Imitative Learning Mechanisms in Robots and Humans

John Demiris and Gillian Hayes

*johnde@aifh.ed.ac.uk, gmh@aifh.ed.ac.uk

**Intelligent Robotics Group, Department of Artificial Intelligence, University of Edinburgh, UK*

Abstract. We do not exist alone. Humans and most other animal species live in societies where the behaviour of an individual influences and is influenced by other members of the society. Within societies, an individual learns not only through classical conditioning and reinforcement, but to a large extent through observation and imitation. This paper presents an analysis of the problem of adding such imitative learning abilities to mobile robots and describes a biologically-inspired architecture we are developing for imitative learning. Our robotic testbed and learning experiments are described and discussed.

1 Introduction

Interest in the field of robot learning has been growing steadily in the last few years. Adding learning abilities to robots offers certain distinct benefits, such as

- Increasing their ability to cope with a dynamic environment where preprogrammed knowledge of the world can become obsolete or is not available at all in the first place.
- Reducing the cost of programming robots to perform specific tasks.
- Increasing their ability to cope with changes in their own specifications, such as sensor drift.

but also offers some important theoretical benefits by "*forcing us to deal with the issue of integration of multiple component technologies, such as sensing, planning, action, and learning*" (Connell and Mahadevan, 1993).

Roboticists have spent a significant amount of time attempting to add learning abilities to robots. Till recently most of the techniques relied on individual learning, where an agent learns by acting on its environment and receiving reinforcement (notable exceptions include (Kuniyoshi, Inaba, and Inoue, 1994), (Mataric, 1994) and (Munch et al, 1994)). In this paper we describe a novel approach to robot learning which complements existing robot learning approaches. Like humans and other animals, robots could learn a significant amount of knowledge by observing and imitating other agents (humans or robots). Learning by imitation has certain desirable characteristics:

- First of all, it speeds up the learning process. A robot could in principle, given sufficient time and energy, learn any task through reinforcement learning (Sutton, 1990), but the presence of an "expert" could be utilised so that existing knowledge from that expert could be passed to the robot. The "expert" could demonstrate how the solution to a task can be achieved, and the learner could learn by observing and/or imitating.
- Learning, in this sense, does not require the expert to spend time teaching the learner robot to perform the task. The expert could go on doing its tasks as usual and the learner could observe and imitate without interrupting the expert.
- No explicit communication is needed between the robots, so this type of learning can still be used in situations where communication can be

costly or dangerous or even impossible (Huber and Durfee, 1993).

In the next section we will review work done on imitation in natural systems including animals and humans. We will then proceed to describe the architecture we are developing in order to implement such mechanisms on mobile robots.

2 Imitation in natural systems

Imitation involves complex and little-understood cognitive operations (Heyes, 1993). Research on imitation in non-human animals has taken place along the following two dimensions:

- What is imitation, and how can it be distinguished from the several other forms of social learning?
- Which animals are capable of imitation and where are these animals with respect to the mental *scala naturae*?

There is wide disagreement in both dimensions above. Imitation has been used as a label for a variety of social phenomena, including ones which could be explained by other simpler processes, such as observational conditioning and instrumental learning, among others. Experiments with rats, pigeons, parrots, octopuses, dolphins, monkeys and chimpanzees have all claimed to have shown the capacity of these animals to imitate. Reviews of the various definitions and of the work in ethological and psychological aspects of social learning include (Whiten and Ham, 1993) and (Galef, 1988).

The situation is rather different with human research. Humans are taken by default to be able to imitate and research is more concerned with **when, why,** and **how** they imitate. One of the most prominent theories on the development of imitative abilities is Piaget's, who advocated that the imitative ability develops in stages. According to his theory (Piaget, 1951), infants begin their life being able to imitate behaviours where both the model's and the infant's action could be compared within the same sensory modality (for example, vocal imitation, or imitation of hand movements when their own hand is also visible). After one year, the infant is able to detect cross-modal equivalences, and imitate acts that require such cross-modal matching (for example, facial imitation: the infant is not able to see itself in order to compare the two acts visually; it relies on proprioceptive information in order to do that). Finally, at about 18-24 months of age, deferred imitation becomes possible, i.e. imitation of acts no longer present at the infant's perceptual field.

