Robotics and Autonomous Systems 37 (2001) 127–143

Robotics and Autonomous Systems

www.elsevier.com/locate/robot

Robot fostering techniques for sensory-motor development of humanoid robots

Adrian Stoica*

Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA

Abstract

The research efforts addressing the control of robot behavior have been polarized; most robots are either fully guided through "strong" programming, or given a few learning algorithms and left alone to explore the world. This paper explores the area in-between, looking at *robot fostering*, referring to techniques by which skills can be transferred to robots through a close interaction with humans. While robot fostering can be the basis of interaction with a variety of robotic shapes, it is most natural and human-friendly when the robot apprentices are anthropomorphic/humanoid. Fostering techniques discussed here include teaching/learning by imitation, teaching by description/explanation, reinforcement, aid and collaboration. The paper illustrates an experiment in teaching/learning by imitation. The human fosters the robot by first imitating its uncoordinated arm movements, thus helping the robot develop its sensory-motor associative system. The human then shows arm movements and the robot visually tracks them; consecutively, the robot is able to learn arm movements by imitation. Fostering techniques, in addition to robot learning/acquisition techniques and more efficient man–machine interaction are considered key elements contributing to the nascence of a new research field, *developmental robotics*, which would focus on the robotics counterpart of human cognitive and motor development. © 2001 Published by Elsevier Science B.V.

Keywords: Robot; Development; Humanoid; Imitation; Fostering

1. Introduction

Humans extract information from the surrounding environment and act upon the environment, transforming it for their benefit. For some tasks they introduced robots as intermediates. Robots are defined by their relationship with the environment and the humans.

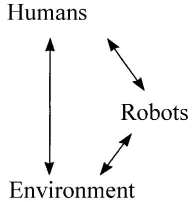

Robots receive orders and report on their interaction with the environment from which they extract information and upon which they act. Thus, robots can be seen as artifacts that:

- *extend human capability* for interacting with the environment (e.g. through them we can "see" on Mars);
- *replace humans* in some of their roles in this interaction.

The shape of the robot is chosen to fit the environment and the task to be performed. There are situations in which a non-human size and/or shape is not only desirable, but in fact necessary. For example, a worm-like shape is more appropriate than human shape for small robots that would burrow to penetrate ice on Europa; this is an *extension* of a capability since humans could

* Tel.: +1-818-354-2190, fax: +1-818-393-4272.
E-mail address: adrian.stoica@jpl.nasa.gov (A. Stoica).

PII: S0921-8890(01)00154-3

have never performed this task directly. In other cases, where humans have already been performing the tasks, the choice of the robot form is more subtle. By specifically defining the roles in which the humans are to be *replaced*, one can design non-anthropomorphic, tailored solutions that are more efficient than humans. For example, industrial robots on the fabrication lines are more efficient solution than humanoid robots at handling machine-customized tools. It should be noted, however, that those robots function in fully artificial, structured environments, mainly doing repetitive tasks. When roles and tasks previously performed by humans are very broad, environments in which they operate are human-oriented, and interactions with humans are a primary factor, anthropomorphic designs may offer some advantages.

The focus of this discussion is the development of robots dedicated to assist or substitute for humans in some of their roles. While "extension" robots may have various shapes, "substitution" robots can greatly benefit from anthropomorphism. This paper will concentrate on anthropomorphic/humanoid robots.

The main objectives of this paper are: (1) to argue for the need of humanoid robots, (2) to introduce the concepts and bring justification for robot fostering and developmental robotics, and (3) to provide an example on fostering humanoid robots to learn motor skills by imitation. The paper is organized as follows. Section 2 discusses the need for humanoids, and introduces developmental robotics. Section 3 discusses fostering techniques for cognitive and motor development. Section 4 presents an example in which the robot acquires eye–arm coordination and arm movements/patterns/skills. Section 5 summarizes and presents the conclusions.

2. Humanoid development

2.1. Anthropomorphic robots and humanoids

"*My research is not just in function, but in shape. In thirty years, in the twenty-first century, I think that human form will be essential in robots. In factories, which are for work, robots can be of any shape, but the personal robot, or "My Robot" as I call it, will have to exist in a regular human environment and be able to adjust to humans*." (Ichiro Kato of Waseda University, the "father" of WABOT-1 first biped walking robot — 1973.)

