An imitation mechanism for goal-directed actions

George Maistros and Gillian Hayes
Institute of Perception, Action, and Behaviour
Division of Informatics
University of Edinburgh
5 Forrest Hill, Edinburgh, EH1 2QL
{georgem, gmh}@dai.ed.ac.uk

5.4.2001

Abstract

This paper proposes a mechanism for imitating hand-object interactions like grasps, manipulations, etc. It is jointly based on neurophysiology and robotics. The aim of the mechanism is both the task goal and a reasonable degree of accuracy in the task itself. In addition, it attends to achieve both the problem of *how* to imitate and *when*. The test platform consists of two simulated robots and an object that they interact with. The results presented here show both successful imitation of familiar actions and poor imitation of unfamiliar interactions. It is intended that this mechanism is used to control the movements of real robots that would imitate other robots or humans.

1 Introduction

Imitation has been an invaluable mechanism to primates and can also be as such to artificial agents. However, whereas much research has focused on learning by demonstration and imitative learning, only the problem of *how* to imitate is usually addressed (acquiring the motor skills to perform an action). Equally important to *how* is *where* (in which context) and *when* to imitate. Recent neurophysiological data shed more light on the above questions.

Experiments on *macaque* monkeys show that neurons in the rostral part of inferior area 6 (area F5), namely *canonical* and *mirror* neurons, are involved in the tight coupling of perception and motor control (Rizzolatti et al., 1988). Mirror neurons, for example, are active both when the monkey observes a specific movement and when it executes the same movement. The discovery of mirror neurons in particular, had significant implications for robotics for they appear to provide the direct transformations from perception to action. It is postulated that mirror neurons form the fundamental basis for imitation and complex mechanisms like action understanding, learning, social interaction etc.

Among the first to be inspired by these neurophysiological data were Demiris et al. (1997) and Demiris (1999) who introduced a model of active and passive imitation. Matarić et al. (1998) developed a model that considers dancing behaviours such that ADONIS, a simulated humanoid robot, could perform a dance (Macarena). Fagg and Arbib (1998) developed FARS, a complete model of motor control, with the mirror system at its core providing the essential visuomotor transformations.

The idea, however, of the tight coupling of perception and action was originally introduced in the psychology field by Gibson (1986). He believed that perception is expressed in terms of action. He introduced the concept of *affordances* according to which the affordances of an object (we will only be concerned with objects at present) are those actions that one may apply to the object — the ones that the object *affords*. A mug, for example, affords a hand grasp (for picking up), a move (to approach the mouth), a mouth grasp (for drinking), etc.

Note that, this concept of affordances was found to express the main cortical input to F5 neurons (Rizzolatti et al., 2000). This means that the brain area that provides this input to F5, extracts affordances

Proc. TIMR 01 - Towards Intelligent Mobile Robots, Manchester 2001.
Technical Report Series, Department of Computer Science, Manchester University, ISSN 1361 - 6161.
Report number UMCS-01-4-1. http://www.cs.man.ac.uk/csonly/cstechrep/titles01.html

from visual information and primes (activates) F5 neurons accordingly. For instance, if a mug is observed, this area will only prime the F5 neurons concerned with mug grasping/moving, mug drinking etc. In fact, in the FARS model, the module that provides input to the mirror system is clearly implemented in terms of affordances.

Demiris and Hayes (1998) and Demiris (1999) introduced an imitation mechanism according to which the mirror system was expressed as an internal behavioural repertoire where each internal behaviour module would try to actively imitate demonstrated behaviours. Demiris and Hayes (1999) extended this idea developing a learning by imitation architecture which could also add new behaviours if imitation using the current repertoire was poor.

The imitation mechanism proposed here continues the line of thought of the above architecture with the main difference that perceptual modules are introduced separately from motor modules while the connections between them express the very function of mirror neurons (we will mention this later on). By doing so, we come closer to the fundamental underlying mechanisms of imitation and address more directly the questions of *how*, *where*, and *when*.

2 Neurophysiological background

Gallese et al. (1996) and Rizzolatti et al. (1996) performed single neuron experiments investigating further the visual and motor properties of the F5 neurons. All canonical and mirror neurons have similar motor properties. They discharge when the monkey actively performs hand and/or mouth movements (i.e. they are responsible for the motor coding of those actions). Their main difference, however, lies in the visual stimuli that trigger them. While canonical neurons are triggered by the *mere* presentation of 3D *objects*, mirror neurons discharge during the presentation of *hand* or *mouth interactions* with *objects*.

F5 mirror neurons

Mirror neurons are highly selective to the kind of action and to finger configurations that trigger them (either visually or motorically). So far, the effective goal-directed actions that have been reported are grasping, manipulating, and holding.

