

Imitation as a dual-route process featuring predictive and learning components: a biologically-plausible computational model¹

Yiannis Demiris and Gillian Hayes
Institute of Perception, Action and Behaviour
Division of Informatics
University of Edinburgh
5 Forrest Hill, Edinburgh, EH1 2QL
y.demiris@ic.ac.uk, gmh@dai.ed.ac.uk

X.1 Introduction

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 on its own, through classical conditioning and reinforcement, but to a large extent through its conspecifics, by observation and imitation. Species from rats to birds to humans have been observed to turn to their conspecifics for efficient learning of useful knowledge. One of the most important mechanisms for the transmission of this knowledge is imitation.

At the heart of the ability to imitate lies a mechanism that matches perceived external behaviours with equivalent internal behaviours of its own, recruiting information from the perceptual, motor and memory systems. This mechanism has been shown to be present even in newborn infants, which have been observed to imitate the facial gestures of their caretakers. In humans, malfunctions of this mechanism, surfaced as an inability to imitate, have been used as detectors of pathological disorders including autism and some forms of apraxia. This chapter presents a computational model of this mechanism.

Why is this an interesting problem? From an engineering perspective, designing an architecture that equips robots with the ability to imitate will allow the possibility for learning through demonstration. A human demonstrator can show an example of the task and the robot can learn by imitating the human. This will give people unfamiliar with robot programming the ability to teach robots to perform tasks. From a scientific perspective, research on imitation spans several disciplines including neurophysiology, psychology, psychophysics and pathology. The available data are often found at completely different levels of description, from neural recordings to behavioural data from human neuropathological examinations (for reviews, see (Carey, Perrett, Oram, 1997, Schaal, 1999)). Computational modelling has the potential to integrate data from several disciplines in a common platform. The need for very precise descriptions so that mechanisms can be implementable on computational and robotic platforms illuminates gaps in theories, and allows research to focus on filling these gaps. Even more importantly, computational modelling enables the development of predictions, which can be an important tool for directing further experiments.

In brief, this chapter offers the following contributions:

- It introduces a distinction between passive and active imitation, to distinguish between approaches where the imitator goes through a “perceive - recognise – reproduce” cycle (passive imitation) and the motor systems are involved only during the “reproduce” phase, and the approaches where the imitator's motor systems are actively involved even during the perception process (active imitation).
- It develops a computational architecture inspired by Meltzoff's Active Intermodal Matching mechanism (Meltzoff & Moore, 1997), hypothesised to underlie *infant* imitation. The architecture (that belongs to the “passive” category) is capable of imitating and acquiring any demonstrated movement that is within the capabilities of the imitator, but its “passive” characteristics do not correlate well with some of the biological data available for *adult* imitation.

¹ Chapter 13, “Imitation in Animals and Artifacts”, K.Dautenhahn & C. Nehaniv (eds), pp. 327-361, MIT Press 2002

First author now at Imperial College, Electrical and Electronic Engineering Department, y.demiris@ic.ac.uk

- To overcome the disadvantages of the passive architecture above, a novel, distributed imitation architecture with “active” properties is developed. The novelty of this architecture lies in that the same motor structures that are responsible for the *generation* of a movement are recruited in order to perform movement *perception*. Imitation becomes an active, predictive process: instead of going through a passive “perceive - recognise – reproduce” cycle, the imitator actively generates possible behaviours in parallel, executes them on internal forward models (internal simulators, or predictors) and selects among them based on the quality of the predictions they offer with respect to the states of the on-going demonstration. However the disadvantage of this route is that it is not capable of imitating demonstrated movements not already present in the imitator's repertoire.
- In order to get the best of both worlds, the two architectures above are combined into the final dual-route architecture: known movements are imitated through the active route; if the movement is novel, evident from the fact that all internal behaviours have failed to predict adequately well, control is passed to the passive route which is able to imitate and acquire the demonstrated movement.
- Computational experiments are performed that demonstrate the ability of the architecture to imitate, as well as acquire, a variety of movements including unknown, partially known, and fully known sequences of movements. They also reveal the inability of the architecture to match demonstrated movements with existing equivalent ones of its own, when they are demonstrated at speeds unattainable by the imitator.
- Finally, the developed architecture is proposed as a model of primate movement imitation mechanisms. A comparison is performed between the characteristics of the architecture and biological data on human and monkey imitation mechanisms. It is shown that they correlate well, thus offering possible explanations for the biological data. Perhaps more importantly, the computational experiments offer *testable* predictions regarding the behaviour of the biological mechanisms.

X.2 On Passive Imitation

The potential of imitation to ease the robot programming process was recognised by robotics researchers who realised that instead of going through lengthy and complex programming, robots could learn how to perform various tasks by observing a human demonstrator. Research by (Ikeuchi & Suehiro, 1992, Suehiro & Ikeuchi, 1992, Kuniyoshi, Inaba & Inoue, 1994, Hovland, Sikka & McCarragher, 1996, Kaiser & Dillmann, 1996, Kang & Ikeuchi, 1997, Yeasin & Chaudhuri, 1997) has successfully used human demonstration to program robots to perform assembly tasks. The techniques that have been utilised to achieve this differ, but the philosophy is essentially the same: the imitation process proceeds serially through the three stages of perception (visual systems), recognition (memory systems), reproduction (motor systems). There isn't substantial interaction between the three stages, and the motor systems are only involved at the final reproduction stage.

Approaches by (Hayes & Demiris, 1994, Dautenhahn, 1995, Demiris & Hayes, 1996, Billard, 1999) in the mobile robotics domain have attempted to follow a different approach by trying to devise imitation mechanisms that will work directly *without* a recognition stage. This line of work is relatively new, but it makes an important distinction: the imitator is not imitating because it is understanding what the demonstrator is showing, but rather, it is understanding it *because* it is imitating. Imitation is used as a mechanism for bootstrapping further learning and understanding.

The distinction between the two approaches is new in the field of robotics but not in psychology. Researchers studying imitation in infants have made a similar distinction while formulating hypotheses regarding the mechanisms underlying early infant imitation. (Meltzoff & Moore, 1977) first reported young infants, between 12 and 21 days old in the original report, being able to imitate both facial and manual gestures, including tongue protrusion, mouth opening and lip protrusion. The experimenters suggested that the infants are able to represent visual and proprioceptive information in a form common to both modalities. These results were against the popular belief of the time, that infants are only capable of imitation after 8-12 months from birth, and that imitation abilities are a result of the infant's cognitive development. Various hypotheses regarding the mechanisms underlying this phenomenon were compared by (Meltzoff & Moore, 1989), including the “innate release mechanism (IRM) model” which postulates that the demonstrator's behaviour simply triggers and releases equivalent fixed-action-patterns (FAPs) by the infant. The IRM model relies on the existence of a set of FAPs, but there isn't a precise specification of what this set is (Meltzoff & Moore, 1989). IRM was judged to be an unlikely candidate for two reasons:

- The range of actions imitated was wide, which would mean that the infant would have to have a large number of FAPs in its repertoire.
- The fact that the infants attempt to and succeed in improving the quality of the imitated act (Meltzoff, 1981).

