Robot's first steps, Robot's first words ...*

Aude Billard & Gillian Hayes

Department of Artificial Intelligence
University of Edinburgh, U.K.
audeb@dai.ed.ac.uk

Abstract

Communication is a desirable skill for robots. We describe a method of how these skills could be learned. An architecture of connectionist models is developed to allow the learning of a simple non-grammatical language. Concepts, in the form of a word and its meaning, and numerical quantities are understood separately. A physical system of two mobile autonomous robots, teacher and student, is built to implement the architecture. Concepts relative to proprio perception, such as the agent's movements, and extero perceptions as objects in the environment, are taught. Extension of the architecture to cover the understanding of quantities is tested by teaching the notion of distance. Teaching is done via a teacher agent to a student. Movement imitation is the means by which the language is transmitted. While following and, thus, imitating the teacher's movements, the student agent learns to associate the teacher's words with its own sensor perceptions.

1. Introduction

Most living societies communicate, from insects to humans. In simple species, communication abilities are often built in. However, in more complex creatures this capacity develops during their life. Studying systems that would develop such an ability is a necessary and natural step towards the construction of robots showing more complex capacities.

Communication is a very desirable skill for robots to possess. We expect communicating robots to be able to carry out a wider range of tasks than is possible for dumb robots. For example, tasks requiring coordination within a group of robots would be easier if the robots were able to inform one another about what they were doing. Whilst fixed communication protocols can be designed, we expect that a flexible language which is acquired through a learning process will provide more adaptability towards the tasks the robots can carry out, the types of robots engaged in those tasks, and the human individual programming the robots. It will provide us with a way to teach robots after they have already been built, as well as (eventually) a way of understanding their internal function, by having them explain it. Consequently, interest in language acquisition for and by robots is growing.

A learning architecture is developed to enable learning of a basic non-grammatical language. Learning is based upon a general ability to correlate proprioceptive and exteroreceptive sensory information. Concepts, in the form of a word and its meaning, and numerical quantities are acquired separately. The learning of the language includes understanding of the meaning of a word and memorisation of its elements, which, in turn, allows its use. The architecture is composed of two connectionist models, an associative memory and a recurrent network. The associative memory, a Willshaw network augmented with a confidence factor notion, allows concept understanding by the association of words received as radio signals with sensory observations of the external environment surrounding the agent and its own proprioceptors. A recurrent network provides the understanding of quantity, by associating radio sequence length with particular sensory measurements.

The architecture is implemented in a physical system of two non identical mobile autonomous robots, teacher and student. By matching its movements with those of the teacher the student is led through a series of situations in which it is to learn. When the teacher detects the presence of the student, it starts to teach concepts related to its movements (such as "stopped", "moving", "turning left"), and related to objects it can perceive in the environment (such as "recharging station"). The architecture also allows the learning of quantities (such as distance moved in terms of the number of wheel rotations of the student/teacher). Once the confidence factors of the student's associative net are sufficiently high, i.e. it is confident about what it has learnt, it starts to transmit its own radio signals associated with what it is currently sensing. This is allowed by the bi-directional nature of the Willshaw net.

Four sets of experiments are conducted, where concepts relative to the internal state of the agents and to the agents' environment are taught. A very basic language is so developed, in which distinctions between the notion of concepts and the quantity associated with the concepts was made possible. The agent learns to differentiate between actions like *being stopped* compared to *moving* and to associate a word with recognition of an object (flashing lamp) in its environment. Conceptual and quantity notions are distinguished in

*In the proceeding of GALA97 conference, Groningen, Assembly on Language Acquisition, Edinburgh, April 4-6 1997.

learning the notion of *distance*. It is important to note that since the two robots have differing designs, the radio signals transmitted by the teacher and received by the student are associated with the particular sensory situation of the student. Thus, the student develops the language in terms of what it “sees”, “hears” and “feels”. It is not necessary for the teacher’s morphology to be identical with that of the student, merely for them to be able to detect one another, have similar movement capabilities, and a learning architecture. In addition to the physical experimentation, theoretical studies and simulation tests showed the robustness of the learning architecture towards noisy sensory signals. While developing the learning architecture, emphasis was put on simplicity and generality so that it would not be specific to the robot’s type nor to the environment used for the implementation.

