Grounding communication in situated, social robots

Aude Billard* Kerstin Dautenhahn†

5/9/97

Keywords
Social learning - Imitation - Communication - Embodiment - Hilly environment

Abstract

This paper discusses the usefulness of communication as a social skill for embodied robotic agents. We study a teacher-learner situation in a 'meaningful' (hilly) environment. The learner uses the teacher as a model, i.e. learning to communicate means in this case that the learner tries to achieve a similar 'interpretation' of the environment to that of the teacher, on the basis of its own sensory-motor interactions. A simple imitative strategy (following and keeping-contact) is used as the social context in which associative learning takes place. A specific scenario ('mother-child') is proposed and implemented as an example for a situation in which the ability to communicate is advantageous for an individual robot. The communication skill develops during the experiment using a non-supervised teaching approach. Physical experiments to test the validity of the learning architecture are reported. Results are successful in terms of the learning capacities, but they point out the limitation of imitation as a means of learning. However, our implementation of 'social bonding' as a following strategy turns out to be a means to get insight into the other agent's perceptions and, as such, could be a useful starting point towards embodied communication in a social, and ecological 'meaningful' context.

1 Introduction

Communication is a social skill and as such would be desirable for artificial autonomous agents that may be expected to interact with other agents or humans. Following a bottom-up artificial life approach, this work investigates grounding and development of communication. In our work, explicit communication between two agents is an interactive process which is embedded in a social context. Communication in our experiments is not a built-in behaviour but it develops through external teaching, using a teacher-learner setup. We design a social context and a behavioural architecture to control the robots, but we do not select the combination of features which a robot has to learn. Learning to communicate means in our experiments that one robot, the learner tries to achieve a similar 'interpretation' of a situation to that of a second robot, the teacher. Agents can learn to agree on a common 'word' (or 'name') for an object or a situation if they 1) can distinguish it from other sensory information and 2) if both agents interpret the environment simultaneously. Learning to communicate is implemented in our experiments as part of general multiple sensory association processes, produced by DRAMA, the D(ynamic) R(ecurrent) A(ssociative) M(emory) A(rchitecture). [2].

A simple social relationship ('spatial bonding') on the basis of imitative behaviour, namely bidirectional following, is defined in our experiments. A teacher robot teaches a learner robot while the learner

Proc. Conference "Towards Intelligent Mobile Robots", TIMR 97, Manchester 1997.
Technical Report Series, Department of Computer Science, Manchester University, ISSN 1361 - 6161.
Report number UMCS-97-9-1. http://www.cs.man.ac.uk/csonly/cstechrep/titles97.html

*Department of Artificial Intelligence at the University of Edinburgh, 5 Forrest Hill, EH1 2QL Edinburgh, United Kingdom, audeb@aisb.ed.ac.uk.
†Department of Cybernetics at the University of Reading, PO Box 225 Reading, RG6 6AY United Kingdom, K.Dautenhahn@cyber.reading.ac.uk.

is following the teacher and thus imitating its movements. Both agents, then, share the same physical world in which they act and interact. The learner associates its sensor data with the simultaneously arriving communicative signals (words) which the teacher emits. The setup does not require that both agents use the same sensory data or even the same sensors in order to characterise a word. A concrete situation is interpreted on the basis of the learner's own sensory-motor interactions. For our experiments we use two robots which have a different shape and different 'morphology'.

Many robotics studies use neural networks for controlling motor-sensor correlated tasks (see [2] for a brief review of it). In our approach, we use a neural network in order to address the problem of how to attach the meaning to a word, e.g. how can one agent's perception and description of the world make sense for another agent that is physically different. In opposition to related robotics studies ([9], [16], [12]), the learning of the language is not implicit in the mechanism and the learning architecture which we use. A simple 'following' strategy allows implicitly the agents to share the same physical context. Learning to agree on a common vocabulary results from a general capability to form multiple sensory associations, where communication is regarded as simply another sensor modality.

Our approach differs from previous works on agent communication and social behaviour ([10], [14], [13]) since it combines the following aspects. 1) The specific experiments which we describe in this paper study the development and usefulness of communication skills in conjunction with a social relationship, rather than studying how communication can be used to establish a relationship. 2) We study communication between heterogeneous robots, on the individual level, rather than studying communication processes within homogeneous societies of agents (system's level). Imitation is a way to learn about another agent's interpretation of the world, even if the individuals' perception (e.g. due to a different spatial perspective) differ. Imitation is a very complex social learning mechanism in the animal kingdom (see [7] and [5]). Related work on imitative learning with robots can be found in [6] and [8]. There, imitative learning controllers based on movement matching mechanisms are developed and implemented in experiments where, in the first case, a robot imitates head movements of a human demonstrator, and in the second case, learns an assembly task by observing a Human's actions. In this paper we use a simple imitative strategy, namely bidirectional following, not as the learning mechanism but as a means for learning to occur by enabling the agents to share a common situation. We proposed the basic idea to use spatial bonding (following) in order to create a common context and allow social learning originally in [5].