Recent experiments by developmental psychologists have challenged this account. Work by Meltzoff and his colleagues[1] has shown that infants are able to imitate a variety of facial and head movements (Meltzoff and Moore, 1989) very soon (minutes) after birth. Based on these experimental findings, Meltzoff postulated the existence of a mechanism, **Active Intermodal Matching**, which he argues underlies early infant imitation, and is present at birth. When infants see an act they use proprioceptive information from their own body movements to iteratively match the visual information they perceive. The ability of perceiving equivalences across different sensory modalities has also been demonstrated in other sensory cases including audio-visual and visual-tactile matching among others (Meltzoff, 1990).

Whether such a mechanism exists and what its exact nature and function in infants is, is still an issue of intense debate among developmental psychologists. In the next few sections we will analyse our attempts to design and implement such a mechanism in mobile robots, with two targets in mind:

- To equip mobile robots with the ability to map visually perceived actions with equivalent actions of their own. This will open the possibilities of learning through imitation with all the advantages described earlier.
- Make explicit the requirements for such abilities to be present.

3 Imitation in artificial systems

We view imitation in natural and artificial systems as a process that involves three different issues:

- The matching mechanism: How does an agent match actions perceived with equivalent actions of its own?

[1] also replicated by several other researchers in a variety of experimental environments and conditions - see (Meltzoff and Moore, 1992) for a complete list

- How is the mechanism developed? What is its starting point, why does it change later, and how does it change?
- How does the agent use this mechanism to refine/improve its behaviour/knowledge over time?

In this paper we will focus on the architecture we are developing in order to deal with the first question: how can an agent match actions perceived with equivalent actions of its own. We will then proceed to how can such mechanism be used to learn new knowledge and skills.

3.1 The matching mechanism

We are designing and implementing a movement matching mechanism inspired by Meltzoff's recent experiments in early infant development. There are five modules in our architecture (figure 1): visual preprocessing, proprioceptive analysis, establishment of relationships, movement analysis, and movement matching.

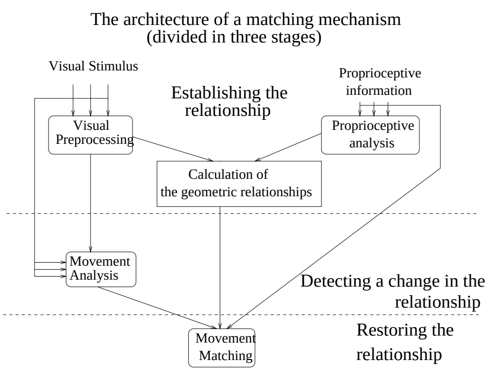

Figure 1: Architecture of the matching mechanism

Visual preprocessing

Within this module the following steps are taken:

- The demonstrator is detected either by their external observable characteristics or by focusing on the candidate regions where movement is observed.
- A vector is extracted which represents the current **posture** of the demonstrator. The complexity of this vector depends on the complexity of the movement one wants to imitate. For example, simple tasks such as following another robot require only a simple 2D vector representing the translation of the demonstrator robot in the 2D plane, while for imitation of head rotations one needs a 3D vector indicating the current direction that its head is facing in 3D space.

Proprioceptive analysis

Within this module we attempt to derive the imitator's own posture vector but this time using the imitator's proprioceptive information (i.e. the robot's current encoder values).

Establishment of relationships

In this module, the relationship between the demonstrator's and the imitator's posture is calculated. Depending on the task, simple geometric relationships are computed, such as the distance between the two agents and the relative angles (pan and tilt) between the two vectors, computed in an imitator-centred coordinate framework.

The three modules above deal with the first level of the imitation process: the establishment of the geometric relationship between the bodies of the demonstrator and the imitator. The next two modules deal with detecting a change in this relationship, and the imitator's attempts to restore it.