In the following , we describe the learning architecture (section 2) and its implementation in a physical experiment (section 3).

2. The architecture

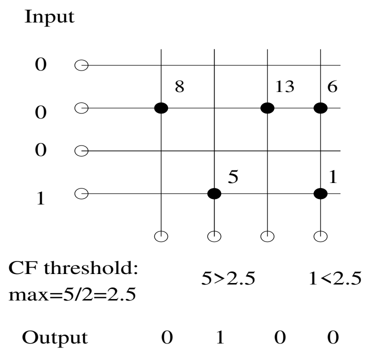

Figure 1: **The training phase (example)**: The patterns are presented as arrays of binary (0/1) inputs. The learning stage begins with all confidence factors equal to zero. The confidence factors associated with each intersection nodes are incremented or decremented by one unit, whenever the two lines are activated or only one is activated and the other zero, respectively. Note: The numbers associated with the confidence factors were chosen randomly for the example (for the testing phase).

Learning is based upon the ability to correlate different sensory inputs. The architecture is composed of two connectionist models, an associative memory and a recurrent network. The associative memory, a **Willshaw network** (Willshaw & von der Malsburg 1976) augmented with a confidence factor notion, allows concept understanding by the association of radio words with sensory observations. 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. During training, sensory outputs are presented at both ends of the

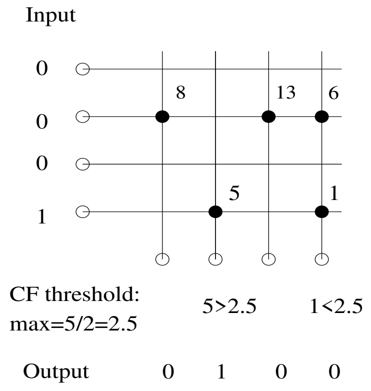

Figure 2: **The testing phase**: The recall of a memorised pattern depends on the confidence factor value for each output unit. An input pattern is presented to the net. For each output column, only the activated matrix nodes corresponding to an active input unit line are considered and their confidence factor added for the same output unit. Only output units whose confidence factor is greater than half the maximal confidence factor for all outputs will be activated.

network. The internal structure of the network is then updated, reflecting correlations between both input sources. Consequently, whenever one input is presented at one end of the network, its correlated output is automatically retrieved at the other end.

In the experiments which are reported here, the architecture is implemented to allow the learning of a language by associating radio ‘words’ with sensory observations. The characteristic of our associative memory allows not only the understanding of the radio ‘words” meaning but also their memorisation and finally the ability to use them in turn. Figures 1 and 2 show a simple example of training and testing phases of the associative network.

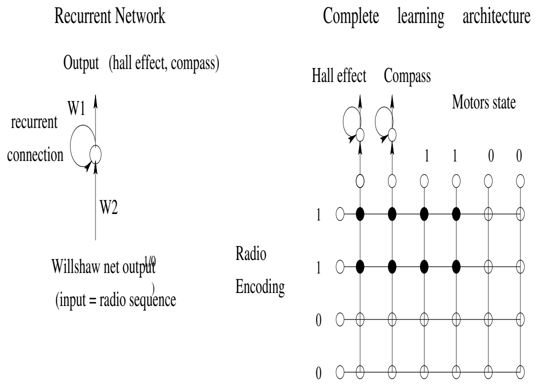

Figure 3: The learning architecture is made of a combination of a Hebbian associative memory and one-node recurrent networks on the output nodes.