We first report briefly some physical experiments to test the validity of the learning architecture. Following the 'social robot' line of research, we then describe a specific scenario ('mother-child'), whose physical implementation is reported in the second part of this paper. This scenario gives an example of a situation where the ability to communicate can be advantageous for the individual. A physical implementation is carried out in a system of two mobile autonomous robots, teacher and learner.

2 The 'Mother-Child' Scenario

The ability to communicate can be advantageous for an individual animal or human in many situations. An individual autonomous robot, that would be expected to interact in society, could then also benefit from such a skill. The 'mother-child' scenario is a simple example of a situation where the ability to communicate increases the survival chances of the agent. Both terms (Mother-child) are used in a metaphorical sense in order to characterise a situation where a) one robots is the model for another robot's behaviour, b) the teacher robot knows the right words for useful places and situations in the environment, and c) the learner robot is supposed to have an instinct to stay close to the teacher robot, i.e. the learner can benefit from the knowledge of the teacher.

The scenario is composed of two parts. In the first stage, the 'mother' robot teaches the 'child' robot while the 'child' is following and imitating her movements. Teaching concerns locations in the environment. The 'child' associates its sensor data with the simultaneously arriving words emitted by the 'mother'. Later, 'child' and 'mother' are separated. The 'child' robot then uses its knowledge of the 'mother"s vocabulary to find her while she is announcing her location. Thus, this can be regarded as the first step of a 'child' developing an understanding for the language of the 'mother'. We use the term 'language' in this paper in an metaphorical sense, referring to bit-strings ('words') as communicative signals used in the specific contexts which we investigate. We do not intend to make any claims about human language.

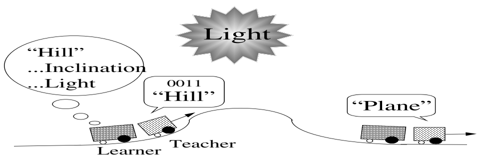

Figure 1: **Visualization of the 'Mother-Child' scenario"**

The scenario was implemented with two mobile autonomous robots, the 'mother' and the 'child', in a hilly area. Firstly, 'mother' and 'child' wander around in the area, following each other. Following and keeping contact is actively maintained by learner and teacher. While moving around, going up and down the hill, the 'mother' robot names the places it travels through. The 'child' learns to associate distinct 'names' with the hilly and flat area. 'Words' or 'names' are simply bit-strings which are transmitted via radio-link, they do not possess a meaning in itself and are not interpreted as such. In a second stage, 'child' and 'mother' are separated. Then, because the 'child' needs crucially its 'mother' (following a predefined instinctive behaviour), it immediately begins to look for her. Unfortunately, because of sensory limitation, it can not see her and thus begins a random search that may be long and may not succeed. Meanwhile, the 'mother' is still naming the locations she goes through. The 'child', hearing and understanding the 'mother''s messages, uses them as supplementary information for finding her. In the 'mother-child' scenario, the ability of understanding the language of the 'mother' is advantageous for the 'child' because it gives supplementary information to find its 'mother' whom it crucially needs (following an instinct).

3 The system architecture

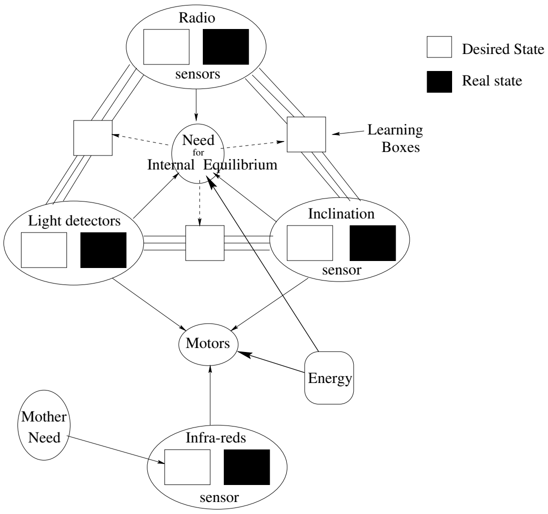

Figure 2: **The behavioural architecture**: A fixed quantity of energy is shared between two possible actions, motor activations and learning. Actions are motivated proportionally to sensory associated motivational factors. Motivational factors are activated by comparing the agents' desired and real sensory states. Desired states are produced by two built-in desires, the *'mother' need* and the *need for internal equilibrium*.