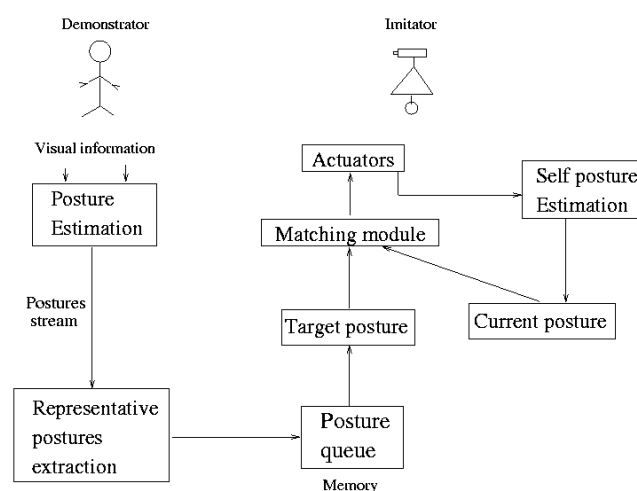
(Meltzoff & Moore, 1983, Meltzoff & Moore, 1989) put forward the “Active Intermodal Mapping” hypothesis which postulates that the infants use the demonstrator's states, perceived visually, as a target against which to direct their own body states, perceived proprioceptively. This hypothesis is particularly attractive in the case of facial or head movements for which the infant has no other way of knowing the state of its own body other than proprioception. The existence of a mechanism that matches stimuli between different modalities has also been advocated by (Maurer, 1993), but while Meltzoff's AIM mechanism appears to be activated as a choice made by the infant, Maurer argues that the infant's intermodal matching of stimuli is a by-product of what was termed neonatal “synesthesia”: the infant *confuses* input from the different senses. The infant, it is argued, does not register the modality that the stimuli appeared in but rather it responds to changes in the stimulation's intensity summed over all of the undifferentiated sensory modalities. Synesthesia is hypothesised to be a normal stage of early infant development: it is argued that the primary sensory cortex is not very specialised in infants, but with development it becomes so, the senses become more differentiated, and “true” intermodal matching develops. Whatever the exact mechanism is, the ability of the infant to match stimuli between modalities is well documented, and has been demonstrated between other modalities in addition to the visual/proprioceptive cases mentioned earlier, for example tactual/visual intermodal matching (Meltzoff, 1981, 1993).

At this stage it is useful to draw parallels between this work and the assembly and mobile robot imitation work mentioned earlier. There are a lot of commonalities between the passive imitation model in assembly robots and the IRM model in infants. Both rely on the existence of a set of predefined action patterns, which are triggered after the perception and classification of the visual input. This set, at least in the robot work is fixed, and frequently tuned to the requirements of the task in hand.

The mobile robot imitation work (Hayes & Demiris, 1994, Dautenhahn, 1995) is closer to the AIM hypothesis model, since the robots do not attempt to recognise the type of action performed by the demonstrator, but imitate directly. However there is a difference between AIM and the approach followed by the mobile robot researchers: the robot imitators do not attempt to match the demonstrator's state with their own (as AIM suggests), but usually achieve it by trying to maintain a quantity constant. For example, in (Hayes and Demiris, 1994) where a robot learns how to negotiate a maze by imitating the movements of another robot, the imitator robot simply tries to maintain the distance between itself and the demonstrator robot constant.

(Demiris and Hayes, 1996) presented a computational architecture that follows the AIM model more closely, and demonstrated it in the context of imitation of head movements by a robotic head (Demiris et al, 97). The details of this architecture have been presented elsewhere (Demiris and Hayes, 96, Demiris et al 97), but the essential parts are shown in figure 1:

Figure 1: the passive imitation architecture



The visual stimuli from the visual perception modules are fed into the posture estimation module, which at each iteration estimates and outputs the current postural state of the demonstrator (posture is defined here as the set of angles between all connected body parts of the agent). This posture stream is stored in memory after being filtered so only the “representative postures” (the postures that are sufficient to define the movement sequence to be reproduced) are retained. The postures are then fed into the movement matching module which outputs the motor commands needed to match these postures with equivalent postures by the imitator, perceived proprioceptively.

Experiments performed utilising this architecture on a robot head in the context of imitation of head movements performed by a human demonstrator have shown (Demiris et al, 97) that the architecture is capable of imitating any kind of demonstrated movement that the hardware of the imitator system can afford. It does so by having very low requirements on needed information: the postures of the demonstrator, perceived visually, and those of the imitator, perceived proprioceptively. The choice of posture as the unit of representation is not arbitrary. The postures of the demonstrator and imitator are always well defined and computable. In addition, postures have a high biological significance: animals frequently use them for communicative purposes (Bruce & Green, 1990, Groothuis, 1993), e.g. threat, appeasement and mating postures, and human body language. Their biological significance might even have led to the development of specialised feature detectors that respond selectively to postures. Indeed, work by Perrett and his colleagues has shown the existence of cells in the superior temporal sulcus (STS) area of the monkey's brain, that respond when a demonstrator assumes certain postures, for example, cells that are responsive to specific head views (Perrett et al, 1990, 1991). Cells selective to other body parts have also been reported (Tanaka93), see also (Carey, Perrett and Oram, 1997) for an overview.

Similarly, as (Tillery, Soechting and Ebner, 1996) point out, physiological studies throughout the somatosensory system have revealed discharges related to own static limb postures: unit activity is usually monotonically related to changes in joint angle. The degree of accuracy of a proprioception-based estimate of the static posture is not completely determined, and it has been shown that it improves when visual information about the state of the body part is also available (Desmurget et al, 1995), or is even being partially substituted by it when proprioception is not available in deafferented patients (Ghez and Sainburg, 1994). (Scott & Kalaska, 1995) demonstrated that cell activity in the monkey motor cortex is highly sensitive to changes in arm posture even if the resulting hand trajectory remains similar.

Finally, evidence that, at least some type of movements are controlled on the basis of a joint angular error has been provided by (Desmurget & Prablanc, 1997) who have shown that three-dimensional upper-limb movements are controlled via a mechanism that is comparing an estimate of the current postural state with a target value.

By relying on information known to exist in the human brain and requiring only an intermodal matching mechanism that is known to be within the capabilities of infants, the passive architecture above manifests itself as an attractive model for the infant imitation abilities. Could it be a universal model for movement imitation for later ages too? There are two issues that are against this. First, by virtue of its design, there is no concept of known and novel movements: all demonstrations are processed and imitated through the same mechanism. In addition, there is a clear separation between perception and action: the motor system is involved only at the late stages of imitation. Both these aspects have been challenged by recent biological data, and in particular human brain activation data, that indicate that actions are processed differently if they are known to the imitator than if they are novel, and that the motor system is already actively involved during the perception phase of the imitation. In the next section, a different type of architecture will be introduced that tackles these issues, and explains these biological data better. However, it would be premature to dismiss this “passive” architecture as invalid. Later on in this chapter, this architecture will be combined with the “active” architecture of the next section: it will be used as a learning component in what will be the final dual-route active-passive imitation architecture.

X.3 Active Imitation

Having seen the advantages and disadvantages of the passive imitation approach, this section will describe work towards the development of an architecture that tightly couples the perception and the generation of an action. The concept of internal forward models will be introduced, and the imitation architecture will subsequently be developed as a parallel set of behaviours paired with forward models. Using a dynamics simulator of a thirteen degrees of freedom robot it will be demonstrated how such an architecture can be used to generate an action as well as perceiving it when generated by others.

Definitions

The architecture that will be described makes extensive use of the concepts of behaviours and of forward models. A forward model of a controlled object (a “plant” as it is known in the control literature) is a function that, given the current state of the plant and a control command to be applied on it, outputs the predicted next state. Also, for the purposes of this work, a behaviour is defined as a function that, given the current state of the plant and the target goal(s), outputs the control commands that are needed in order to achieve or maintain the goal(s). Target goals might be implicit or need to be made explicit. For example, for a pick-object behaviour, the target object to be picked up must be stated explicitly and fed to the behaviour, while for a head-nodding-yes behaviour, the target goal (i.e. moving the head downwards) is already defined implicitly and the current state is enough to determine the motor commands needed to execute this behaviour. A behaviour is similar to what is known in the control literature as an “inverse model”, however, contrary to behaviours, inverse models do not usually utilise feedback about the current state, but output commands in a feed-forward manner. The boundary between behaviour and inverse model however, is not a rigid one since, as (Wolpert & Kawato, 1998) pointed out, “even control strategies, such as feedback control, which do not explicitly invoke an inverse model can be thought of as implicitly constructing an inverse model”.

Combinations of forward and inverse models have been used for various applications such as arm trajectory formation (Wada & Kawato, 1993) and supervised learning (Jordan & Rumelhart, 1992) among others. Internal forward and inverse models have also been hypothesised to exist in the human brain (Wolpert, Miall & Kawato, 1998), where they are utilised for a variety of tasks including sensorimotor integration (Wolpert, Ghahramani & Jordan, 1995), and motor control (Miall & Wolpert, 1996, Wolpert & Kawato, 1998).