The notion of quantity is encoded by repeating the same radio word that encodes for the particular concept, a number of times proportional to the quantity carried by the concept (for example, 20 radio words correspond to 2 wheel rotations). Simply adding a one unit recurrent network at the end of each output unit (see figure 3) allows us to solve this simple learning task. The network consists of only two connections, a recurrent connection on the unit and a connection from the output unit of the associative memory to the recurrent unit. Learning is achieved once the correct weights for the connections are found (weights are updated following a classical backpropagation algorithm), i.e. the sequence formed by the repetition of the same radio concept is associated with a particular value for a sensor measurement. Because of the simplicity of the learning task, it is easy to guess what the values of the connections should be. The recurrent connection corresponds to the addition function, its weight should then be equal to one. The second connection corresponds to the multiplication function. Its weight should be equal to the proportionality factor between the radio sequence and the sensor output value.

Theoretical studies and simulation tests have been conducted to show the viability of the architecture, which showed its robustness towards up to 30 percent of noisy data (see (Billard 1996)). Noisy data consist in our case of the incorrect teachings, that is the situations where the teacher's teaching do not correspond to the student actual observation. The model was implemented in experiments with physical mobile LEGO robots, which we describe in the following section.

3. The experiments

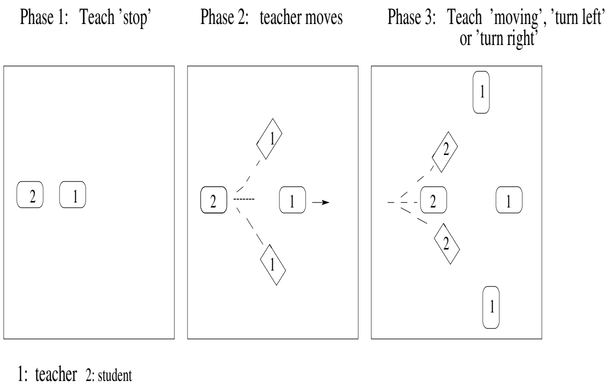

Figure 4: Movements of the robots during the teaching phase.

Four sets of experiments are conducted, where concepts relative to the internal state of the agents and to the agents' environment are taught. Two experiments are carried out to teach movement related concepts. The agent learns to differentiate between actions of *being stopped* compared to *moving* and *being stopped* compared to *turning to the left* or *turning to the right*. Teaching is done via a teacher agent to a student. Movement imitation is the means by which the language is transmitted. Teacher and student follow each other around a given environment. While following the teacher, the student implicitly imitates the teacher's movements. When the teacher detects the presence of the student, it starts to teach concepts related to its movement (such as "stopped", "moving", "turning left"). The student then associates the teacher's words with its simultaneous perception of its own actions. Figure 4 shows a schematic description of the teaching phase. In (Billard 1996)), we report other experiments where the agent learns to recognise an object (such as a flashing lamp) and learns the notion of distance.

A set up of two mobile LEGO robots was used to carry out the experiments (see picture 5). The robots are provided with bumpers and infra-red sensors for obstacle avoidance. They also have light sensors for detecting each other, each carrying a bright light. A behaviour-based architecture is used to direct the robots' behaviours. A set of simple behaviours such as *recharging* and *avoiding obstacles* are designed to provide the agent with basic self-sufficiency. More complex behaviours, such as *curiosity*, *learning* and *speaking* are added to give the agent the ability to communicate. The bidirectional following of the two agents is enabled by using a simple phototaxis behaviour (The two robots carrying a bright light). Figure 6 shows a schematic description of the system.

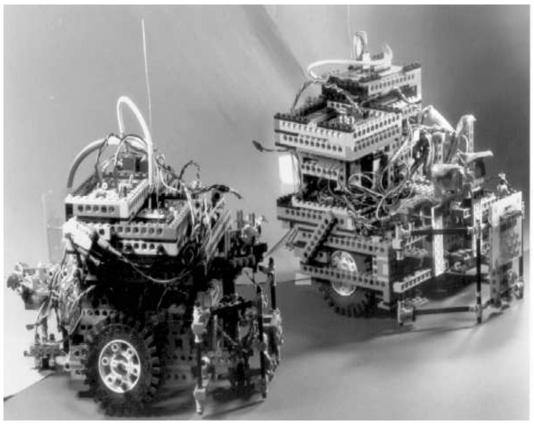

Figure 5: From left to right: the teacher (Ulysses) and the student (Kazam).