The behavioural architecture is an implementation of DRAMA, the general control D(ynamic) R(ecurrent)

A(ssociative) M(emory) A(rchitecture) we developed previously in [2]. Figure 2 represents schematically the structure of the architecture used in the experiment[1].

The DRAMA architecture proposes a model of control architecture for autonomous learning robots. General mechanisms are defined to combine *life-long* learning with built-in basic behaviours. Learning is dynamical and continuous. It is based upon the ability to correlate sensory stimuli. An associative memory network (learning box), an improved version of the original Willshaw network[15] , is defined between each pair of sensors. In each cycle, the connections of the associative memory are updated following an Hebbian rule in order to represent the actual correlation between the two sensory input state. An important advantage of the Willshaw network and of associative memory networks in general over other kind of connectionist models is that learning is bidirectional. There is no distinction between inputs and outputs as in a feedforward network for instance. Consequently, whenever one input is presented at one end of the network, its correlated output is automatically retrieved at the other end. This bidirectionality allows easily transitive learning and thus chaining reasoning. Adding a few mechanisms to this architecture can result in potentially complex behaviours. The general ideas of the architecture are presented below together with their implementation in this experiment.

Desired versus real sensory states The control architecture of the behaviour described in [2] introduces the notions of *real* and *desired* sensory states. The *desired* sensory state is determined by a built-in behaviour, say an instinct[2]. It represents what the agent wishes to sense but not necessarily what it really senses, which is the *real* sensory state. Comparison between desired and real sensory states is used as a motivation for acting. The agent acts (either moves or learns) in order to satisfy its desires if they do not match reality.

In our implementation, two built-in desires are defined, the *'mother' need* and the *internal equilibrium need*. The need of the 'mother' is a survival instinct, namely the drive to perceive a high frontal infra-red reflection. This principle was successful because the 'mother' robot was the only obstacle in the arena. The need for internal equilibrium provides the motivation for learning (and additionally for moving). Internal equilibrium is achieved when the real sensory stimuli at both ends of the associative network correspond to the predicted, i.e. desired outputs. If this is not the case, a motivation for learning is activated. The motivation intensity is proportional to the Hamming distance between the two sensory vectorial states (desired and real). The confidence factors associated to the nodes of the network are updated proportionally to the motivation factor[3].

Motivation versus Energy potential The motivational factors compete for the energy distribution. We define the energy potential of the agent as a fixed value which is shared between all possible actions of the robot. The actions in our experiment are moving (activation of the two motors) and learning (updating of the structure of the learning architecture). Actions are activated by motivational factors.

In our experiment, the desire for an infra-red detection of the 'mother' produces a motivation for moving. Motors left and right are activated proportionally to the motivation. The robot runs in the direction where the infra-red sensor reading is the closest to the desired value, which corresponds to a high infra-red measurement. As we said previously, the desire for an internal equilibrium creates a motivation for learning. Because the need for equilibrium is linked to the state of the learning system, the motivational factor associated with it should favour the action of learning over that of moving. However, it was considered that small movements in parallel to a tuning of the learning could be advantageous in some cases in order to cope with the experimental noise. Because the *'mother' need* is concerned with the survival of the individual, its motivational influence on the motors' activation is privileged over the one of the *need for equilibrium* in the energy sharing process. The activation of the motors due to the *'mother' need* can consume up to a maximum of 80% of the energy available. The remaining energy (at minimum 20%) is divided up by the actions to satisfy the need for equilibrium, that is in first instance learning and in second instance motors' activation. In the experiment, the sensory associations to be learned referred to two light sensors and one inclination sensor (to describe the environment) and a radio receiver

[1]Note: The control architecture described here was implemented only in the 'child' robot. The behaviour of the 'mother' was defined as a simple behaviour-based architecture, with two behaviours *following* and *talking*.

[2]In future experiments, this could be modified and determined by the architecture itself, as for instance defining them as motivational factors which change dynamically.

[3]For reason of simplification, we updated the nodes with a constant factor in the experiment but we will apply this idea in further experiments.

(to generate the communication signals). The motivational factors associated with the light detectors and inclination were set to influence twice as much the action of learning as moving. However, because the radio reception was almost perfect, a tuning in terms of movements was judged unnecessary. Thus, the motivational factor associated with the radio system influences only the learning and not the motor activation.