In addition, there is a relationship between the visual stimulus (observed action) and the motor stimulus (performed action). As such, mirror neurons have been classified according to their *visuomotor congruence* (Gallese et al., 1996). About 30% of them are *strictly* congruent, about 60% of them are *broadly* congruent, and nearly 8% of them are *non*-congruent. In strict congruence there is a correspondence between the observed and executed action both in terms of general action (e.g. grasp) and the type of action (e.g. precision grip). In broad congruence there is a loose correspondence between the observed and executed actions. For example, the neurons discharge when observing different kinds of grips and when executing a single kind of grip. In non-congruent neurons the visual stimulus bears no relation to the motor stimulus.

Experiments on the temporal relation of the discharge of mirror neurons found a variety of behaviours; neurons firing for the whole of the interaction, during early or late preshaping, and discharges starting shortly after contact with the object (Rizzolatti and Fadiga, 1998).

F5 canonical neurons

Mirror neurons share their motor properties with canonical neurons. The canonical neurons of a monkey are most effectively triggered when the monkey performs grasping, manipulating, or placing actions. Their visual properties, though, are different; they discharge at the mere presentation of (rather than hand or mouth interaction with) 3D objects.

Again, there exists a relation between motor and visual stimuli of canonical neurons — albeit somewhat differently defined. It appears that those canonical neurons that fire when the monkey performs a specific interaction, say precision grip, with an object, also discharge at the view of objects that, in fact, require similar precision grip to be seized.

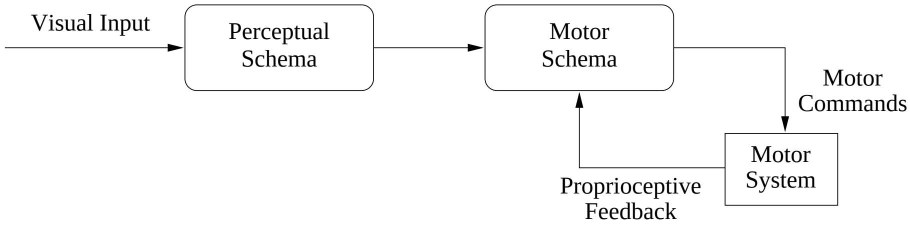

Figure 1: The basic building block of the schema network: a perceptual schema coupled with a motor schema.

3 The Mechanism

The proposed imitation mechanism employs the *Schema Theory* (Arbib and Cobas (1990) specifically designed the Schema Theory to express brain functions) to express the very function of the F5 neurons. According to the Schema Theory, perceptual and motor schemas are connected to form a network. This network receives input from perceptual schemas (in the form of perceptual structures) while its output is given from motor schemas (in the form of motor commands). Other kinds of schemas may be introduced for coordination, although this is not always necessary. The actual architecture of the network and the individual implementation of schemas is subject to the desired functionality.

The proposed mechanism forms a schema network using only perceptual and motor schemas. The basis of the network is the perceptual-motor schema pairing of figure 1. The perceptual schema holds the perceptual structure of a fraction of an action and the paired motor schema holds the motor structures necessary for the execution of that fraction of the action (Maistros and Hayes, 2000). Both the perceptual and the motor structures are expressed in terms of a sequence of robot articulation joint angles together with position and orientation information of the object(s) in peri-personal space. Schemas have individual confidences and only confident enough schemas get to act at any time. The main differences between the mechanism in Demiris (1999) and this one is that here, there is no Forward Model, and that a perceptual module is introduced separately from — yet connected to — the motor module. This separation is one that gives the temporal independence to our mechanism; each perceptual schema has to 'recognise' when it is best for the corresponding motor schema to start (or stop) imitating. Furthermore, this connection between schemas can directly reflect the relation between perceptual and motor stimuli found in F5 neurons (visuomotor congruence).

As such, to execute an action stored in such a pair, the perceptual schema compares the visual input with the stored perceptual structure and it calculates a measure of this match (its confidence). This confidence is then passed onto the motor schema. If this confidence is high enough, the motor schema would then output the necessary motor commands together with its own confidence (the confidence from the perceptual schema). The motor commands are calculated using a Proportional-Integral-Derivative module (Demiris and Hayes, 1999) that considers the current joint angles of the robot (proprioception feedback) and the current target joint angle (from the stored motor structure).

A complete schema network is illustrated in figure 2. It consists of several of the perceptual-motor schema pairs operating in parallel. All perceptual modules receive the visual input and produce confidence values that are passed into the corresponding motor modules. Then, highly confident motor modules produce candidate motor commands which are output into the competition module. The competition module, in turn, selects the motor command associated with the highest confidence value to be sent to the motor system for execution. If there are no candidate motor commands then the imitator tries to maintain its current postural position.