The architecture

The fundamental structure of the architecture is a behaviour paired with a forward model (figure 2). In order to execute a behaviour within this structure, the behaviour module receives information about the current state (and, optionally, of the target goal(s)), and it outputs the motor commands that it believes are necessary to achieve or maintain the implicit or explicit target goal(s). The forward model provides an estimate of the next state(for time $t+1$) which is fed back to the behaviour, allowing it to adjust any parameters of the behaviour (an example of this, as will be shown in the next section (implementation), is adapting the gains of the PID controller used to implement a behaviour in order to achieve different movement speeds).

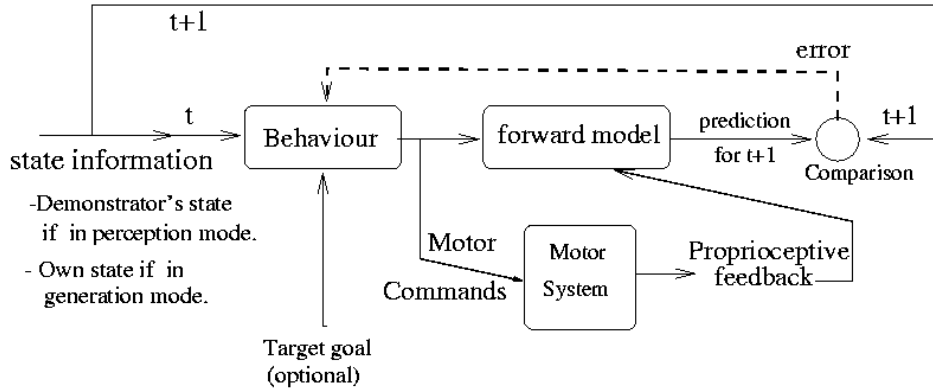


Figure 2: the architecture's basic building block, a behaviour paired with a forward model.

More importantly, the same structure can be used in order to match a visually perceived demonstrated behaviour with the imitator's equivalent motor one. This is done by feeding the demonstrator's current state as perceived by the imitator to the behaviour modules and having it generate the motor commands that it would output *if it was in that state and wanted to execute this particular behaviour*. The motor commands are inhibited from being sent to the motor system. The forward model outputs an estimated next state which is a prediction of what the demonstrator's next state will be. This prediction is compared with the actual demonstrator's state at the next time step. This comparison results in an error signal which can be used to increase or decrease the behaviour's confidence value, which is an indicator of how confident the particular imitator's behaviour is that it can match the demonstrated behaviour.

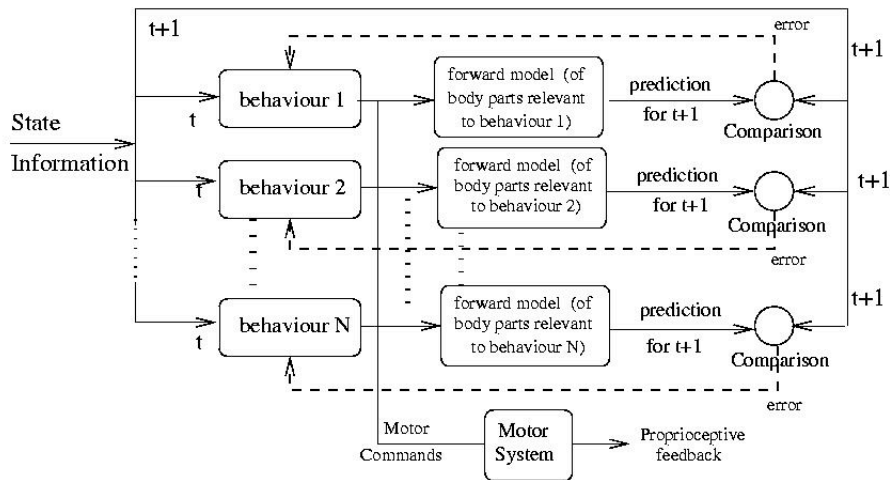


Figure 3: the complete active architecture, with multiple behaviours-forward models operating in parallel.

Figure 3 shows the complete architecture which consists of several of the structures that were described above, operating in parallel. When the demonstrator executes a behaviour, the perceived states are fed into the imitator's available behaviours which generate motor commands that are sent to the forward models. The forward models generate predictions about the demonstrator's next state which are compared with the actual demonstrator's state at the next time step, and the error signal resulting from this comparison affects the confidence values of the behaviours. At the end of the demonstration (or earlier if required) the behaviour with the highest confidence value, i.e. the one that is the closest match to the demonstrator is selected.

Implementation of the architecture

This section presents and analyses the results of implementing the architecture above on a dynamics simulator of a thirteen degrees of freedom robot. The results show that the architecture is capable of

correctly selecting the appropriate behaviour even when the demonstrator and the imitator have different dynamics. As behaviours, we implemented various movements involving moving both arms from the rest position to various positions. To allow for easy comparison, compact description and graph display of the results, we used an “alphabet”, the postures of the international standard semaphore code (ISSC) as the final positions that the arms reach (implementing here a total of 26 behaviours). For example, in figure 5 the eight windows are a snapshot of behaviours which currently show movement towards letters E, N, L, K, I, E, T, and E respectively. The experiments involved single movements (“letters”) and sequences of movements (“words”).

A thirteen degrees of freedom simulated robot was constructed and its dynamics simulated using the DynaMechs dynamical simulation library (McMillan, Orin & McGhee, 1995). “Bouncer” (figure 4) has three degrees of freedom at the neck joint, three at each shoulder, one at each elbow, and one at each wrist.




Figure 4: Bouncer, the thirteen degrees-of-freedom simulated robot

Bouncer operates under the effect of gravity, so if no forces are applied to a joint, the connected body part moves toward the ground. Movement is also subject to friction at the joints. The full experimental platform (figure 5) consists of two simulated robots, a demonstrator and an imitator, with similar body structure. In some of the experiments that were conducted, the dynamics of the two robots were the same, while in other they were different. The imitator is allowed to read the demonstrator's postural states (joint angles) in a crude simulation of visual capabilities. To account for the fact that in the case of the imitator, real vision and proprioception never result in perfectly correct values of the visually perceived demonstrator states and proprioceptively perceived imitator states, uniformly-distributed random noise is added to both of them before they become available to the imitator.




Figure 5: the complete experimental platform includes a demonstrator and an imitator (top left and right), and windows for displaying the behaviours that the imitator is currently considering.

Behaviours were implemented as sets of representative body postures, coupled with proportional-integral-derivative (PID) controllers (Astrom & Hagglund, 1996), that drive the controlled body part(s) through these key postures that constitute the behaviour. The PID controllers, one for each of the thirteen controlled joints, output torque values which are calculated by summing three terms: the proportional, integral, and derivative terms respectively, all of which are based on the error between the current state and the target state. If we take $e(t)$, the error at time t , to be the difference between the target state and the plant's state at time t , then the formula is as follows (for more details see (Demiris, 1999)):

$$T(t) = k_p e(t) + k_i \int e(t) dt + k_d \frac{de(t)}{dt}$$

where K_p , K_i , K_d are gains coefficients which were determined experimentally, and are allowed to adapt (within limits) in order to cope with different movement speeds.

An important issue to note is that the PID controllers perform best within a specific range of gain values; if they are not tuned within this range, they perform sub-optimally or might even lead to a destabilization of the controlled plant. The gain parameters K_p , K_i and K_d of all behaviours start having their optimal settings, but are allowed to adapt in order to reduce the prediction error between the anticipated states produced from internally executed behaviours and perceived demonstrator states. A simple adaptation mechanism was implemented for this: at each iteration, if the prediction for a joint angle value proves to be different from the actual value, the corresponding gains for the PID controller that controls that joint angle are increased or reduced (depending on whether the prediction underestimated or overestimated the actual value) by a small constant amount. However, although the gain parameters are allowed to fluctuate, (experimentally determined) upper and lower bounds are imposed in order to prevent the controller from outputting very high torque values and destabilising the plant. As it will be demonstrated later, this renders the perception of particular instances of some behaviours impossible.