We are now 30 years later, and the date when we think humanoids will be around has been pushed by some studies for another 30 years or so [12,19]. Nevertheless, in recent years, due to research advances in robotics and some impressive demos such as that of the Honda robots, the idea of building humanoids becomes more commonly accepted, and a series of yearly international meetings has started [8]. This section presents some key arguments for humanoids and addresses the degree of anthropomorphism that is needed or useful for human substitution robots. Where direct human substitution is not important other shapes may be more efficient and will continue to be developed. The arguments arise in the following areas: (a) adaptation to human-dedicated environments/artifacts (including habitats, transportation systems, and tools), (b) interaction with (and acceptance by) humans, (c) efficiency of teaching/programming, and (d) testbeds for human-related studies (intelligence, prosthetics, interfaces, theories of behavior).

2.1.1. Adaptation to human-dedicated environments/artifacts (including habitats, transportation systems, and tools)

Environments designed and built for humans have the imprint of human shape. Buildings have stairs and elevators; trains and airplanes have seats and narrow corridors. All these impose constraints on the shape. Anthropomorphism could be the simplest solution for human-substitutes functioning in these environments and cooperating with humans. Human environments alone may not require human shape if robots are not required to be human-substitutes. Cats and spiders can live in human environments. But they do not change light bulbs, lay out the table, clean the house, and cannot carry humans in their arms to rescue them from a fire.

The designs of future habitats, particularly in space, could be modified, of course, if the environments we build put inconvenient pressure on the shapes of human-substitutes. But, at least as an intermediary step before we redesign the world, is it is not simpler to first develop a robot that fits this world, for which we already have a model available? Interestingly, environments themselves will become more and more intelligent. Our concept of robotics may entirely change as perception and intelligent systems become ubiquitous.

Human-substitution also implies the very important capability of using all the same artifacts/tools as humans do. Robots may be able to use other tools as well for performing the same tasks; this should not be perceived in the limiting sense, so long as retain back compatibility with human tools. Is anthropomorphism required for this? That is less obvious, but anthropomorphism could be a good starting point.

2.1.2. Interaction with, and acceptance by, humans

2.1.2.1. Interaction. Human interaction with robots will be easier if the robots are humanoid. The more humanoid the robot, the easier it will be for a human to intuitively understand its limitations and capabilities, to plan its actions, and to communicate directions clearly. Ideally, interacting should be so natural that even a child could easily utilize robot assistance.

"For a human-level intelligent robot to gain experience in interacting with humans, it needs a large number of interactions. If the robot has humanoid form, then it will be both easy and natural for humans to interact with it in a human-like way. In fact it has been our observation that with just a very few human-like cues from a humanoid robot, people naturally fall into the pattern of interacting with it as if it were a human. Thus, we can get a large source of dynamic interaction examples for the robot to participate in. These examples can be used with various internal and external evaluation functions to provide experiences for learning in the robot. Note that this source would not be at all possible if we simply had a disembodied human intelligence. There would be no reason for people to interact with it in a human-like way." [4].

2.1.2.2. Acceptance. "One of the most delicate and important factors to take into consideration for the success of service robots relates to the psychological aspects and to the implementation of techniques for human–robot interaction in "unprotected" and "unstructured" environments such as a house" [5].

Humans have a tendency to develop affinities based on resemblance. We can relate better to a chimpanzee than to a snake. Similarly, we find it easier to interact with a humanoid than with a large insect-like robot.

One should mention, however, that beyond the to anthropomorphization of the robots, some studies and theories such as the theory of Social Responses to Communication Technologies, indicate that on a more fundamental level, people's interaction with computers are identical to those between other human beings [28]. The recent field of interactive robotics, which includes personal robotics and service robotics [28], will play an important role in developing appropriate human–robot interaction means.

2.1.3. Efficiency of teaching/programming

Human intelligent behavior derives in part from interaction with the external environment. Attempting to create similar robot behavior may require similar interaction, i.e. similar ways of gathering information and perceiving and acting upon the environment, and that may require similar shape.

Human-oriented teaching of robots has important advantages. Humans teaching each other make great use of teaching by demonstration. This is more efficient than describing the movements in symbols/words, although these can also help the instruction. For example, in a context where the operator can visualize a target but the exact coordinate values of the position are not known, guiding via conventional software control is too complicated but analogic teaching can help [18]. For controlling complex motions a teaching pendant or a joystick are not as efficient as teleoperation in a master–slave configuration. The most efficient teleoperation is when the master and slave are identical; hence for a human it would be most natural and efficient to control an anthropomorphic robot.