The significance of the schema network, although not yet thoroughly investigated, is threefold. First it reflects the high selectivity of F5 neurons, second it reflects their visuomotor congruence, and third their temporal independence.

Selectivity is expressed through perceptual structures representing specific actions that those structures would more confidently recognise, should the need be.

Visuomotor congruence however, is expressed by the variety of connections between perceptual and motor schemas. Each connection would express a motor schema 'reacting' to a perceptual schema 'seeing'

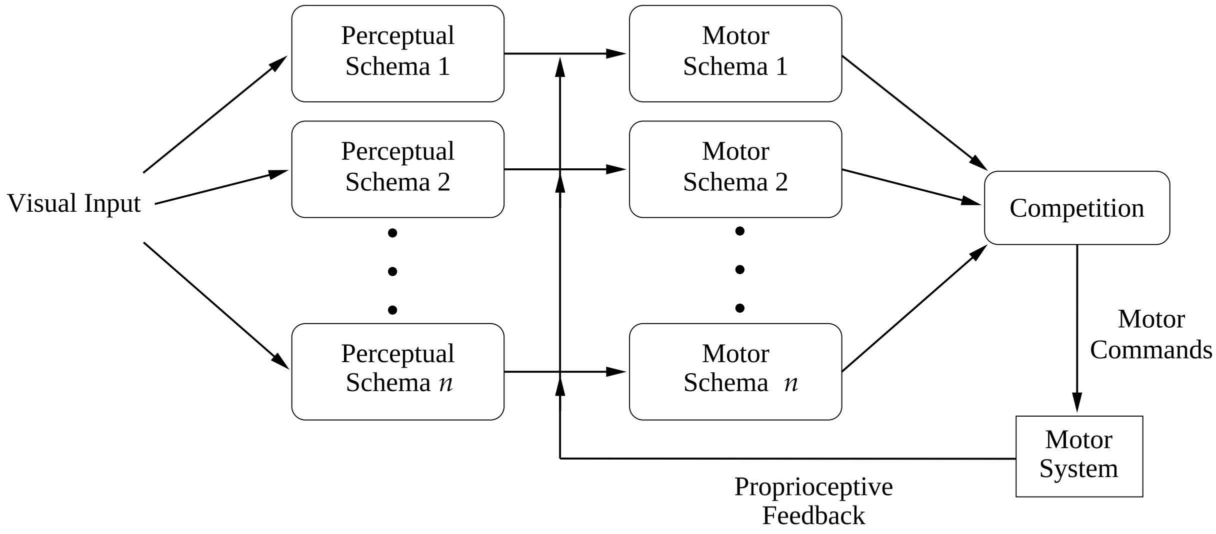

Figure 2: A schema network: several perceptual-motor schema pairs send candidate motor commands to the competition module which gives the output of the network.

an action. Hence, one may have specific perceptual schemas (representing finger prehension) connected to similarly specific motor schemas (again finger prehension) as already suggested above, or general perceptual schemas (representing grasps) connected to specific motor schemas (representing whole hand prehensions only). Notice that the former corresponds to strict visuomotor congruence, the latter to broad congruence. A further representation of broad congruence is to have a perceptual schema connected to multiple motor schemas (expressing a recognised action that is "known" to be achieved in different ways), or inversely, multiple perceptual schemas connected to one motor schema (different recognised actions are "known" to be achieved in only one way).

In fact, we postulate that by connecting perceptual schemas expressing only 3D objects to motor schemas representing, say, the grasps related to those objects, one may reach the function of canonical neurons.

Lastly, this mechanism permits perceptual-motor schema pairs to be defined to be temporally independent. No pair is told explicitly when to start, or finish, or for how long to be active. Thus, motor schemas have to produce candidate motor commands when perceptual schemas best represent the demonstrated action. The following experiments investigate the selectivity and most importantly the temporal independence of this mechanism.

4 Experiments

The implementation of the proposed imitation mechanism has been tested on a demonstrator-imitator scenario. The dynamics of two eleven degrees of freedom robots, a demonstrator and an imitator, were calculated using the DynaMechs dynamic simulation library (McMillan et al., 1995).

Each robot has three degrees of freedom at the neck joint, three at each shoulder, and one at each elbow (see Figure 3). The motor control of the demonstrator is explicitly controlled by a hand coded set of via points whereas the imitator is controlled by the schema network. The only inputs to the imitator's network are visual perception and proprioception. The imitator is allowed to 'look' at the joint angles of the demonstrator (visual perception), and 'feel' its own joint angles (proprioception). As this is a rough approximation to real vision and proprioception, noise is added onto the joint angles before the imitator can 'look' or 'feel' them.