The DynaMechs simulation package (McMillan, Orin & McGhee, 1995) includes libraries for simulating rigid body dynamics, and these were used in order to implement the forward models. The procedure involves four steps:

- (a) Applying the forces supplied by the behaviour, taking into account the current state (joint positions and velocities) of the robot.
- (b) Calculating all the forces exerted (including joint friction and gravity) and the inertias that are present in each joint.
- (c) Calculating the resulting accelerations recursively for each body part starting from the torso and moving towards the wrist.
- (d) Calculating the new state (joint positions and velocities).

Although in the experiments reported here, the forward models are directly coded in, they can also be learned by randomly generating motor commands, and using the resulting actual state as the target output state for the forward model, in what is sometimes called “motor babbling” (Bullock, Grossberg & Guenther, 1993, Jordan & Rumelhart, 1992) which is considered an important stage in the development of infants (Meltzoff & Moore, 1997, Meer, Weel & Lee, 1995).

Experimental results

The first set of experiments reported here consists of the demonstrator performing a behaviour that was composed of a single action and the imitator observing it having a set of behaviours in its repertoire. The number of behaviours is not important since the behaviours run in parallel and they are independent of each other (only their confidences need to be compared, a simple computation performed at the end of the demonstration). Experiments with six behaviours are shown for graph clarity reasons.

Figure 6 show an example confidence graph for an experiment where the demonstrator was executing the behaviour [R], while the imitator had in its repertoire behaviours [A, B, C, D, E, R]: the imitator's




Figure 6: Confidences of imitator's internal behaviours [A, B, C, D, E, R] when demonstrator executes behaviour [R];

behaviours start initially by having the same confidence, zero, and end up having a confidence correlated with their similarity to the demonstrated behaviour. Since the demonstrator is demonstrating the movement towards reaching the letter [R], the imitator's [R] behaviour gets the highest confidence, well above zero. All the other behaviours end up well below zero. Since all the behaviours in the first iteration assume the posture of the demonstrator, initially, and for several iterations, they all receive positive reinforcement, because they all look plausible at the onset of the movement. It is only after a few iterations (i.e. after the demonstrated movement has advanced) that some of the behaviours look less plausible (i.e. their predictions are very different from the actual demonstrated states), and their confidence levels begin to reflect that.

Similar results were obtained with a variety of different demonstrator and imitator dynamics, and with behaviours implementing sequences of movements ("words"). The full set of experimental results can be found in (Demiris, 1999).

The architecture described in this section only considered the cases where the demonstrated action or sequence of actions were already in the imitator's repertoire. If the demonstrated action is not in the imitator's set of known actions, it will neither be recognised or imitated. In the next section we deal with these cases, by blending the active imitation architecture of this section, with the passive one of the previous section, with the final dual-route architecture able to imitate and learn novel behaviours.

X.4 Combining Active and Passive Imitation Routes

In the experiments described in the previous section, the demonstrator does not perform any movements that are not already in the imitator's repertoire. One of the imitator's behaviours always ends up having positive confidence and is selected as the one to be imitated. However, if the demonstrator performs a behaviour that the imitator does not know (figure 8), all behaviours end up with negative or zero confidence and none of them is judged as suitable for imitation. This can be used as a trigger condition in order to attempt to learn the demonstrated behaviour. Note that in the work reported here, learning is the process of acquiring a behaviour, either its trajectory specifications or the motor commands needed to achieve it. Learning as used here does not imply generalisation or adaptation to different circumstances or any other processes as used in the field of machine learning (Shavlik & Dietterich, 1990).

The solution that is proposed in this section is to use the passive architecture of section 2 in order to imitate any movements not already in the imitator's repertoire. Figure 7 demonstrates this at the high level.

The representative postures that are extracted while the unknown behaviour is demonstrated are stored, and together with a PID controller, form new behaviours which are added to the imitator's set.

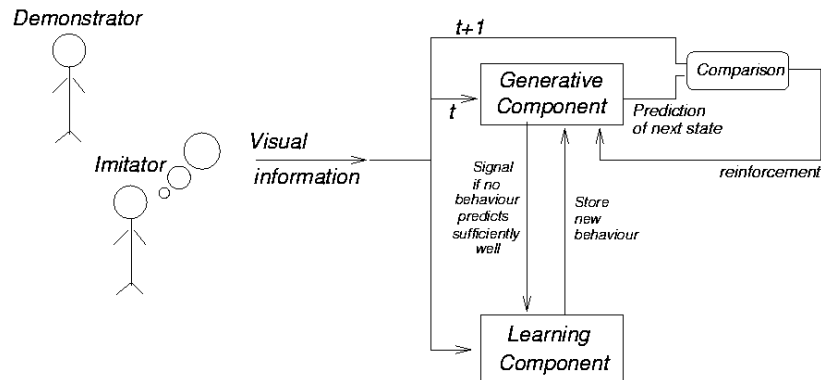


Figure 7: the complete dual-route architecture featuring generative-predictive and learning components.

Figure 8 shows the confidences plot of the imitator's behaviours [A-F] while the demonstrator executes 'R'. All of the imitator's behaviours end up with a negative confidence value; concurrently with the internal execution of candidate behaviours, the passive route was extracting the representative postures of the demonstrated movement. Since none of the behaviours performed adequately during the demonstration, the extracted representative postures, together with a PID controller, formed a new behaviour ([learned-R]) which was added to the imitator's repertoire.




Figure 8: Confidences of imitator's internal behaviours [A, B, C, D, E, F] when demonstrator executes behaviour [R];

The experiment is now repeated with the imitator equipped with the [learned-R] behaviour. Figure 9 shows the confidences plot of the imitator's internal behaviours while the demonstrator executes [R], where [learned-R] does end up with positive confidence. This demonstrates the ability of the architecture to learn new behaviours through the passive route and utilise them through the active one.




Figure 9: Confidences of imitator's internal behaviours [A, B, C, D, E, Learned-R] when demonstrator executes behaviour [R];

Experiments were also done with sequences of movements, covering the cases where all or some of the demonstrated components of the sequence are known, and the learning of the sequence is required, with equally favourable results (Demiris, 1999).

X.5 A model of primate imitation mechanisms

In the previous sections, two imitation architectures were presented, a passive and an active one, and a combination of them was subsequently developed. In this section, the biological plausibility of the combined dual-route architecture will be examined by proposing it as a model of primate action imitation mechanisms. First, a set of criteria that a model must meet in order to be useful will be presented, followed by an analysis of the model based on these criteria.

On criteria for useful models

Despite extensive use in the Artificial Intelligence, Cognitive Science and Artificial Life fields of the word "model", there are surprisingly few attempts to develop a set of criteria with respect to what a proper and useful model should be like (notable exceptions, frequently from scientists in other fields, include (Lehman, 1977, Rothenberg, 1989, Webb, 1993). By considering programs as the computational embodiments of a theory of how the modelled system might work, a substantial amount of theoretical work from the philosophy of science literature (Popper, 1972) on the possible criteria concerning the adequacy and proper form of a theory can be useful here. Having as a primary target the development of useful models, the following criteria were set for use in this work:

- A model should be clear on what it is a model of. This allows the incorporation of a number of pieces of evidence regarding the structure and behaviour of the modelled system as test sets for the plausibility of the model. A model should be accurate with respect to this evidence. The degree of accuracy, and the range of evidence that it agrees with can be used as a measure of the generality of the model.
- By virtue of being designed as an analog of the modelled system, a model should provide possible explanations for the data available about the modelled system.

- A model should be able to generate testable predictions. This is considered important in order to establish the scientific usefulness of the model, and demarcate it from useless exercises in computer programming.