This setup produced the behaviour designed for the 'mother-child' scenario. When the 'child' is moved away from the 'mother', the motivational factor produced by the 'mother's need activates the motors at the maximal excitation (up to 80% of the energy available). In parallel, the 'mother' sends radio encodings corresponding to her position in the arena (on the hill or on the plane) that does not correspond to the position of the 'child'. The need of equilibrium between the 'child's actual sensory measurement of light and inclination and the desired one stimulated by the 'mother''s radio messages creates a motivation for both learning and moving. The energy consumed by the motors jumps to about 90% of the energy available, so that about 10% of energy is left for learning. This is, however, not sufficient to activate an update of the associative network, the activation of the learning function being thresholded to occur with a minimum of 15% of the energy. In consequence, the 'child' does not update its knowledge about the communication vocabulary when hearing the 'mother''s messages, which would result in a relearning of the message. Instead, it uses them to search for the 'mother'.

Software We programmed the robots in PDL language (developed at the Vrije Universiteit Brussels, [11]). It is based on a parallel processing principle. Each process acts on global quantities that are updated only at the end of each cycle. The final value of the quantity is the result of update of each process. Consequently, the observable behaviour of the robot results from the combination of all processes. The *following* behaviour of one robot by the other one was programmed as a combination of ten processes acting on the motors quantities in function of the bumpers and the infra-red sensor measurements. The *following* was bidirectional. Each robot was provided with a set of three infra-red sensors, on the back of the teacher and on the front of the learner that enabled them to 'see' each other. The motors were given values proportional to the sensor measurement of the infra-red detectors. The closer the following robot, the slower it moves and inversely for the followed robot. The difference between the side detectors and the middle one was used to control the two independent driving wheels. A set of six bumpers (frontal for the learner and in the back for the teacher) was used for close alignment. This resulted in very efficient alignment and following behaviour of one robot behind the other.

4 Experiments

Firstly, we describe experiments carried out to test the learning capacity of the architecture. Learning consists of multiple sensory association between radio, inclination and lights sensors. The teacher robot teaches the learner robot while the learner is following it. The learner associates its observations with the teacher's words. Teaching concerns locations in the environment. The stage of the experiment consists of a rectangular arena (1.8mx4.9m), in the middle of which a hill (1mx1m) is located. Walls (30cm heigh) surround the arena. A light source is positioned 42cm above the top of the hill. Thus, when the robots are climbing up the hill, they sense an increase of light intensity. Figure 3 shows a schematic view of the setup of the experiments with the hill and the two robots.

The learner learns to distinguish between going up and down the hill and moving on the plane by associating these actions with different words. In parallel, it learns to associate an increase of light intensity with going up the hill.

The robots' bodies are made of *fishertechnik* components. The robots are equipped with different sensors, a radio transceiver, infra-red and light sensors, inclination sensors (one is used in the experiments to measure the front-back body axis). The robots have two motors and one passive wheel. The robots have different oval shapes, surrounded by a ring of 16 contact sensors. The teacher is 30cm wide and 32cm long, the learner is 33cm long and 27cm wide. They have a height of 23 cm and a weight of 2.2-2.5 kg. Figure 4 shows two photos of learner and teacher robots crossing the hill. We use radio transmission as the communication channel. The learner robot associates radio words which are encoding three situations: going up and down the hill and moving on the plane. It associates on the one hand the radio words with changes in inclination and light intensity, and on the other hand changes in light intensity with variations of inclination.

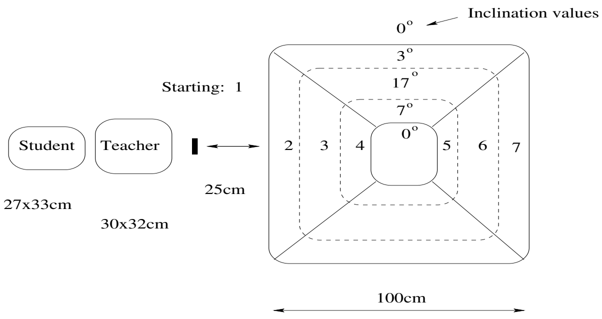

Figure 3: **Setup of the experiments**: A learner robot follows a teacher robot, going up and down a hill. The robots' sizes and shapes are different. The hill stands in the middle of the area. The four faces of the hill are symmetrical. The top of the hill is flat and the inclination is gradual. The steepest inclination is equal to about 17 degrees centigrade. 7 control positions are defined for testing the learning results.