Most current methods of extracting movement data rely primarily on sensors attached to joints and have important limitations. Giving robots vision to watch the movements themselves could significantly increase their capabilities. Humanoids can watch human body movements, e.g. arm movements, and then imitate them. In order to imitate the arm movement, the robot must have the necessary ability to transform images of the human arm into commands for its own arm. This visuo-motor coordination can be learned, as demonstrated, e.g., in [23]. There, an anthropomorphic robot learned to control its arm and then imitate the 3D movements of a master arm. The approach can be extended to other parts of the body.

An interesting possibility enabled by anthropomorphism is to have robots learn from videos/movies of humans.

In addition, according to the arguments of Johnson and Lakoff, the shape of our bodies is critical to the representations that we develop and use for both our internal thought and our language. If we are to build a robot with human-like intelligence, then it must have a human-like body in order to be able to develop similar sorts of representations [4].

Opinions differ about whether non-human entities can ever develop human-like intelligence. It is easy to imagine a video simulant displaying human-like intelligence, but how do we get this trait into real hardware and software? Experience seems necessary, but what kind of experience can an immobile computer have? A computer can simulate virtual life in a virtual reality; however, doing so in a realistic manner still requires solving most of the problems of robotics. Physical bodies may not be essential for artificial intelligence, but they would at least be convenient. Certainly the more humanoid the robot, the easier it would be to give it useful human behaviors.

2.1.4. Testbeds for human related studies (prosthetics, interfaces, theories of behavior)

Humanoids are potentially the best real-world models of humans; consequently they could provide the most efficient testbed for learning about humans. Prosthetics, ergonomics, and safety testing are among the first disciplines that could make use of human-like robots. Theories of child development, concept formation, motor behavior, intelligence, etc. find an ideal testbed in humanoids. Another important aspect is the effect of such a challenge (building a humanoid) in domains such as Artificial Intelligence. Now that IBM's Deep Blue has won a competition with the world's human chess champion, new challenges need to be formulated to drive artificial intelligence research.

2.2. Hot jobs for humanoids

Two "areas of employment" are seen as the most promising from the perspective of "human substitutes": earth jobs in the areas of robot assistants, service robotics, and hazard rescue and jobs in the area of collaborative space exploration. Engelberger [7] predicted that service robotics will outstrip industrial robotics sometime early in the 21st century. While in 1994, the industrial robot industry shipped about 65,000 robots, the market prediction for elderly care robots alone amounts to millions [7]. We are at the dawn of a new era in robotics. Many households now have a personal computer; not far in the future it may become common to have a personal robot. Response to unexpected hazards such as smoke, fire, steam, floods, and radiation, in which robots would perform rescue missions in human habitats, appear as well to be high pay-off application.

2.2.1. Attack on disabilities

An *IEEE Spectrum* article [26] cited US Census Bureau statistics indicating that 49 million people in US were in some way disabled. Nearly half of these, i.e. almost a 10th of the US population have a severe disability in which a physical shortfall is coupled with a mental illness such as Alzheimer disease. Humanoid robots can play a major role in acting as personal assistants for people affected by these disabilities. In addition, byproducts of the development of anthropomorphic systems are likely to benefit the human rehabilitation process.

2.2.2. Human–machine partnership on planetary outposts

Long duration, affordable and productive human presence in space will require a seamless human–machine partnership in which collaboration is the key. Humanoid robots are expected to play an important role in future human populated space colonies, as well as on Earth. They could assist astronauts during missions. They could build facilities prior/between human visits.

2.3. Developmental robotics: growing robot cognitive and motor skills

Current robots depend largely on being programmed in the same way computers are programmed. This is awkward for in field programming, e.g. teaching the robot new ways of solving problems or manipulate objects directly in the real-world workplace, in particular in space. The robots lack the capability of being taught easily, in a human-oriented way. New research efforts address human-centered techniques of teaching/programming robots, which would

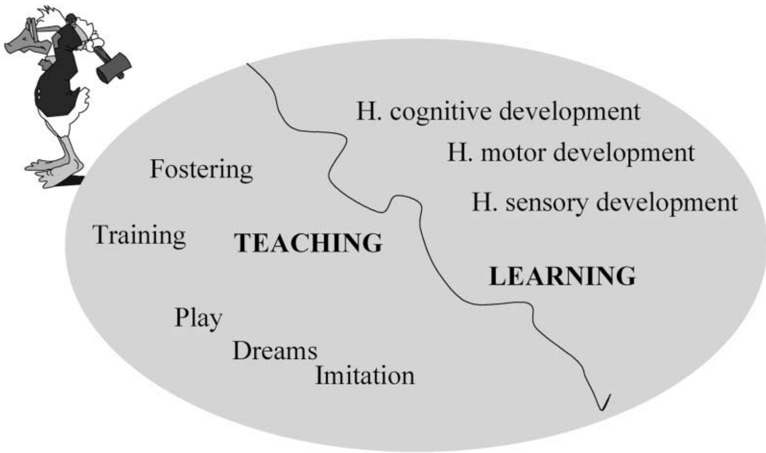

Fig. 1. Teaching and learning ensemble.