The architecture as a model of primate imitation mechanisms

The first criterion of the ones described earlier requires an explicit statement on what the architecture is a model *of*. The aim of this section is to propose the dual-route architecture described in section 4 as a model of primate imitation mechanisms, and describe evidence from imitation research in primates that can be used as test sets for the plausibility of the model. Why specifically target primates? The main reason for this is that, in contrast to “lower” animals where the majority of research has concentrated on whether a certain animal is capable of imitation or not, there is a sufficient amount of data with respect to the underlying mechanisms of primate imitation to make a computational model possible.

The validation data

This section presents neurophysiological, psychological and brain activation data pertaining to issues important to the approach adopted in this work. Firstly, human brain activation data are presented, followed by neurophysiological data on “mirror neurons” found in the premotor areas of the monkey brain. The interplay between observing, imagining, performing and imitating a movement is the unifying theme of the psychophysical data presented right after, and the presentation of the validation data is concluded with the examination of data available on human imitation capabilities following brain damage, focusing on resulting apraxia disorders.

Activation of brain structures in humans

In humans, several experiments have investigated the interplay between action generation and action perception. (Fadiga et al, 1995) stimulated the motor cortex of human observers and recorded the motor evoked potentials (MEPs) from hand muscles, utilising the assumption that if action observation activates the premotor cortex (as it does in monkeys), this activation should induce an increase of the motor evoked potentials elicited by the magnetic stimulation of the motor cortex. They found a significant increase of the MEPs when subjects observed movements, and additionally the patterns of muscle activation was very similar to the pattern of muscle contraction present during the execution of the same action, i.e. the increase was present only in those muscles that are active when the human subjects executed these actions.

A different set of experiments with human subjects used Positron Emission Tomography (PET) brain scanning as a way of mapping the brain regions whose activations are associated with the observation of hand actions (Decety et al, 1997), as well as mental rehearsal (Decety et al, 1994) (similar to what was termed “internal generation” in the architecture of section 3).

In (Decety et al, 1994) normal subjects were asked to either passively observe movements of a virtual hand grasping objects or to imagine their own hand grasping objects, presented through a virtual reality system. Their brain activity during these conditions was mapped. The results demonstrated that cortical and sub-cortical motor structures were activated both during movement observation and movement imagery. It was concluded that consciously representing an action involved a pattern of cortical and subcortical activation that resembles the one observed during an intentionally executed action. It is important to note that during the observation condition, subjects were instructed to watch the movements of the virtual hand “as if it were their own hand” (this is similar to the first step taken by the behaviours in the active imitation route, i.e. internally assuming the observed state of the demonstrator). The importance of the observer's intentions during observation was further examined in (Decety et al, 1997) where subjects observed actions with the aim of either recognising them or imitating them later. The results showed that the pattern of activation was different between the two conditions, suggesting that the motivations and intent of the observer during the demonstration determine (or at least influence) which brain structures will be activated to process the incoming stimuli. (Decety et al, 1997) also examined the effect that the meaning of the observed actions has on the patterns of brain activation during observation. The results were striking: different brain structures were activated when the actions demonstrated were meaningless to the observer than those activated when the actions were known to the observer. This is very interesting since it indicates that knowing or

not the action demonstrated has an influence on the way this action will be processed in order to be imitated.

Mirror neurons in monkeys.

Neurophysiological experiments with macaque monkeys have revealed an important class of neurons in area F5 of the monkey's premotor cortex, which were termed "mirror neurons" (Gallese et al, 1996, di Pellegrino et al, 1992). These neurons were found to become active both when the monkey executes goal-oriented movements, and when it observes the demonstration of similar movements executed by another monkey or a human demonstrator. A variety of mirror neurons were discovered: grasping, tearing, manipulating, and placing objects neurons, among others. Some of the neurons were active only during the demonstration while some others remained active for a while after the end of the demonstrated action. The majority of the mirror neurons are active selectively when the monkey is observing a particular type of action (e.g. grasping), and some of them are highly selective not only to the type of action, but also to the particular way that the action is executed (e.g. grasping with the index finger and the thumb). The distance of the demonstrator from the monkey does not affect the responses of the mirror neurons, and control experiments have ruled out the possibility that the neurons are active simply as a response to particular visual configurations (for example, either of the demonstrator's hand or of the monkey's own hand) since most of them are also active when the monkey executes the action in darkness. Non-biological stimuli (for example, observing a set of pliers grasping the object) do not activate the neurons.

Observation, imagery, actual performance and imitation

Of relevance to the work presented in this chapter are also psychophysical experiments investigating the differences between observing an action, imagining an action and executing that action. (Vogt, 1995) performed a series of studies where subjects learned to reproduce a sequence of cyclical arm movements, either through repeatedly observing the sequence on a monitor, or through mentally or physically rehearsing the sequence. The results were very interesting since they demonstrated that observation or mental or physical rehearsal led to similar improvement in temporal consistency when the subject was later asked to reproduce the observations. Some further experiments (Vogt, 1996) with a short-term memory paradigm where subjects were allowed to observe the model movement only once, showed that timing imitation did not benefit from any further intermediate imitation (imaginary or physical) in the interval between the presentation of the model movement and the point where the subjects were asked to reproduce it. Related results were obtained in "mental chronometry" experiments by (Decety, 1996). Subjects were asked to perform a task either mentally or physically. The movement times required to execute the task were very similar irrespective of the modality of execution (mental or physical). In related sets of experiments (Decety et al, 1991, Wang & Morgan, 1992, Wuyam et al, 1995), subjects were asked to mentally perform tasks that would require different physical effort and found that autonomic responses (cardiac and respiratory activity) during motor imagery paralleled the autonomic responses to physically performing the task.

Brain & cognitive disorders and imitation abilities

Since imitation is a complex task involving the integration of information from multiple brain systems including perception, memory and motor systems, it has been used as a reference task for identifying and assessing various brain and cognitive disorders. In particular (and most relevant to this chapter) it has been used to identify and assess the various forms of *apraxia*, a "neurological disorder of learned purposive movement skills that is not explained by deficits of elemental motor or sensory systems" (Rothi & Heilman, 1997). Apraxia usually results from brain damage (usually in the left hemisphere) and its symptoms vary, giving rise to the various forms of apraxia, which are identified through a series of tests, that involve performance of actions on verbal command, imitation of meaningful and meaningless gestures, and gesture recognition and naming. A type of apraxia of particular relevance here is *visuo-imitative apraxia* (Mehler, 1987). Patients suffering from this apraxia are able to perform meaningful gestures when they are described verbally, or when they are asked to imitate them after a demonstration, but are unable to imitate meaningless gestures (Goldenberg & Hagmann, 1997, Merians et al, 1997). The nature of the demonstrated act, and in particular whether the act is known or not to the imitator, appears to be very important and determines whether or not the patient will be able to imitate it. This correlates well with the brain activation data described earlier, which show that different brain areas are activated depending on the nature of the demonstrated

act, and its meaning to the observer (Decety et al, 1997).

Two additional disorders are also of interest here: *autism* and *imitation behaviour*. Autism is a syndrome which includes abnormalities of social and communicative development, partially characterised by an inability to comprehend the viewpoints of other people (Baron-Cohen, Flusberg & Cohen, 1993). People suffering from autism display severe deficits in imitation and pantomime tasks (Smith & Bryson, 1994), which cannot be attributed to visual recognition memory, motor initiation and basic motor coordination deficits (Rogers et al, 1996). Furthermore, autistic children show deficiencies in empathy and joint attention tasks, as well as an inability to engage in pretend play (Charman et al, 1997). On the other side of the spectrum are patients that suffer from frontal-lobe damage, and display a pathological behaviour that has been termed “imitation behaviour” (Lhermitte, Pillon & Serdaru, 1986). These patients imitate the demonstrator's gesture although they were not instructed to do so, and some times even when told *not* to do so (de Renzi, Cavalleri & Facchini, 1996). An explicit, direct command from the doctor to the patient would stop the imitation behaviour but a simple distraction to a different subject was sufficient to see imitation reappearing, despite the patient remembering what (s)he had been told.