The test scenario is involved with a set of goal-directed behaviours that the demonstrator performs while the imitator, aware only of a subset of them, tries to actively imitate them to the best of its knowledge. The full test set consists of the following five episodes:

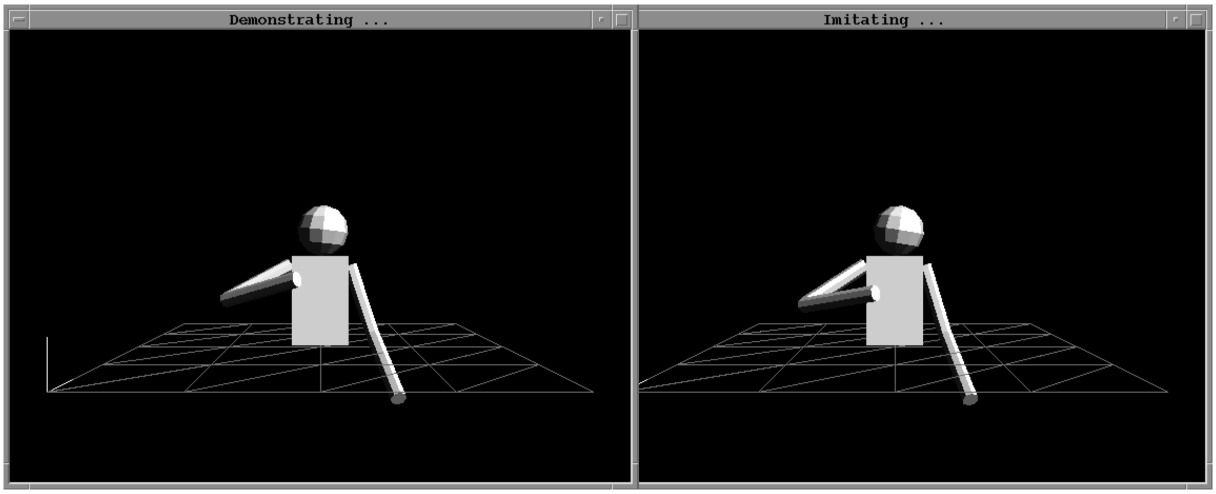

Figure 3: The implementation platform. The imitator, on the right, is currently imitating the demonstrator, on the left, drinking a 'virtual' glass of beer, *impolitely* (see section 4).

(i) Drinking a glass of beer, *politely*: reaching for a solid cube (resting on a surface at waist level), picking it up, bringing it to the mouth, and moving it back towards the surface.

(ii) Drinking a glass of beer *impolitely*: as in (i) save for the elbow which far away from the body throughout the action (see Figure 3).

(iii) Drinking a glass of beer, *less politely*: as in (i) save for the elbow which is a little further away from the body throughout the action (but much less than in (ii)).

(iv) Moving a glass using both hands: reaching for a solid cube (same as in (i)), picking it up using both hands, moving over the surface, placing it back on the surface, and moving hands away (towards starting posture).

(v) Passing a glass over to the other hand: reaching for a solid cube (resting on a surface at waist level), picking it up, moving it towards the other hand (the other hand approaches as well), passing the cube over to the other hand, placing it back on the surface (with the other hand), and moving hands away (towards starting posture).

The choice in these actions serves the need to investigate the recognition-reproduction capabilities of the imitation mechanism. It is important to test whether the schema network can distinguish between quite similar actions as well as between different ones. For this purpose, the imitator's schema network has only actions (i), (ii), and (iv) stored inside it. At the observation of known actions the imitator should successfully recall the correct actions, while unfamiliar actions should yield poor or no imitation.

Implementation of the imitator's network

Actions (i), (ii), and (iv) are first divided into several sequential parts, which are then stored into the imitator's schema network. Actions (i) and (ii) were divided into five parts and action (iv) into four. Each part is stored in perceptual-motor schema pairs which form the full network. In an attempt to avoid duplicate schema pairs, the first part of actions (i) and (ii) — which represent identical moves towards the 'glass' — appear only once in the network. The network consists of thirteen schemas (one approach schema, and four schemas for each of the three action).

Note that this division of actions is indeed arbitrary and assumes no knowledge or structure as in Matarić (2000) where such primitives are generalised and parametrised. We are however considering this perspective and ways to evaluate it.