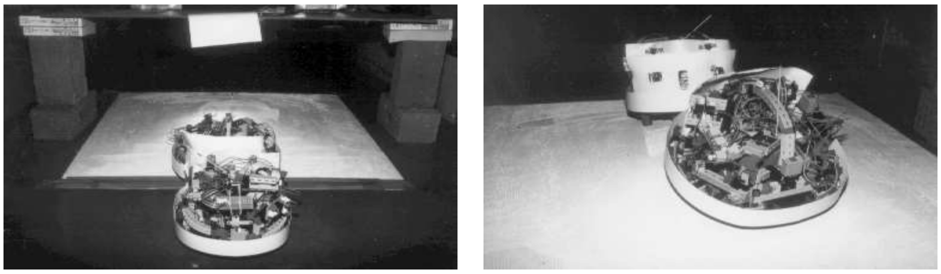

Figure 4: A: The two robots beginning to climb the hill, learner following the teacher. The light source is located behind the white sheet of paper. B: The two robots on top of the hill.

The experiment starts with the two robots aligned one behind the other one, facing the hill, within a distance of 25cm from it. They climb up and down the hill and are stopped on the other side of the hill after having reached the symmetrical position, 25cm behind the hill. They are then turned around to face the hill from the other side and climb again in the opposite direction. During one test, they climb the hill back and forth ten times. The choice of 25cm for the starting positions on both sides of the hill gives an equivalent proportion to each area of teaching.: about 50cm for each inclination zone (plane, upwards inclination and downwards inclination). Seven control positions were chosen in the area to test the result of the learning. At the end of a test, the learner robot is placed immobile on each of the seven positions and its responses (words) are measured. We could in this way document which word has been associated with which position, each position corresponding to a different inclination. The first position is on the flat part of the area, 25cm away from the hill. The six others positions are distributed symmetrically on both sides of the hill. Average values of inclination for each position are given in figure 3.

5 Results and Discussion

Imitation, as the reproduction of another agent's action, is a powerful but limited way of getting insight into the other's world. In our experiment, imitative behaviour is used as a means to develop a vocabulary about an agent's perceptions and environmental states. By following the teacher robot, the learner robot imitates in a simple way the movements of the second robot. Movements occur therefore more or less simultaneously and are quite similar. In addition, if both robots follow each other closely enough, their sensory perceptions of the situation are also similar. But, movement imitation introduces a temporal delay in the world perceptions of the two robots. Moreover, its a basic characteristic of embodied agents that they occupy distinct places in space, causing a spatial displacement ([4]). Learning by imitation is therefore intrinsically imperfect. Two robots can never share the same position in the area, they always have a slightly different ‘perspective’ and their perceptions can never be identical. The teacher’s perception cannot correspond perfectly to the learner’s one, therefore its teachings describes a world that is different from the learner’s one. Consequently, learning is partially erroneous.

If we assume that both robots follow each other perfectly, then the world description given by the teacher becomes the learner’s one as soon as the learner has travelled a distance equal to the teacher’s body length. One could then think that delaying artificially the teacher’s teachings (transmission of signals) could easily solve the problem. This would be a very naive conclusion. Instead, it cannot be guaranteed that the learner will ever have the same sensory perceptions as the teacher: while following each other, the two robots are constantly re-adjusting their alignment one behind the other other. These small sideway movements result in a different orientation of the two robots and so in a different measurement of inclination and light intensity.