Explanations

Involvement of motor systems during perception

The human brain and mirror neuron activation data suggest that there is a motor system involvement during observation of movement. The explanation offered for these data by this work (Demiris, 1999) is that the motor system is activated in order to generate and internally simulate candidate behaviours, and offer predictions regarding the incoming perceptual data from the demonstrator. On a more specific note, the fact that some mirror neurons cease to be active when the demonstration is complete while others continue to be active for a while after the end of the demonstration can be explained if viewed within the composite nature of the organisation of the behaviours: more complex ones can be composed from elementary ones. Upon completion, a behaviour X ceases to be active; however, a behaviour X* which incorporates X as its initial step will continue to be active, since it is still capable of offering further predictions about the demonstrator's future states until X* completes its remaining steps. This suggests that the mirror neurons that cease to be active when the demonstrated action finishes represent that action specifically, while the other class of neurons which remain active represent sequences of actions that incorporate the demonstrated action as their first part.

The active route of the architecture understands an action by internally generating it. The observer does imitate the demonstrated movement internally, even when it does not do so externally. This feature of the architecture could explain why physically imitating a set of demonstrated movements does not aid their later recall (Zimmer & Engelkamp, 1996), as well as why physical rehearsal of a demonstrated behaviour does not lead to any significant differences in the levels of performance improvement from mental rehearsal or mere observation (Vogt, 1995). Since observation, imagery and imitation are done using mostly the same structures (behavioural modules and forward models) the same laws should govern their operation, which explains the mental chronometry data by (Decety, 1996), which indicate that it takes roughly the same time to perform a task mentally or physically.

Influence of content of the demonstrated action

The human brain activation data described by (Decety et al, 1997) indicate that different brain structures are activated during the observation of an action depending on whether the action is known to the observer or not. This is explained by the dual-route nature of the architecture: if the demonstrated act is known to the imitator, then the corresponding behaviour in the active imitation route will be activated. If the demonstrated act is not known to the imitator, then the passive route will be activated in order to extract the representative postures and acquire the demonstrated behaviour. Currently, there are no brain activation data to correlate with the behaviour of the architecture for the cases where the demonstration consisted of sequences of actions, and particularly partially-known sequences.

If the passive route is destroyed, the architecture will no longer be able to imitate any novel behaviours, although, with the active route intact, behaviours that are already known will be successfully imitated. This correlates favourably with the neuropathological data for patients suffering with visuo-imitative apraxia (Mehler, 1987).

Predictions

The computational studies of (Demiris, 1999) revealed limits to what the architecture can perceive, in particular with respect to movement speeds. For example, in figure 10, the demonstrator is executing a behaviour that the imitator does have in its repertoire ([cooler]); however this time the imitator is 400% heavier than the demonstrator, so it is not capable of executing the demonstrated behaviour at the demonstrator's speed levels. As a result, all the behaviours end up with very low (below zero) confidence values.




Figure 10: Confidences of imitator's internal behaviours [Cool, Cook, Cookie, Cooter, Cooler] when demonstrator executes behaviour [Cooler] at speeds unattainable by the imitator.

So, if the demonstration is performed at speeds that cannot be attained by the imitator, the demonstrated actions will not be understood, even if they are in the imitator's repertoire. By projecting this behaviour to that of the mirror neurons described earlier, the architecture offers a testable prediction: a mirror neuron which is active during the demonstration of an action should not be active (or possibly be less active) if the demonstration is done at speeds unattainable by the monkey. A further prediction with respect to mirror neurons has already been hinted at earlier in the explanation section. Mirror neurons that remain active for a period of time after the end of the demonstration are encoding more complex sequences that incorporate the demonstration as their first part. Further investigation through manipulation of the demonstration (adding further actions to it, while retaining the first part) should reveal the exact sequence that the neuron is encoding. Two less easily testable predictions regarding the mirror neurons are: (a) the existence of other goal directed mirror neurons and (b) the trainability of new mirror neurons. The first one predicts that there exist mirror neurons for other goal-directed actions: since perception and generation of an action is so tightly coupled, it can be expected that at least the most important actions in the monkey's repertoire (body postures that convey messages, for example threat postures, facial expressions, among others) should have mirror neurons associated with them. The second one predicts that, since the passive route provides the active route with new behaviours after their demonstration, it should be possible to create new mirror neurons by training the monkey to imitate a demonstrated action.

Discussion

The architecture, if viewed as a model, suggests that when humans and other primates observe a movement with the intent to imitate, they “put themselves in the place of the demonstrator”, and do what they would do if they were in the demonstrator's place. Understanding a demonstrated movement comes from internally generating alternatives and selecting among those, based on the quality of their

predictions. But why predict? Why not wait until the demonstration has finished and classify the result? From an evolutionary perspective, the ability to predict and its adoption during observation might have prevailed since it allows the animal to act/respond to an action of a conspecific before that (potentially non-beneficial to the observer) action has been completed.

The initial step taken by the observer, of putting herself in the position of the demonstrator, seems to be important too. Autistic children who suffer from an inability to do so, as witnessed by their poor performance in empathy, joint attention and pretend play tests, are unable to imitate. Normal children observing a human experimenter demonstrating an act but failing at it (for example, trying to pull apart a dumbbell, but failing due to finger slippage), do imitate the intended action of the human successfully, but do not do so when they see a mechanical device trying to do the same act but failing (Meltzoff, 1995). An explanation for this could be that the children did manage to imagine themselves in the place of the demonstrator when the demonstrator was human but not when it was not of biological nature. As already mentioned earlier, mirror neurons also do not respond when the action (e.g. grasping) is done with pliers, and not by a human arm (Gallese et al, 1996).

The dual route nature of the architecture is interesting too. It was already discussed in the explanations section earlier, that damaging the passive route leads to behaviour similar to that of visuo-imitative apraxic patients. What about the reverse condition? There is currently no evidence for the reverse dissociation, i.e. having the active route destroyed while retaining the passive one intact. This condition would be hard to detect, since known behaviours can still be imitated through the passive route as being novel. However, it is important to note that, essentially, the active route maps the observed movements to the imitator's internal ones, i.e. it serves as a recognition process. Any internal representations associated with these behaviours (including symbolic ones, for example the name of the behaviour, or emotional significance, intentions or any other attributes) can be retrieved through this route (note that it has been suggested (Gallese & Goldman, 1998) that the role of the mirror neurons is to facilitate the detection of the mental states of observed conspecifics by adopting their perspective). There are cases reported related to a disruption to this process: (Rothi, Mack & Heilman, 1986) reported two patients who could imitate demonstrated pantomimes but could not recognise (or discriminate among) them in what is termed as “pantomime agnosia”.

In the experiments reported in this chapter, all behaviours that are present in the imitator's repertoire are activated in order to generate alternatives and offer predictions as to what comes next. For efficiency reasons, it is conceivable that context can be used in order to select among all the available actions the ones that are applicable or at least relevant to the current situation. Although it is still early to speculate about the exact nature of this process, experiments with humans and monkeys have shown that actions applicable to a certain context are retrieved even if no action is required on behalf of the subject. (Rizzolati et al, 1988, Murata et al, 1997) have shown that there are neurons in the area F5 (same area with the mirror neurons) of the monkey's premotor cortex that are active during grasping movements, but are also active when the monkey views a graspable object. The interpretation favoured by the experimenters was that the responses of the F5 neurons represented the description of the presented object in motor terms, i.e. the visual features of the object are automatically translated into a potential motor action (regardless of whether the monkey intended to move or not). In humans, experiments with positron emission tomography have shown (Grafton et al, 1997) that observation of tools activated the premotor areas in the absence of any overt motor demand (it is interesting to note that the additional task of silent tool naming did not result in any additional activation of the premotor cortex, but tool-use naming did). These data indicate that the brain might indeed be using context to reduce the amount of behaviours that will be tried out.

Currently, a single presentation is enough for the architecture to acquire a new behaviour. It is not clear how this relates to primate behaviour but the architecture could be modified so that only frequently-occurring behaviours get acquired, or possibly the most biologically-important to the imitator. Then the passive route would essentially act as a short-term memory that would filter the behaviours letting only some of them through to the long-term memory of the active route.