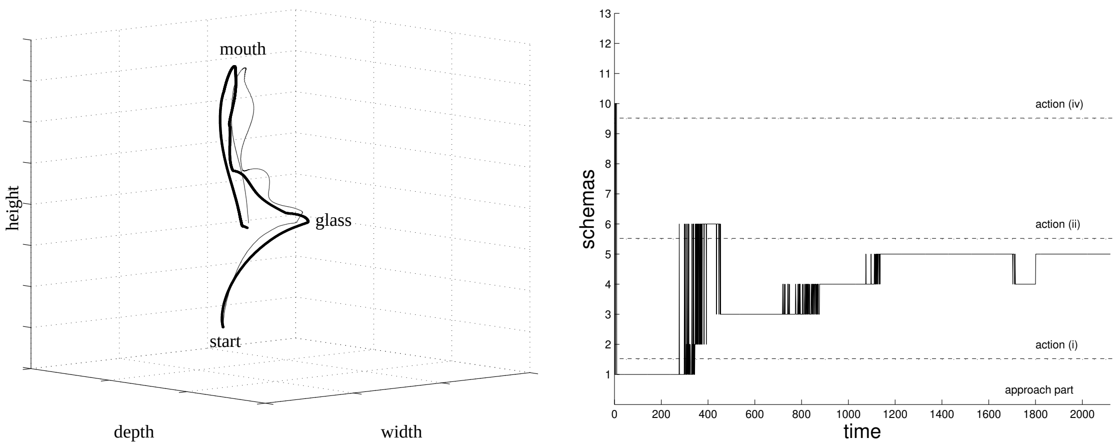

Figure 4: Episode (i): the imitator correctly imitates action (i).

Due to software restrictions, the imitator (and the demonstrator) does *not* act upon objects, although objects *are* represented, recognised, and virtually — yet effortlessly — moved. Thus, to avoid any further confusion, the objects present in an episode do exist but are not displayed. Consequently, there was no real need for the robots to have wrists or fingers. We do expect however to move to a different platform that would allow us to explore the full capabilities of our model.

5 Results

The full set of experiments consists of five episodes where the demonstrator performs actions (i)–(v) while the imitator uses its schema network to reproduce them. Figures 4–8 present the results. Plots on the left show the trajectories of the robot's wrists: the demonstrator's in bold and the imitator's in normal font. The robot's wrists are plotted from the start till the end of the demonstration — even if the imitator has not yet reached its last targets. Both wrists start from the same spatial point and move for the same amount of time. Plots on the right, show which motor schema wins the competition and is therefore the output of the schema network. Schemas are numbered from 1–13: schema 1 represents the approach part of actions (i) and (ii) — and inevitably (iii) and (v); consecutive schemas 2–5 represent the rest of action (i); consecutive schemas 6–9 the rest of action (ii); consecutive schemas 10-13 represent action (iv) in whole. The schemas have been grouped together for visualisation purposes only. Schema 0 corresponds to no winning schema when the imitator maintains its current posture.

Episode (i)

Figure 4 summarises the results of episode (i) (i.e. action (i) was demonstrated). The plot on the left, shows the path of the imitator's wrist which follows quite closely the demonstrator's one. Similarly, the actual behaviour of the imitator correctly reproduces action (i). The plot on the right shows the activity of the network and although it is not what one might expect, it is correct — if not, in fact, desirable.

Shortly after the end of schema 1 (approach part) the output of the network alternates confusingly between schemas 2 and 6. The confusion here is caused by both schemas (almost identically) guiding the imitator to pick up the 'glass'. Hence, although the network alternates between winning schemas, the imitator still receives correct motor commands, either from schema 2 or from schema 6. This phenomenon is also instantaneously observed at the very beginning of the episode between schema 1 and 10 which both take the imitator away from the same state (the starting state).

In fact, this alternation occurs to different degrees in every transition between schemas as is more clearly seen later on in the episode. This is in fact desirable for it allows schemas to get themselves in synchrony to the situation. For instance, at time 700–900, as the winning schema (schema 3) guides the imitator's wrist towards its last target point A, the next schema (schema 4) guides the imitator's wrist