Movement analysis

When the demonstrator is performing an action, the posture vector is changing due to a combination of one or more types of motion: translation and rotation. In this module we are trying to detect what kind of motion the demonstrator is performing. The reliable complete recovery of the whole motion from sequences of 2D images is still a research issue. As we will describe later we are trying to simplify this part by either adding bright LEDs at specific parts of the demonstrator or constraining the type of motions that can be tracked.

Movement matching

The demonstrator's movement is effectively breaking the initial relationship between the two robots. The movement matching mechanism is the module that attempts to restore this relationship. There are effectively two alternatives:

- Knowing what action brings what effect and selecting the one(s) that have the desired effect.

- Learning the appropriate actions by trying different ones and having an error function that the agent is trying to minimize. This way the imitator can *learn* how to imitate.

Note that the imitator does not have to restore the relationship immediately. This can be deferred by keeping track of the relationship changes and restoring them later. Since the restoration occurs using proprioceptive information, the visual stimulus is not required to be present during the restoration.

We will now describe two experiments we are performing in order to instantiate and refine this architecture.

3.2 Experiments

The two experiments we have undertaken so far concern the imitation of translating movements in 2D space and the imitation of human head movements. In this section we present the experimental setup and describe these experiments.

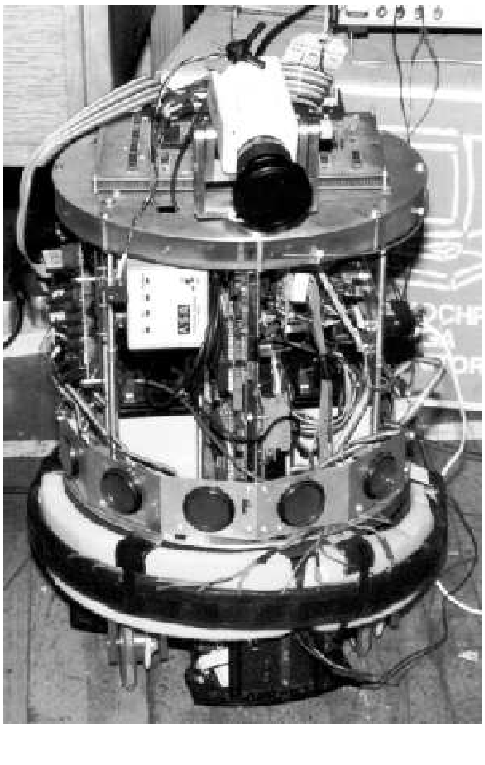

Figure 2: The learner robot in our 'following' experiments

3.2.1 Imitation of movements in 2D space

In the experiments described in (Hayes and Demiris, 1994) we investigated movement matching between a learner robot and a teacher robot moving on the floor of our lab. The learner matches the teacher's movements in order to learn how to traverse a maze. The learner[2] (figure 2) follows the teacher[3] (figure 3) along narrow corridors in a training maze. While it is doing that it associates the wall configuration it perceives with the actions it is performing as a consequence of its imitation of the teacher. At the end of the training phase, the learner is placed in a new maze (different from the training one) and successfully traverses it by using the rules it learned during the training (production rules of the form *if wall configuration then action(s))*. For a detailed description of these experiments see (Hayes and Demiris, 1994). We are currently attempting to investigate how well this approach generalises in the more unconstrained environment of our lab having a human as the teacher to be followed.

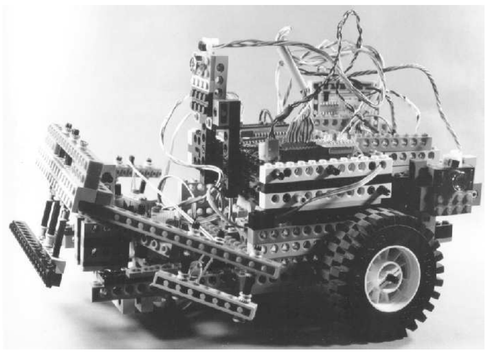

Figure 3: The teacher robot in our 'following' experiments

Since the problem of calculating the motion of an object when the observer is also moving is a difficult issue (Thompson and Pong, 1990) we have added bright LEDs on top of the teacher robot, which are tracked instead. The relationship between the two robots is one of **distance**. When the teacher robot moves this relationship is altered. The movement matching mechanism is essentially trying to maintain it by moving in such a way that the LEDs stay

[2] This consists of a RWI B12 mobile robot base, a Transtech Transputer board with 4 transputer units, IR and sonar sensors, a monochrome camera capable of 256 grey levels, and a radio modem.