X.6 Epilogue

In this chapter, a computational architecture for equipping robots with the capability to imitate was proposed, and subsequently proposed as a model of primate imitation mechanisms. The architecture has an active and a passive route: within the active route, the imitator mentally places itself in the place

of the demonstrator and internally executes ('imagines') candidate behaviours, eventually selecting among them based on the accuracy of their predictions regarding the demonstrator's incoming states as they are being perceived as the demonstration unfolds. If there are no behaviours that can predict sufficiently well, the passive imitation route learns the demonstrated behaviour, and adds it in the imitator's behavioural repertoire.

Computational experiments were performed instantiating this architecture using a dynamics simulator of a thirteen degrees of freedom robot, and showed the architecture to be able to imitate known behaviours, as well as acquiring new ones and successfully utilising them later. The architecture was also proposed as a model of primate imitation mechanisms and its characteristics, as explored by the computational experiments, were correlated with data on biological imitation generating a number of explanations and predictions. Completely understanding the underlying mechanisms of imitation is still a distant goal, and utilising computational and robotic architectures as models of real neural systems is a difficult and relatively new endeavour, but the biological data that are available indicate that the behaviour displayed by this architecture is towards the right direction.

References

- K. J. Astrom and T. Hagglund, 1996: PID Control, *The control handbook*, pp. 198-209, W. S. Levine (ed.), IEEE Press.
- S. Baron-Cohen, H. Tager-Flusberg and D. J. Cohen, 1993: *Understanding other minds: perspectives from autism*, Oxford University Press, Oxford, UK.
- A. Billard, 1999: *DRAMA: a connectionist model for robot learning: experiments on grounding communication through imitation in autonomous robots*, Ph.D Thesis, University of Edinburgh Scotland, UK.
- V. Bruce and P. Green, 1990: *Visual perception: physiology, psychology and ecology*, Lawrence Erlbaum Associates, Hove and London, UK.
- D. Bullock, S. Grossberg, and F. H. Guenther, 1993: A self-organising neural model of motor equivalent reaching and tool use by a multijoint arm, *Journal of Cognitive Neuroscience*, 5:4, pp. 408-435.
- D. P. Carey, D.I. Perrett and M. W. Oram, 1997: Recognizing, understanding and reproducing action, *Handbook of Neuropsychology*, Vol. 11, pp. 111-129, F. Boller and J. Graftman (eds.), Elsevier Science.
- T. Charman, J. Swettenham, S. Baron-Cohen, A. Cox, G. Baird and A. Drew, 1997: Infants with autism: an investigation of empathy, pretend play, joint attention and imitation, *Developmental Psychology*, 33:5, pp. 781-789.
- K. Dautenhahn, Getting to know each other - artificial social intelligence for autonomous robots, *Robotics and Autonomous Systems*, vol. 16, pp. 333-356.
- J. Decety, 1996: Do imagined and executed actions share the same neural substrate?, *Cognitive Brain Research*, vol. 3, pp. 87-93.
- J. Decety, M. Jeannerod, M. Germain and J. Pastene, 1991: Vegetative response during imagined movement is proportional to mental effort, *Behavioural Brain Research*, vol. 42, pp. 1-5.
- J. Decety, D. Perani, M. Jeannerod, V. Bettinardi, B. Tadary, R. Woods, J. C. Mazziota and F. Fazio, 1994: Mapping motor representations with positron emission tomography, *Nature*, 371, pp. 600-602.
- J. Decety, J. Grezes, N. Costes, D. Perani, M. Jeannerod, E. Procyk, F. Grassi and F. Fazio, 1997: Brain Activity during observation of actions: Influence of action content and subject's strategy, *Brain*, Vol. 120, pp. 1763-1777.

- J. Demiris, 1999: *Movement imitation mechanisms in robots and humans*, Ph.D Thesis, University of Edinburgh, Scotland, UK.
- J. Demiris and G. M. Hayes, 1996: Imitative learning mechanisms in robots and humans, *Proceedings of the 5th European workshop on learning robots*, pp. 9-16, Bari, Italy.
- J. Demiris, S. Rougeaux, G. M. Hayes, L. Berthouze and Y. Kuniyoshi, 1997: Deferred imitation of human head movements by an active stereo vision head, *Proceedings of the 6th IEEE international workshop on robot human communication*, pp. 88-93, Sendai, Japan, IEEE Press.
- M. Desmurget and C. Prablanc, 1997: Postural control of three-dimensional prehension movements, *Journal of Neurophysiology*, vol. 77, pp. 452-464.
- M. Desmurget, Y. Rossetti, C. Prablanc, G. E. Stelmach and M. Jeannerod, 1995: Representation of hand position prior to movement and motor variability, *Canadian journal of physiology and pharmacology*, pp. 262-272, vol. 73.
- L. Fadiga, L. Fogassi, G. Pavesi and G. Rizzolatti, 1995: Motor Facilitation during action observation: a magnetic stimulation study, *Journal of Neurophysiology*, vol. 73:6, pp. 2608-2611.
- V. Gallese and A. Goldman, 1998: Mirror neurons and the simulation theory of mind-reading, *Trends in Cognitive Sciences*, 2:12, pp. 493-501.
- V. Gallese, L. Fadiga, L. Fogassi and G. Rizzolatti, 1996: Action Recognition in the Premotor Cortex, *Brain*, Vol. 119, pp. 593-609.
- C. Ghez and R. Sainburg, 1994: Proprioceptive control of interjoint coordination, *Canadian journal of physiology and pharmacology*, vol. 73, pp. 273-284.
- G. Goldenberg and S. Hagmann, 1997: The meaning of meaningless gestures: a study of visuo-imitative apraxia, *Neuropsychologia*, 35:3, pp. 333-341.
- S. T. Grafton, L. Fadiga, M. A. Arbib and G. Rizzolatti, 1997: Premotor cortex activation during observation and naming of familiar tools, *Neuroimage*, vol. 6, pp. 231-236.
- T. G. G. Groothuis, 1993: The ontogeny of social displays: form development, form fixation and change in context, *Advances in the study of behavior* vol. 22, pp. 269-322.
- G. M. Hayes and J. Demiris, 1994: A robot controller using learning by imitation, *Proceedings of the 2nd International Symposium on Intelligent Robotic Systems*, pp. 198-204, Grenoble, France.
- Hovland, Sikka & McCarragher, 1996, Skill acquisition from human demonstration using a hidden markov model, *Proceedings of the 1996 IEEE international conference on robotics and automation*, pp. 2706-2711, Minnesota, IEEE Press.
- K. Ikeuchi & T. Suehiro, 1992, Towards an assembly plan from observation - part I: assembly task recognition using face-contact relations (polyhedral objects), *Proceedings of the IEEE international conference on robotics and automation*, pp. 2171-2177, Nice, France, IEEE Press.
- M. Jordan and D. Rumelhart, 1992: Forward Models: Supervised Learning with a Distal Teacher, *Cognitive Science*, vol. 16, pp. 307-354.
- M. Kaiser and R. Dillmann, 1996, Building elementary robot skills from human demonstration, *Proceedings of the IEEE international conference on robotics and automation*, Minnesota, USA, IEEE Press.
- S. B. Kang and K. Ikeuchi, 1997: Toward automatic robot instruction from perception - mapping human grasps to manipulator grasps, *IEEE transactions on robotics and automation*, 13:1, pp. 81-95.

