceptions. For theses experiment, we built and used a robot whose shape and features look similar to that of a familiar doll. We tested the reaction of children to the robot by having them playing a simplified teaching game.

Results demonstrated the validity of the set-up, namely the learning architecture and the imitative strategy, to teaching a robot complex sequences and combinations of sensor and actuator inputs. In particular, tests with children suggest that the imitative and communicative behaviours of Robota makes it an interesting toy for children. Moreover, if Robota was provided with more complex actuator capabilities, it would also make an interesting robotic platform for research on human-robot interaction and especially on robot learning from demonstration and by imitation.

8 Acknowledgement

We like to thank anonymous reviewers for helping us to improve a previous version of this paper. Many thanks to Dr. Hamish MacLeod, Mrs Cameron and the children of the nursery of the Psychology department at the University of Edinburgh. Many thanks to Junichi Takiguchi and his son Kyo. A very big thanks to Alexander Colquhoun and Robert MacGregor, technicians of the DAI, for their precious help in constructing the electronic parts of the doll. A very special thanks to Auke Jan Ijspeert for his help in constructing the LEGO body of the doll robot. Aude Billard is supported by a personal grant from the Swiss National Science Foundation.

References

[1] Bandura, A. & all. (1961), 'Transmission of aggression through imitation of aggressive models', *Journal of Abnormal and Social Psychology*,**63**, pp. 575-82.

[2] Billard A. & Dautenhahn K., 'Experiments in social robotics: grounding and use of communication in autonomous agents', (1998), submitted to *Adaptive Behaviour Journal*, special issue on 'simulation of social agents', Dautenhahn K. Guest editor.

[3] Billard A. & Dautenhahn K. (1998), 'Grounding communication in autonomous robots: an experimental study', to appear in *Robotics and Autonomous Systems*, special issue on 'Quantitative Assessment of Mobile Robot Behaviour'. Edited by U. Nehmzow, M. Recce, and D. Bisset.

[4] Billard A. & Dautenhahn K. (1997), 'Grounding communication in situated, social robots', Proc. of TIMR 97, Towards Intelligent Mobile Robots Conference, (Tech. Rep. UMCS-97-9-1), Dept of Computer Science, Manchester University.

[5] Billard A. & Hayes G. (1998), 'DRAMA, a connectionist architecture for control and learning in autonomous robots'. To appear in *Adaptive Behaviour* journal, Vol. 7:2, MIT Press.

[6] Billard A. & Hayes G., (1997), 'Transmitting Communication Skills through Imitation in Autonomous robots'. Proceedings of EWLR97, Sixth European Workshop on Learning Robots, Brighton, July 97. *Lecture Notes in Artificial Intelligence*, A. Birk & J. Demiris editors, Springer-Verlag publishers.

[7] Cooke, S. R. et al., (1997), 'Delayed and Real-Time Imitation of Complex Visual Gestures', Proceedings of the International Conference on Vision, Recognition, Action: Neural Models of Mind and Machine, Boston University, May 28-31, 1997.

[8] Demiris J. & Hayes G., (1994), 'A Robot Controller Using Learning by Imitation', Proceedings of the 2nd International Symposium on Intelligent Robotic Systems, Grenoble, France, July 1994.

[9] Garton A. F., (1992), 'Social interaction and the development of language and cognition', *Essays in developmental psychology*, Laurence Erlbaum Associates, Publishers (USA).

[10] Dautenhahn K., (1998), 'Grounding Agent Sociality: The Social World is its Own Best Model', Proceedings of the 14th European Meeting on Cybernetics and Systems Research, EMCSR'98, pp. 779-784.

[11] Dautenhahn K. (1995), 'Getting to know each other – artificial social intelligence for autonomous robots', Robotics and Autonomous Systems, **16**, pp. 333-356.

[12] Klingspor V., Demiris J., and Kaiser M., (1997), 'Human-Robot-Communication and Machine Learning', *Applied Artificial Intelligence Journal* 11, pp. 719-746, 1997.

[13] Kuniyoshi K., Inaba M. and Inoue H., (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, Dec. 1994.

[14] Lock E.D. (1978), 'Action, Gesture and Symbols: the emergence of language', Academic Press, Harcourt Brace Jovanovich publishers.

[15] Meltzoff A, (1996), 'Towards a developmental cognitive science. The implication of cross-modal matching and imitation for the development of representation and memory in infancy', in *The development and Neural Basis of Higher cognitive Functions*, Adele Diamond (ed.). Annals of the N. Y. Acad. of Sci., **608**.

[16] Schaal S., (1997). "Learning from demonstration." In: M.C. Mozer, M. Jordan, and T. Petsche (eds.), Advances in Neural Information Processing Systems 9, pp.1040-1046, Cambridge, MA: MIT Press

[17] Whiten A. & Ham R., (1992), 'On the nature and Evolution of Imitation in the Animal Kingdom: Reappraisal of a Century of Research', in *Advances in the Study of Behaviour*, **21**,pp. 239-83.