[3] A Motorola-68000 based Lego robot equipped with infrared and bumper sensors.

in the same position in the image.

3.2.2 Imitation of human head movements

We are currently attempting to instantiate our architecture in a head movement imitation experiment. We are utilising a RWI robot (shown in figure 4) that has a pan-tilt head and a colour camera, and a network of Intel Pentium/486 microprocessors.

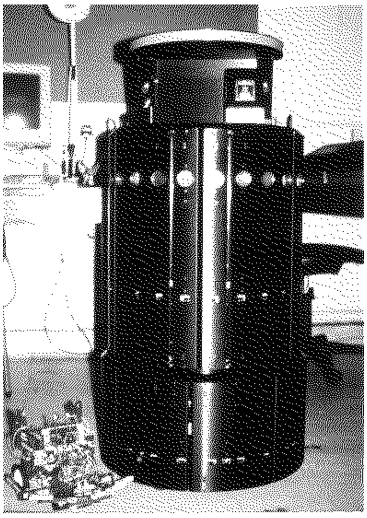

Figure 4: The robot we use for our head movement matching experiments

The human demonstrator sits in front of the robot, and the robot attempts to locate their face (figure 5). Locating human faces in an unconstrained environment is a particularly hard problem and there have been several attempts at solving it including (Brunelli and Poggio, 1993) and (Sung and Poggio, 1994). We use a neural network approach very similar to the one described in (Rowley, Baluja, and Kanade, 1995) where a part of the image is classified as a face or a non-face after being preprocessed (to correct illumination) and resized to pass through the neural network filter. This network (320 input units, 1 hidden layer, 1 output unit) has been trained with a back-propagation variant using the Stuttgart Neural Network Simulator and a collection of 1918 face examples (274 faces with 7 examples of each face, in front view plus views artificially rotated clockwise in the frontal plane by -9, -6, -3, 3, 6, and 9 degrees, and 9562 non-face examples, generated artificially from images that do not contain any faces at all. The complete face detection system [4] scans the image at several scales, and takes several minutes to scan a whole image.

To cope with our real-time requirements, we have constrained the system by asking the demonstrator to sit at a specific distance from the camera so the image need be searched only at a few specific scales, and by asking the demonstrator to start each head movement by looking straight at the camera. The speed can also be improved by asking the demonstrator to move their head before the search starts, performing frame subtraction on the sequence of images and concentrating only on the region where movement is observed.

Figure 5: Locating faces in an unconstrained environment - the dark area indicates the window where a face has been found

Once the face has been located, the part of the image that contains it is searched using template matching to locate the eyes, nostrils, and outer corners of the mouth. The facial normal is calculated using the planar method described in (Gee and Cipolla, 1994). Thus we know the direction in which the face is looking with respect to the camera. The direction in which the camera is pointing with respect to the robot's front is also calculated using the robot's head encoder values, but for simplicity we generally have the camera initially looking straight ahead with respect to the front of the robot's body and parallel to the floor. The angles between the camera and the facial normal can then be computed. The eyes and the mouth edge corners are tracked from frame to frame. As the demonstrator's head rotates[5] the facial normal changes and so do the angular relationships between the heads of the demonstrator and the imitator. The change is

[4]developed by Arturo Espinosa-Romero as part of his PhD thesis in the AI department of Edinburgh University

[5]Currently only rotational movements are allowed.

picked up by the movement matching mechanism which activates the commands that restore these angular relationships. There are currently two behaviours which are activated (one for restoring pan changes and one for restoring tilt changes), and these stay activated until the imitator's proprioceptive information indicate that the corresponding angular relationship has been restored.