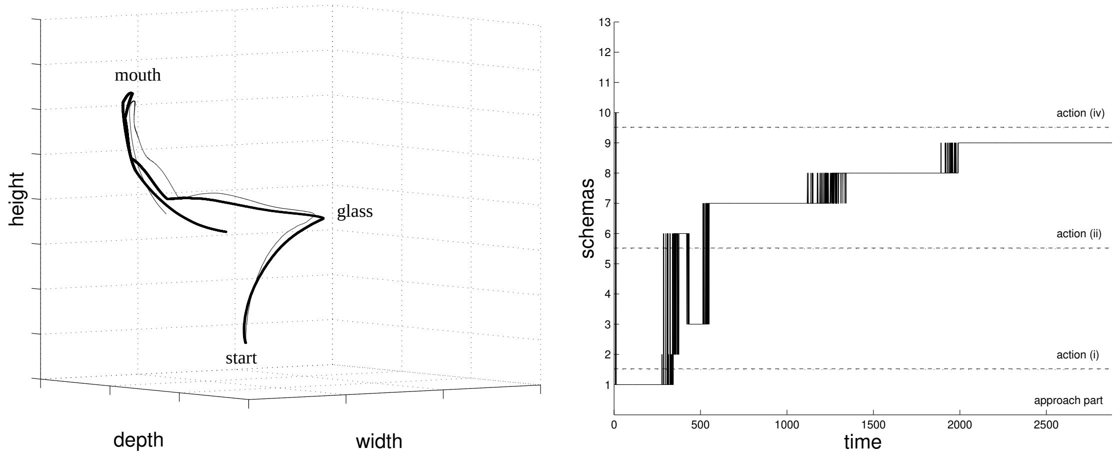

Figure 5: Episode (ii): the imitator correctly imitates action (ii).

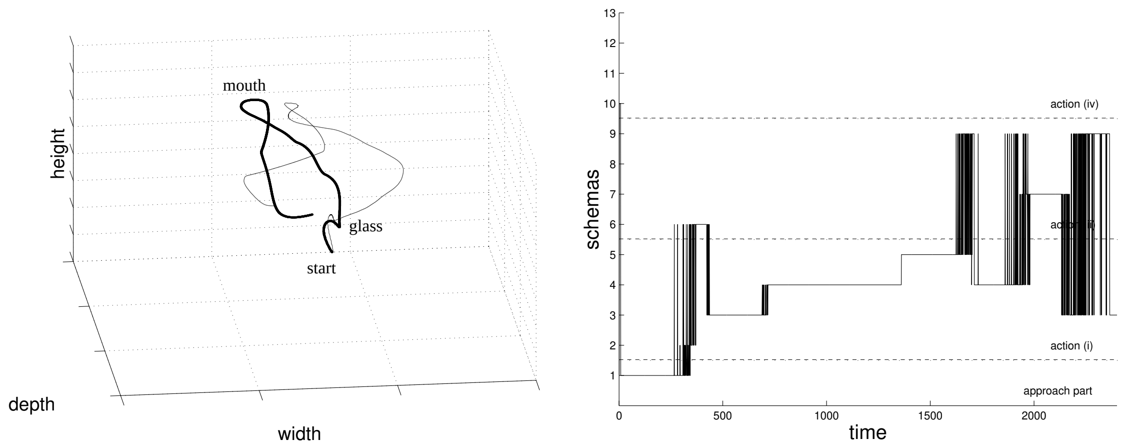

Figure 6: Episode (iii): the imitator correctly imitates part of the similar, yet unfamiliar, action (iii).

towards its first target point, which is also A. Thus, as seen in the trajectory plot, this does not interfere with the path of the imitator's wrist or with its actual behaviour.

One last thing to note is that around time 1700 the output of the network switches from schema 5 to 4 and then back to 5. This is only a minor mistake which little affects the imitator's behaviour. While schema 4 is responsible for moving the 'glass' upwards closer to the mouth, schema 5 moves the 'glass' downwards towards the surface. Therefore, as the two paths come closer to each other, the network is unable to distinguish from a single (and hence static) visual input the direction that the demonstrator is moving to. Still, however, this ambiguity is quickly resolved, causing only a small delay which is easily covered in the following time steps.

Episode (ii)

Figure 5 refers to the results of episode (ii). Again, the imitator's wrist follows closely the path of the demonstrator's wrist, although it is lagging a bit behind. This is mostly due to a minor error early in the episode, where schema 3 becomes the winning schema for a short while. (This was mostly due to the fact that schemas 3 and 7 are represent actions that start from similar states). However, this only introduces a short delay which has little effects on the behaviour of the imitator, save for a delay which