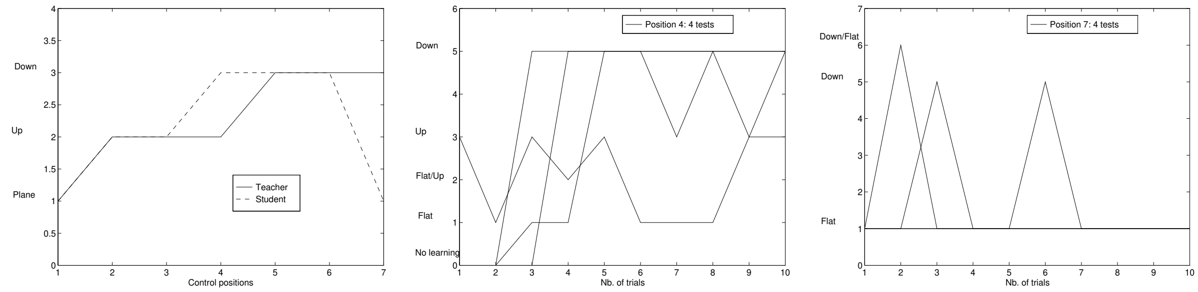

Figure 5: **Left**: The learner robot learns to associate three radio words with different values of inclination (seven control positions stand for seven different inclination positions). It distinguishes between zones of upwards (‘up’), downwards (‘down’) and plane (‘flat’) inclination. However, because teacher and learner’s perceptions are not simultaneous, teaching is imperfect. Consequently, learner’s and teacher’s sensory zones attached to the same radio concept are shifted. The critical positions are positions 4 and 7 (top and bottom of the hill respectively, cf. figure 3). The graph shows the learning results in average over 5 tests, where one test consists of ten trials (1 trial is one way climbing the hill). **Middle, right**: Middle and right figures show the evolution of the learner’s concept learning for the two critical positions 4 (top) and 7 (bottom of the hill) along the trials over 5 tests. Learning first fluctuates between the different taught concepts (‘down’, ‘flat’ and sometimes ‘up’ for position 4 and ‘down’ and ‘flat’ for position 7). It finally converges after about 5 to 10 trials to the most frequently taught concepts (‘down’ for position 4 and ‘flat’ for position 7), different from what the learner really senses (‘up’ for position 4 and ‘down’ for position 7).

The learning architecture is robust against noisy data (where the noise is the number of bad teachings against the number of good ones). But, unfortunately, the imperfections of the teaching due to this time delay cannot always be balanced by a large number of trials. Some concepts simply cannot be correctly learned in this experiment. The critical positions in our experiment are the top and bottom of the hill. In the situation where both robots are climbing up the hill, the teacher usually reaches the top of the hill while the learner is still climbing the hill. The teacher is then either standing flat at the top or beginning to go down on the other side, therefore it teaches most often either the word ’flat’ or the word ’down’ and very seldomly the word ’up’ that is what the learner really senses. Similarly, when they are both going down the hill, the teacher reaches the flat surface at the bottom of the hill while the learner is still climbing down the hill. Its teaching is then ’flat’ which never corresponds to the learner’s sensor data (down).

Consequently, learner’s and teacher’s sensory zones attached to the same radio concept are shifted one from the other one as shown experimentally (see figure 5 left). The learner associates the concept ’down’, already while it is still reaching the top of the hill (control position 4) and the concept ’flat’ when it is reaching the bottom of the hill (control position 7). Figures 5 middle and right show the evolution of the learning for these two critical positions. Learning first fluctuates between the different taught concepts (‘down’, ‘flat’ and sometimes ‘up’ for position 4 and ‘down’ and ‘flat’ for position 7). It finally converges after about 5 to 10 trials to the most frequently taught concepts (‘down’ for position 4 and ‘flat’ for position 7), different from what the learner really senses (‘up’ for position 4 and ‘down’ for position 7).

Simultaneously, the learner also learns to associate the increase of the environmental light intensity

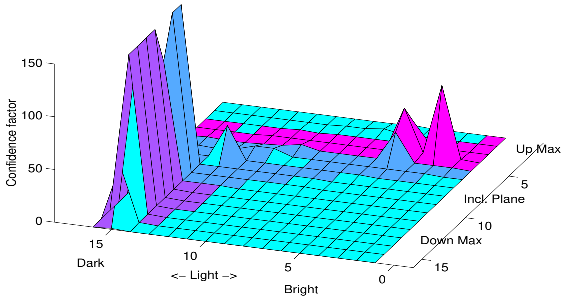

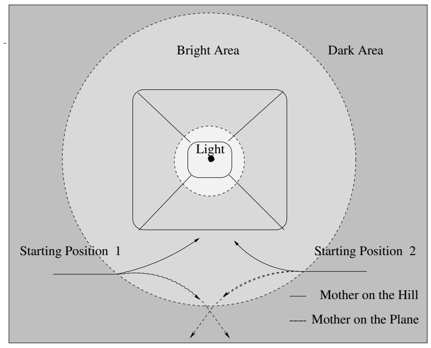

Figure 6: **Left picture** shows the result of triangular association of light intensity, inclination measurements and radio words: A progressive increase of the light intensity is associated with the increase of inclination in upwards direction. Downwards and flat inclination data are associated with low values of light intensity. Graphics legend: the Z axis represents the value of confidence factor associated with the matrix nodes between light and inclination measurement, represented accordingly on the X and Y axis. Colours legend: There are four distinct zones representing the results of association with radio words. The largest and brightest zone corresponds to a zone where no association has ever been made between radio concepts and light and inclination measurements. Then, association with the three radio concepts (flat - down - up) is represented respectively by order of increasing darkness. The radio association corresponds to figure 5. **Right picture** shows a schematic view of the 'Mother-Child' experiment. A light source is positioned on the top of the hill. Firstly, the 'mother' teaches the 'child' to distinguish between the hilly and plane zones corresponding to bright and dark zones. Later, the 'child' tries to find the 'mother', that is saying her location, by directing itself in relation to the light.