4 Imitation as a learning mechanism

We propose using imitation as a way for robots to learn new skills. By matching the behaviour of the demonstrator using the movement matching mechanism described, and associating the environment that the imitator is perceiving with the actions it is performing due to its imitation of the demonstrator, the imitator can learn what is a good action to perform under certain circumstances. This could cut-down the search space that an agent would need in order to learn these skills using reinforcement learning considerably.

For example, our 2D movement matching experiments have shown (Hayes and Demiris, 1994) that a robot can learn how to negotiate the different type of corners in a maze using imitation. The teacher robot knows how to deal with the different types of corners, while the learner one initially does not - it merely knows how to imitate the teacher. When the teacher performs the correct action, the learner imitates it but also associates the action it is performing with the environment it perceives - in this case, the current wall configuration detected through its infrared sensors.

Experiments on head movement matching pave the way for learning experiments involving more than 2D translatory motion matching, such as movements involving robot arms in imitation tasks. Other possibilities for learning by imitation include learning the meaning of certain words that are communicated between the agents, by associating the sounds that are communicated with the distance and angle to a food source that the teacher knows but the learner does not (Billard, 1996). (Moukas and Hayes, 1996) report such experiments where the meaning of the words is communicated through a bee-like "dance" (a movement pattern). The learner then follows the teacher to a food source associating the distance and the angle from home with the visual language (the dance) perceived.

Kerstin Dautenhahn in (Dautenhahn, 1995) also presents a very interesting proposal: imitation as the basis for developing artificial social intelligence for autonomous robots. She uses imitation as a means of recognizing other robots and learning movement patterns. Her approach shares common ground with our work in (Hayes and Demiris, 1994) in that the imitator first follows the demonstrator and then analyses the movement that it made in order to derive its conclusions. Her work is based on a very interesting experimental habitat, the "Huegellandschaft", a hilly surface (some of the hills have recharging stations on top of them), which allows her to ingeniously avoid the binary distinction between space occupied by obstacles and plain surface. In this habitat, robots with different physical capabilities attempt to follow others, and by associating the energy consumption rates (dependent on the steepness of the surface) with the movements they make, they learn to distinguish between good and bad demonstrators and learn new movement patterns.

5 Conclusion and future work

In this paper we have presented a novel approach to robot learning which complements current learning approaches. The existence of other agents in the learner's environment is utilized as a source of information. By imitating the actions of other agents, the learner can quickly learn new behaviours that could be beneficial to it.

We have designed an architecture which deals with the fundamental problem of matching actions perceived visually with equivalent actions of the imitator. Our architecture views the problem as a dynamic process where the learner participates in an active role since actions are viewed as relationship changes. This has the additional significance that this process can become *bidirectional* with the robot imitator being able to detect (as humans are able to) when *it is being imitated*: relationship changes caused by *its own actions* are being restored *by the other agent*.

We have also described our robot experiments that are implementing this approach in two different types of movements, involving translation and rotation respectively. We have argued that imitation can be used as an effective way for learning through other agents what to do when.

Our next experiments aim at increasing the complexity of the tasks to be imitated by including movements that combine translation and rotation.

Apart from doing experiments that use this mechanism in order to learn new skills and knowledge, we are interested in developing an architecture that will redirect this new knowledge back into the imitation mechanism with the aim of improving its utility. Imitation is not *always* a good idea, since what is a beneficial for the teacher is not necessarily beneficial for the learner. We are developing an additional module that takes into consideration any reinforcement that the learner received during previous similar imitation acts and will stop the imitation process if the result is not beneficial for the learner.

We believe that the architecture and the experiments we described in this paper will allow robots to learn autonomously from other agents through imitation, and make better use of their social environment.

Acknowledgements

John Demiris is supported by a Teaching Assistantship from the Department of Artificial Intelligence of the University of Edinburgh, and a Junior Scientist Fellowship from the European Science Foundation. Gillian Hayes is supported by a lectureship from the University of Edinburgh. Thanks to John Hallam for his useful comments, and Simon Perkins for his comments and frequent technical advice. Special thanks to Arturo Espinosa-Romero for designing, implementing, and training the neural network for face detection, and for lots of useful comments on the difficulties of processing facial images. Also, thanks to an anonymous reviewer for many useful suggestions in improving this paper.