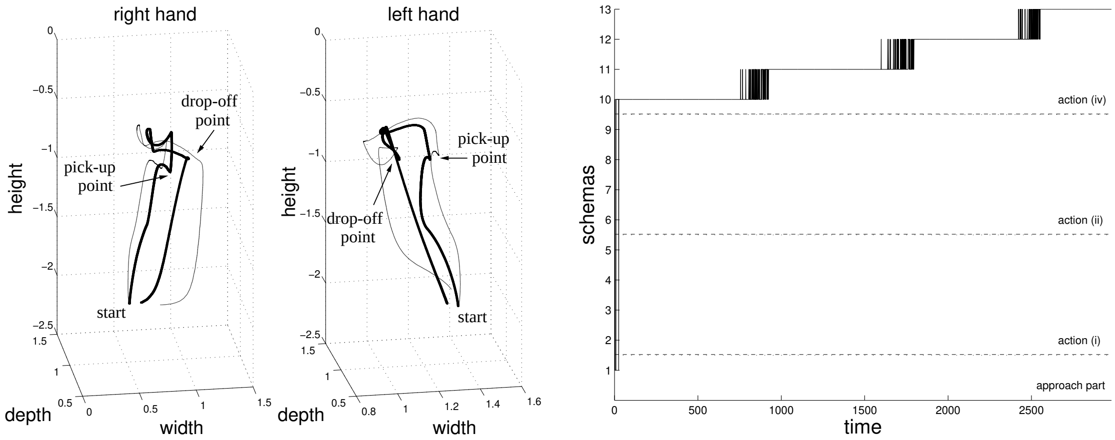

Figure 7: Episode (iv): the imitator correctly imitates action (iv).

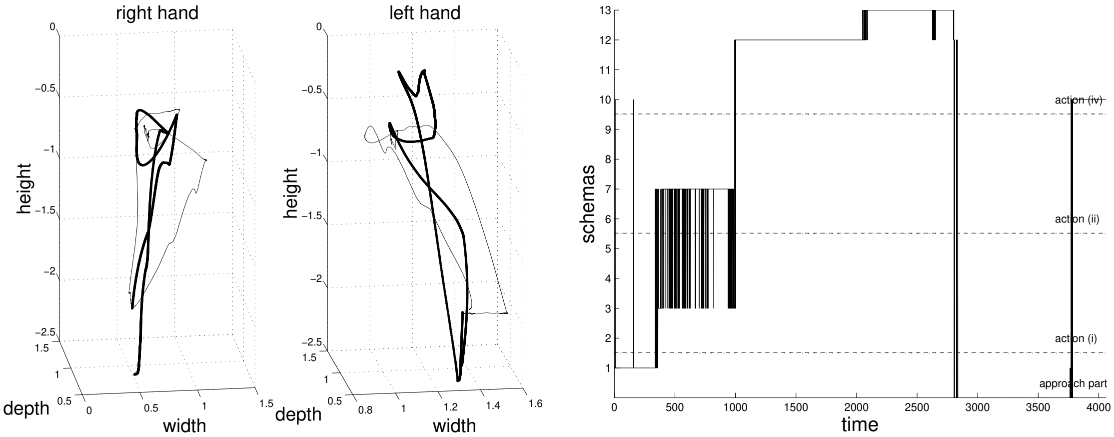

Figure 8: Episode (v): the imitator fails to imitate unfamiliar action (v).

is not covered even by the end of the episode. Should the imitator be allowed to move for a little longer (pursuing its final target states), it would quickly reach the demonstrator's final state.

Episode (iii)

Figure 6 refers to the results from episode (iii). Here, the imitator's performance is significantly poorer than in previous episodes where the demonstrated actions were familiar. Since action (iii) is partly similar to action (i) — and thus partly familiar — part of the output of the network resembles the first half of episode (i).

However, in the second half of the episode, the network mainly alternates between schemas 5 and 9. These schemas guide the imitator to move the 'glass' towards the surface, albeit through different postures. Such output of the network reflects the fact that the postures demonstrated are somewhat 'in between' those represented in these schemas. This confusion in the second part of the episode produces a significant error in the behaviour of the imitator that is also reflected in the trajectory plot, although the two trajectories do have a similar shape.

Episode (iv)

Figure 7 summarises the results from episode (iv). The trajectory plots have been separated for illustration purposes. As observed in the trajectory plots, action (iv) was imitated at a reasonable accuracy. It was observed, though, that while the 'glass' was moved, it was not always kept horizontal (an object that is picked up with both hands usually needs to be kept horizontal). However, this would involve much hand coordination, that is not explicitly coded into the motor schemas. Note that we are considering the use of virtual forces to coordinate hands and avoid collisions, usually among body parts (as was used in Matarić et al., 1998).

Episode (v)

Figure 8 summarises the results from episode (v). Again the trajectory plots have been separated for illustration purposes. As seen in the trajectory plots, action (v) was poorly imitated — if at all — which easily indicates unfamiliarity. Notice that while the demonstrated action is overall badly imitated, the right hand which is better represented in our schema network (more right-handed than left-handed actions), performs considerably better than the left hand. Lastly, although action (v) is unfamiliar, schemas 12 and 13 are active during time 1000–2700. However this is only because at that point action (v) resembles action (iv); the passing of the 'glass' from one hand to the other, action (v), appears similar to holding the 'glass' with both hands over the surface, action (iv).