- Y. Kuniyoshi, M. Inaba and H. Inoue, 1994: Learning by Watching: Extracting Reusable Task Knowledge from Visual Observation of Human Performance, *IEEE Transactions on Robotics and Automation*, 10:6, pp. 799-822.
- R. S. Lehman, 1977: *Computer simulation and modelling: an introduction*, Lawrence Erlbaum Associates Inc, New Jersey, USA.
- F. Lhermitte, B. Pillon and M. Serdaru, 1986: Human autonomy and the frontal lobes. Part I: imitation and utilization behavior: a neuropsychological study of 75 patients, *Annals of Neurology*, 19:4, pp. 326-334.
- D. Maurer, 1993: Neonatal synesthesia: implications for the processing of speech and faces, *Developmental neurocognition: speech and face processing in the first year of life*, B. de Boysson-Bardies et al (eds.), pp. 109-124, Kluwer Academic Publishers, Netherlands.
- S. McMillan, D. E. Orin and R. B. McGhee, 1995: DynaMechs: an object oriented software package for efficient dynamic simulation of underwater robotic vehicles, *Underwater robotic vehicles: design and control*, pp. 73-98, J. Yuh (ed.), TSI Press.
- A. L. H. van der Meer, F. R. van der Weel and D. N. Lee, 1995: The functional significance of arm movements in neonates, *Science*, vol. 267, pp. 693-695..
- M. F. Mehler, 1987: Visuo-imitative apraxia, *Neurology*, vol. 37, pp. 129.
- A. N. Meltzoff, 1981: Imitation, intermodal coordination and representation in early infancy, *Infancy and epistemology: an evaluation of Piaget's theory*, G. Butterworth (ed.), pp. 85-114, Harvester Press.
- A. Meltzoff, 1995: Understanding the intentions of others: re-enactment of intended acts by 18-month-old children, *Developmental Psychology*, 31:5, pp. 838-850.
- A. N. Meltzoff and M. K. Moore, 1977: Imitation of facial and manual gestures by human neonates, *Science*, vol. 198, pp. 75-78.
- A. N. Meltzoff and M. K. Moore, 1983: Newborn infants imitate facial gestures, *Child development*, Vol. 54, pp. 702-709.
- A. N. Meltzoff and M. K. Moore, 1989: Imitation in newborn infants: exploring the range of gestures imitated and the underlying mechanisms, *Developmental psychology*, 25:6, pp. 954-962.
- A. Meltzoff and M. K. Moore, 1997: Explaining facial imitation: a theoretical model, *Early Development and parenting*, 6:2, pp. 157.1-14.
- A. S. Merians, M. Clark, H. Poizner, B. Macauley, L. J. G. Rothi and K. M. Heilman, 1997: Visual-imitative dissociation apraxia, *Neuropsychologia*, 35:11, pp. 1483-1490.
- R. C. Miall and D. M. Wolpert, 1996: Forward models for physiological motor control, *Neural Networks*, vol. 9:8, pp. 1265-1279.
- A. Murata, L. Fadiga, L. Fogassi, V. Gallese, V. Raos, and G. Rizzolatti, 1997: Object representation in the ventral premotor cortex (area f5) of the monkey, *Journal of Neurophysiology*, 78, pp. 2226-2230.
- G. di Pellegrino, L. Fadiga, L. Fogassi, V. Gallese and G. Rizzolatti, 1992: Understanding motor events: a neurophysiological study, *Experimental Brain Research*, vol. 91, pp. 176-180.
- D. I. Perrett, M. H. Harries, A. J. Mistlin, J. K. Hietanen, P. J. Benson, R. Bevan, S. Thomas, M. W. Oram, J. Ortega and K. Brierley, 1990: Social signals analysed at the single cell level: someone is looking at me, something touched me, something moved!, *International Journal of Comparative Psychology*, 4:1, pp. 25-55.

D. I. Perrett, M. W. Oram, M. H. Harries, R. Bevan, J. K. Hietanen, P. J. Benson, and S. Thomas, 1991: Viewer-centred and object-centred coding of heads in the macaque temporal cortex, *Experimental Brain Research*, 86, 159-173.

K. R. Popper, 1972: *The logic of scientific discovery*, Hutchinson, London.

E. de Renzi, F. Cavalleri and S. Facchini, 1996: Imitation and utilisation behaviour, *Journal of Neurology, Neurosurgery and Psychiatry*, vol. 61, pp. 396-400.

G. Rizzolatti, R. Camarda, L. Fogassi, M. Gentilucci, G. Luppino, and M. Matelli, 1988: Functional organisation of inferior area 6 in the macaque monkey – part ii: area f5 and the control of distal movements, *Experimental Brain Research*, 71, pp. 491-507.

S. J. Rogers, L. Bennetto, R. McEvoy and B. F. Pennington, 1996: Imitation and pantomime in high functioning adolescents with autism spectrum disorders, *Child Development*, 67:5, pp. 2060-2073.

L. J. Gonzalez-Rothi, L. Mack and K. M. Heilman, 1986: Pantomime agnosia, *Journal of Neurology, Neurosurgery and Psychiatry*, vol. 49, pp. 451-454.

J. Rothenberg, 1989: The nature of modelling, *Artificial intelligence, simulation and modelling*, L. E. Widman, K. A. Loparo and N. R. Nielsen (eds.), pp. 75-92, John Wiley and Sons, NY, USA.

L. J. G. Rothi and K. M. Heilman, 1997: *Apraxia: the neuropsychology of action*, Psychology Press, East Sussex, UK.

S. Schaal, 1999: Is imitation learning the way to humanoid robots, *Trends in cognitive sciences*, 3:6, 233-242.

S. H. Scott and J. F. Kalaska, 1995: Changes in Motor Cortex Activity During Reaching Movements With Similar Hand Paths but Different Arm Postures, *Journal of Neurophysiology*, 73, pp. 2563-2567.

J. W. Shavlik and T. G. Dietterich, 1990: *Readings in Machine Learning*, Morgan Kaufmann Publishers, San Mateo, California USA.

I. M. Smith and S. E. Bryson, 1994: Imitation and action in autism: a critical review, *Psychological Bulletin*, 116:2, pp. 259-273.

T. Suehiro and K. Ikeuchi, 1992: Towards an assembly plan from observation - part II: correction of motion parameters based on fact contact constraints, *Proceedings of the IEEE/RSJ international conference on intelligent robots and systems*, pp. 2095-2102, Raleigh, USA, IEEE Press.

K. Tanaka, 1993: Neuronal mechanisms of object recognition, *Science*, vol. 262, pp. 685-688.

S. I. Helms Tillery, J.F. Soechting and T. J. Ebner, 1996: Somatosensory Cortical Activity in Relation to Arm Posture: Nonuniform Spatial Tuning, *Journal of Neurophysiology*, 76:4, pp. 2423-2438.

S. Vogt, 1995: On relations between perceiving, imagining and performing in the learning of cyclical movement sequences, *British Journal of Psychology*, vol. 86, pp. 191-216.

S. Vogt, 1996: Imagery and perception-action mediation in imitative actions, *Cognitive Brain Research*, vol. 3, pp. 79-86.

Y. Wada and M. Kawato, 1993: A neural network model for arm trajectory formation using forward and inverse dynamics models, *Neural Networks*, vol. 6, pp. 919-932.

Y. Wang and W. P. Morgan, 1992: The effect of imagery perspectives on the psychophysical responses to imagined exercise, *Behaviour Brain Research*, vol. 52, pp. 167-174.

B. H. Webb, 1993: *Perception in real and artificial insects: a robotic investigation of cricket phonotaxis*, PhD Thesis, University of Edinburgh, Scotland, UK.

D. Wolpert, Z. Ghahramani and M. Jordan, 1995: An Internal Model for Sensorimotor Integration, *Science*, vol. 269, pp. 1880-1882.

D. M. Wolpert and M. Kawato, 1998: Multiple paired forward and inverse models for motor control, *Neural Networks*, vol. 11, pp. 1317:1329.

D. M. Wolpert, R. C. Miall and M. Kawato, 1998: Internal models in the cerebellum, *Trends in Cognitive Sciences*, 2:9, pp. 338-347.

B. Wuyam, S. Moosavi, J. Decety, L. Adams, R. Lansing and A. Guz, 1995: Imagination of dynamic exercise produced ventilatory responses which were more apparent in competitive sportsmen, *Journal of Physiology*, vol. 482, 713-724.

M. Yeasin and S. Chaudhuri, 1997: Automatic robot programming by visual demonstration of task execution, *Proceedings of the international conference on advanced robotics (ICAR)*, pp. 913-918, Monterey, CA, USA, IEEE Press.

H. D. Zimmer and J. Engelkamp, 1996: Routes to Actions and their Efficacy for Remembering, *Memory*, 4:1, pp. 59-78.