The experiments showed that the student learned successfully to distinguish between the concepts *being stopped* and *moving* and to learn a word for describing the flashing lamp and for the notion of distance. Understanding of the concepts was demonstrated by the student using the words correctly. However, experiments to learn concepts associated to directional

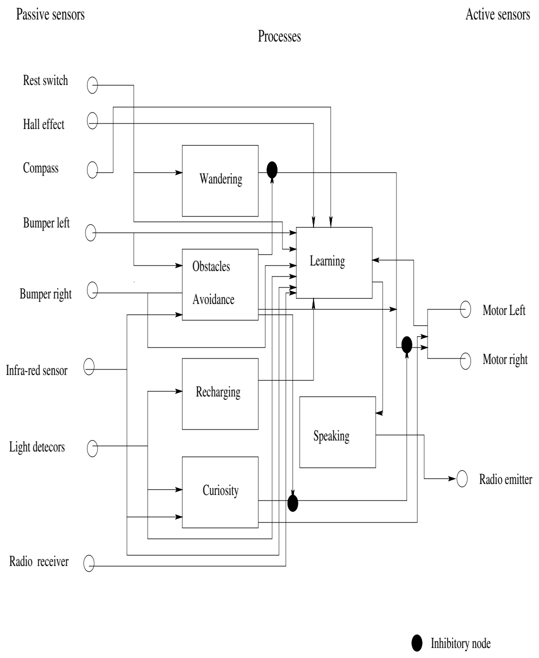

Figure 6: A behaviour-based architecture is used to direct the robots' behaviours.

movements such as *turning to the left* and *turning to the right* as well as experiments to learn the notion of numerical quantity contained in the distance notion have not been successful. Failure was due to the difficulty of carrying out a statistically reasonable number of tests and also to the difficulty in keeping the sources of disturbance within an acceptable margin without constraining the experiment too strongly. For a more detailed description of the experiments and a fuller discussion of the results, the reader should refer to (Billard 1996)). Work is continuing in order to improve the complexity of the system (Billard & Hayes 1997)).

4. Conclusion

This work reported physical experiments with autonomous robots that learn to develop a knowledge of a basic non-grammatical language. Teaching is done from a teacher robot to a student robot by means of a movement imitation mechanism. The student agent learns a vocabulary for describing proprio and extero perceptions. Learning is based upon associations between multiple sensor information. A meaning is attached to a words, i.e. a radio signal, in terms of the agent's sensor perception Understanding of a notion of quantity, independently of a particular concept, is made possible, by encoding the quantity value in the repetition of the word attached with the concept. Learning consists of finding the correct relationship between the repetition duration and the actual sensor measurement. Successful experiments contributed to show the robustness of the architecture towards noise and thus demonstrated its usefulness for physical implementations.

5. Acknowledgements

Thanks to the technicians for their essential support. The research reported here was supported by a grant from the LEMO company and the Swiss Research National Fund.

References

Billard, A. (1996), "allo kazam, do you follow me?' or learning to speak through imitation for social robots', *MSc thesis, Technical Report 43, Dept. of AI, University of Edinburgh*.

Billard, A. & Hayes, G. (1997), 'Learning to communicate through imitation in autonomous robots', *Proceedings of 7th International Conference on Artificial Neural Networks, Lausanne (CH)*.

Willshaw, D. & von der Malsburg, C. (1976), 'How patterned neural connections can be set up by self-organization.', *Proceedings of the Royal Society of London B, 194:431-445*.