6 Discussion

Our approach is jointly based on neurophysiology and robotics; inspired by a set of neurons that possibly form the fundamental basis for imitation, we propose a biologically-oriented mechanism implemented on a schema network that brings together perception and action.

Our aim in robotic imitation is not to achieve perfect imitation, but to come closer to task-level imitation, without however attending solely to the task goal. We hope to provide robots with perceptual and motor skills that are spatially as well as temporally independent.

We have tested our implementation of the network on a set of experiments that explore some of its capabilities and have exposed both strong features and potential weaknesses. The results suggest that the network is able to imitate familiar actions at a reasonable degree of accuracy, while unfamiliar actions produce poor imitation which is in most cases easily detectable. However, as is the case in much of the robotic imitation literature (Demiris and Hayes, 1999; Matarić, 2000), it is both hard and challenging to evaluate the performance of imitative systems.

References

Arbib, M. A. and Cobas, A. (1990). Schemas for prey-catching in frog and toad. *From animals to animats (SAB)*, 1:142–151.

Demiris, J. (1999). *Movement Imitation Mechanism in Robots and Humans.* PhD thesis, Institute of Perception Action and Behaviour, University of Edinburgh.

Demiris, J. and Hayes, G. M. (1998). Active imitation. DAI research paper; no.936, Dept Artificial Intelligence, University of Edinburgh.

Demiris, J. and Hayes, G. M. (1999). Active and passive routes to imitation. In *Proceedings of the AISB symposium on imitation in animals and artifacts.* Edinburgh, Scotland, UK.

Demiris, J., Rougeaux, S., Hayes, G. M., Berthouze, L., and Kuniyoshi, Y. (1997). Deferred imitation of human head movements by an active stereo vision head. In *Proceedings of the 6th IEEE International Workshop on Robot Human Communication*, pages 88–93. IEEE Press. Sendai, Japan.

Fagg, A. H. and Arbib, M. A. (1998). Modeling parietal-premotor interactions in primate control of grasping. *Neural Networks*, 11(7/8):1277–1303.

Gallese, V., Fadiga, L., Fogassi, L., and Rizzolatti, G. (1996). Action recognition in the premotor cortex. *Brain*, 119:593–609.

Gibson, J. J. (1986). *The ecological approach to visual perception.* Lawrence Erlbaum Associates.

Maistros, G. and Hayes, G. M. (2000). An imitation mechanism inspired from neurophysiology. In *Proceedings of EmerNet: Third International Workshop on Current Computational Architectures Integrating Neural Networks and Neuroscience.* Durham, UK.

Matarić, M. J. (2000). Getting humanoids to move and imitate. *IEEE Intelligent Systems*, pages 18–24.

Matarić, M. J., Williamson, M., Demiris, J., and Mohan, A. (1998). Behavior-based primitives for articulated control. In Pfeifer, R., Blumberg, B., Meyer, J.-A., and Wilson, S. W., editors, *Proceedings: From Animals to Animats 5, Fifth International Conference on Simulation of Adaptive Behavior (SAB-98)*, pages 165–170. Durham, UK.

McMillan, S., Orin, D. E., and McGhee, R. (1995). Dynamechs: An object oriented software package for efficient dynamic simulation of underwater robotic vechicles. In Yuh, J., editor, *Underwater Vechicles: Design and Control*, chapter 3, pages 73–98. TSI Press.

Rizzolatti, G., Carmada, R., Fogassi, L., Gentilucci, M., Luppino, G., and Matelli, M. (1988). Functional organization of inferior area 6 in macaque monkey. 2. Area F5 and the control of distal movements. *Experimental Brain Research*, 71(3):491–507.

Rizzolatti, G. and Fadiga, L. (1998). Grasping objects and grasping action meanings: the dual role of monkey rostroventral premotor cortex (area F5). *Novartis Foundation Symposium*, 218:81–103. In book: Sensory Guidance of Movement.

Rizzolatti, G., Fadiga, L., Gallese, V., and Fogassi, L. (1996). Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, 3(2):131–141.

Rizzolatti, G., Fogassi, L., and Gallese, V. (2000). Cortical mechanisms subserving object grasping and action recognition: A new view on the cortical motor functions. In Gazzaniga, M., editor, *The New Cognitive Neurosciences*, pages 539–552. Cambridge, MA: MIT Press.