References

Billard A. (1996), "Learning to Communicate through Imitation", *MSc thesis*, Department of Artificial Intelligence, University of Edinburgh, September 1996.

Brunelli R. and Poggio T. (1993) "Face Recognition: Features versus Templates", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 15, No 10, pp. 1042-1052, October 1993.

Connell J. H. and Mahadevan S. (1993), *Robot Learning*, Kluwer Academic Publishers, 1993.

Dautenhahn K. (1995), "Getting to Know Each Other - Artificial Social Intelligence for Autonomous Robots", in *Robotics and Autonomous Systems*, 16, pp. 333-356, 1995.

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. (eds.), Lawrence Erlbaum Associates, 1988.

Gee A. H. and Cipolla R. (1994), 'Determining the gaze of faces in images", in *Image and Vision Computing,* 12(10): 639-647, December 1994.

Hayes, G. and Demiris, J. (1994). "A Robot Controller using Learning by Imitation" In Adam Borkowski and James L. Crowley editors, *SIRS-94, Proc. of the 2nd Intern. Symposium on Intelligent Robotic Systems*, pages 198–204, LIFIA-IMAG, Grenoble, France, July 1994.

Heyes C. M. (1993). "Imitation, Culture and Cognition" In *Animal Behaviour*, **46**, pp. 999-1010, 1993.

Huber M. J. and Durfee E. H. (1993), "Observational Uncertainty in Plan Recognition Among Interacting Robots", in *Working Notes of the IJCAI-93 Workshop on Dynamically Interacting Robots*, France, August 1993.

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

Mataric M. (1994), "Learning to behave socially", in *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behaviour*, MIT Press, August 1994.

Meltzoff A. N. (1990), "Towards a Developmental Cognitive Science: The Implications of Cross-Modal Matching and Imitation for the Development of Representation and Memory in Infancy", in *The Development and Neural Bases of Higher Cognitive Functions*, Adele Diamond (ed.), Annals of the New York Academy of Sciences, Vol. 608, 1990.

Meltzoff, A.N and Moore, M.K. (1989), "Imitation in Newborn Infants: Exploring the Range of

Gestures Imitated and the Underlying Mechanisms", in *Developmental Psychology* , 25:6, pp. 954-962.

Meltzoff, A.N and Moore, M.K. (1992), "Early imitation within a functional framework: the importance of person identity, movement and development", in *Infant Behavior and Development*, 15, 479-505, 1992.

Moukas, A. and Hayes, G. (1996), "Synthetic Robotic Language Acquisition by Observation" in *Proceedings of the Fourth Conference on Simulation of Adaptive Behavior,* Cape Cod, USA, September 1996.

Munch S. , Keuziger J. , Kaiser M. , and Dillman R. (1994), "Robot Programming by Demonstration (RPD) - Using Machine Learning and User Interaction Methods for the Development of Easy and Comfortable Robot Programming Systems", in *Proceedings of the 25th International Symposium on Industrial Robots*, Hannover, 1994.

Piaget J. (1951), "Play, dreams, and imitation in childhood", English translation from the French original by C. Gattegno and F.M. Hodgson, W. Heinemann Ltd, 1951.

Rowley, H. A. , Baluja S. , and Kanade T. (1995), "Human Face Detection in Visual Scenes", *Carnegie Mellon Computer Science Technical Report CMU-CS-95-158R*, November 1995.

Sung Kah-Say and Poggio T. (1994) "Example-based Learning for View-based Human Face Detection", AI Lab Memo 1512 / CBCL Paper 112, MIT, Cambridge, MA,December 1994.

Sutton R. (1990), "Reinforcement learning architectures for animats", in *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behaviour*, J.A. Meyer and S.W. Wilson (eds.), MIT Press, 1991.

Thompson, W. B. and Pong TC (1990), "Detecting Moving Objects", in *International Journal of Computer Vision*, 4, 39-57, 1990.

Whiten A. and Ham R. (1993), "On the Nature and Evolution of Imitation in the Animal Kingdom: Reappraisal of a Century of Research", in *Advances in the Study of Behaviour*, vol. 21.