with the increase of the inclination, a bright light being suspended at the top of the hill (see figure 3). Results of the learning are shown in figure 6 left. We observe that the brightest values of lighting are associated with upwards inclination, while low values of light intensity are associated with downwards and flat inclinations. This corresponds to our expectations since the learner receives the maximal light intensity when it is facing the light, going upwards towards it. When it is in a flat position, either on the plane or at the top of the hill, it does not receive as much light because it does not face it, but it still receives more than when it goes down the hill, turning its back to the light. The breadth of the lighting zone associated with upwards inclination accounts for the large variation of the robot's lighting measurements when climbing up the hill. This variation is due to the sideways movements of both teacher and learner to re-align themselves one behind the other. On the one hand, the learner, by its own movements, changes constantly its orientation towards the light source affecting its measurement of the lighting. On the other hand, the shadow produced by the teacher's body and projected on the learner varies depending on the teacher's movements.

5.1 'Mother-Child' experiment

The ability of learning multiple sensory association, as demonstrated previously, was used for implementing the 'mother-child' scenario, described in section 2. In the first stage, the teacher, the 'mother', wanders around the area with the learner, the 'child', following 'her' and teaches it to distinguish between climbing the hill and being on the plane. The 'child' associates the concept of hill with both a bright light and an upwards inclination. In the second stage, 'child' and 'mother' are separated and the 'child' tries to find the 'mother' when hearing 'her' saying her position on the hill or on the plane by directing itself in relation to the light.

The setup of the experiment is similar to that used for the previous experiment. Figure 6 right shows a schematic view of the experiment with the hill and the light source on top of it. The light emitted by the light source is directed downwards to the hill, forming a cone of light around the hill, brighter at the top and then decreasing progressively. The dark area corresponds to the part of the area outside the light cone. In this area, the robot does not receive direct light but only diffuse reflections from the walls.

The light intensity received by the robot's two detectors is then approximatively the same. This light intensity threshold is important because it determines the reactivity of the robot. The difference between desired and real states for each light detector determines the speed given to the motors. The difference of sensitivity between right and left detector determines consequently the direction of movement of the robot. Therefore, as long as the robot stays in the dark area, it can only move forward, because the difference in light intensity measured by its two detectors is too small to cause turning. However, once it has crossed the boundary between the two areas, it will turn either towards or away from the light source depending on the 'mother"s messages.

In the first stage, 'mother' and 'child' climb the hill up and down, from opposite sides, until radio concepts 'hill' and 'plane' and sensory stimuli (inclination and light) are well associated. An LED display on the learner robot allows the experimenters to observe the progress of the learning process. 'Child' and 'mother' are then separated. The 'mother' is put on the hill far away from the 'child's infra-red detection. The 'child' is put in starting position 1 (cf. figure 6) and its behaviour observed. The scenario is then repeated by placing the 'child' in the symmetrical position 2 and so on with five different positions on each side of the hill. The situation is then reversed. The 'mother' is put on the plane and the 'child' placed again in one of the starting positions and so on. This allows us to test if there is a bias from any direction.

10 tests were done and recorded with a video camera. Tests showed that the experiment was successful. The robot's reactions were correct, but very sensitive to the difference between dark and bright areas. Outside the bright area, the robot keeps strictly moving forward, as explained previously. When entering the bright area only, the robot begins well to turn to the right direction, that is either towards or away from the hill, depending on whether the 'mother' is on the hill or on the plane. Once the region 'hill' is reached, the desire for light is satisfied. The robot slows down, however it does not stop until the desire for inclination and infra-reds are also satisfied. Thus, if it does not find the 'mother' on the hill, it moves slowly around the hill, satisfying partially its desires for both light, upwards inclination and infra-red perception.

6 Conclusion and Outlook

This work is a step towards the construction of autonomous social robots. The social interactions of communication and imitation were studied in a physical experiment with two autonomous robots. A learner robot learned to communicate about 'meaningful' situations in an environment. A 'mother-child' scenario was designed and experimented in order to demonstrate the need for communication in an artificial ecological system.

Imitative behaviour was used to establish a social relationship between two agents. Imitation is a way to get insight into the other's world perception and, as such, was demonstrated to be useful for communication. However, it is an imperfect means of learning. Because they imitate each other, the two agents' perceptions become similar but not identical, thus leading sometimes to erroneous learning. Probably more complex cognitive mechanisms could lead to a better interpretation and use of imitation in order to avoid this intrinsic imperfection. An excellent example for very successful imitation are human imitation games. For instance, the ability of learning sequences could be a way to overcome the problem of time-delay: the learner robot will associate the concept 'top of the hill' correctly by associating it with the next closest *event* (event defined as a change in inclination) and not with the simultaneous sensory measurement.

The DRAMA architecture used for controlling the behaviour of the robot is a general framework of control architecture for social robot learning. Its implementation in the 'mother-child' scenario showed that it allows complex behaviours, involving learning and social skills, without requiring complex computation. Further research will be directed towards more complex experiments, tackling the issue of sequential learning and studying multiple agent interactions in a more complex environment. Computer simulation would accompany physical experiments.

7 Acknowledgement

Thanks to Luc Steels and Gillian Hayes for supporting the experiments.

References

[1] T. Arita and C. E. Taylor, (1996), ‘A simple model for the evolution of communication’, UCLA Commotion Lab Tech Report 013, UCLA, US.

[2] A. Billard and G. Hayes (1997), ’Robot’s first steps, robot’s first words ...’, Proceedings of Groningen Assembly on Language Acquisition, Edinburgh, U.K., April 97.

[3] K. Dautenhahn, (1996), ‘Embodiment in animals and artifacts’, Working Notes AAAI 96 Symposium on Embodied Action and Cognition.

[4] K. Dautenhahn and T. Christaller, (1996), ‘Remembering, Rehearsal and Empathy - Towards a Social and Embodied Cognitive Psychology for Artifacts’, Two sciences of the mind. Readings in cognitive science and consciousness, John Benjamins North America Inc. publisher, Sean O’Nuallain and Paul Mc Kevitt editors, pages 257-282.

[5] K. Dautenhahn, (1995), ‘Getting to know each other – artificial social intelligence for autonomous robots’, Robotics and Autonomous Systems, vol. 16, pages 333-356.

[6] J. Demiris and G. Hayes, ‘Imitative Learning Mechanisms in Robots and Humans’, Proc. of the 5th European Workshop on Learning Robots, Bari, Italy, July 1996.

[7] Galef, B. G. Jr. (1988), ‘Imitation in animals: history, definition and interpretation of data from the psychological laboratory’, in Social Learning: Psychological and Biological Perspectives, T.R. Zentall and B.G. Galef, Jr. (Cds.), Lawrence Erlbaum Associates.

[8] Kuniyoshi Y., Inaba M. and Inoue H. (1994), ‘Learning by Watching: Extracting Reusable Task Knowledge from Visual Observation of Human Performance’, in IEEE Transactions on Robotics and Automation’, Vol. 10, No 6, Dec. 1994.

[9] L. Steels & P. Vogt, (1997) ‘Grounding adaptive language games in robotic agents’, Proceedings of the fourth european conference on Artificial Life, Brighton, U.K., July 97.

[10] L. Steels, (1996) ‘The spontaneous self-organisation of an adaptive language’, S. Muggleton (ed.): Machine Intelligence 15, Oxford University Press.

[11] L. Steels (March 1992), ‘The PDL reference manual’, Memo published at the Artificial Intelligence Laboratory, Vrije Universiteit Brussel.

[12] M. L. Mataric, (1994) ‘Learning to behave socially’, Proc. of the Third International Conference on Simulation of Adaptive Behavior

[13] B. MacLennan,(1991) ‘Synthetic ethology: an approach to the study of communication’, Proc. of the Second Artificial Life Workshop, C. G. Langton, C. Taylor, J. D. Farmer and S. Rasmussen editors, pages 631-658.

[14] G. M. Werner and M. G. Dyer, (1992) ‘Evolution of communication in artificial organisms’, Proc. of the Second Artificial Life Workshop, C. G. Langton, C. Taylor, J. D. Farmer and S. Rasmussen editors, pages 659-687.

[15] D. Willshaw and C. von der Malsburg, (1976), ’How patterned neural connections can be set up by self-organisation’, Proceedings of the Royal Society of London B, 194:431-445.

[16] Yanco, H., (July 1993), ’Talking about the world: Cooperative Robots that learn to communicate’, In Proceedings of the Workshop on Learning Action Models, held at the Eleventh National Conference on Artificial Intelligence, Washington, D.C.