provide a paradigm shift from programming robots in machine-oriented language to teaching/fostering robots in similar ways we teach humans [22]. These techniques target humanoids with similar types of sensory-motor capabilities as the humans. Humanoid robots are the best candidates for learning by imitation, from the demonstration of human motor behaviors, and for being fostered, e.g. in learning to walk.

It appears impossible to program a humanoid from beginning to end. It may be easier instead to grow it like a child. Thus, a new area of research is foreseen, which is referred here as *developmental robotics*. Its aim would be to develop knowledge, methods, and techniques for having the robots, like the humans, develop gradually their cognitive and motor skills from the interactions with humans, other robots, and environment. Very much like child cognitive development [14], robot development would benefit from play, dreams and imitation. Teaching would become as important component as learning. Fig. 1 illustrates the current separation in addressing teaching and learning. While the learning area has received, by comparison, more attention, teaching has been less explored. In developmental robotics these two aspects would be treated as an ensemble.

The idea of developmental robotics has been suggested independently by several researchers, e.g. Stoica [22,25] and Asada [1]. To cite from the earliest reference known to this author: “Investigation in this area could lead to a new direction of robotics research, possibly called *developmental robotics*, aiming at building robots based on mechanisms similar to human cognitive and motor development.” [22, pp. 146].

3. Fostering techniques for cognitive and motor development

The degree to which the human controls robot behaviors tends to polarize toward the extremes. At one end the human is in charge of everything: controlling the robot as a marionette (Fig. 2, left), or feeding its brain with everything one assumes the robot should know (Fig. 2, right).

At the other end, the robot is seldom provided with a set of learning algorithms and left alone in the world to learn by exploration, build maps, make sense of it by itself. This is a very challenging task, and in many respects we throw the robot to the lions (i.e. the dangers in the real, unpredictable world) (Fig. 3).

The midway is to have an active, continuous involvement of a human (or of a teacher robot) during the development of the set of capabilities the robot needs in the world. In the animal world fostering is considered an important component to ensure survival of the species. Interestingly, it is been observed that the more “advanced” a species, the longer the period of immaturity of its offspring — in other words the longer the parents need to foster their children [3]. It is this period when the young ones develop the skills

Fig. 2. Human in charge of everything: (left) controlling the robot as a marionette, or (right) feeding its brain with everything one assumes the robot should know.

that would make them successful in life. The parents act as first teachers taking the young ones through various phases of learning. In time the grown-up will in turn teach others (not seldom themselves learning more through teaching) (Fig. 4).

3.1. Phases of learning

Humans learn by themselves or from others. In the initial learning phase they may learn movements without any information or control from others, while later they may learn under total guidance and control. When the learner does not know what controls to give to the muscles, he can learn by exploration — children learn their first movements this way. For example, during learning eye–hand coordination they randomly flail their hands and record the perceptions they get for the applied controls, associating perceptions with actions. (Piaget called this mechanism *circular reaction*). Adults may also learn by exploration, e.g. when they first move in an unknown environment, such as water or snow, especially when their limbs change

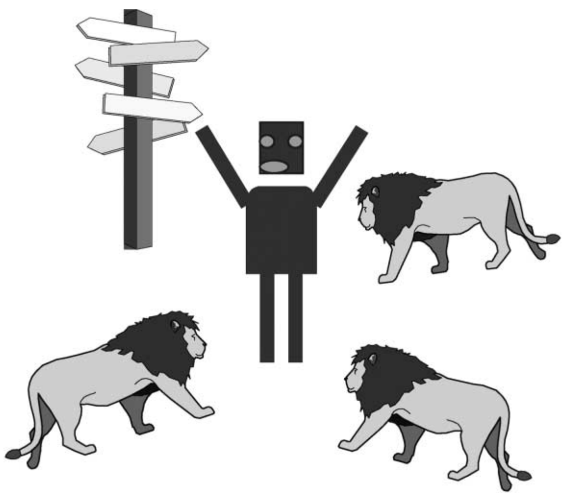

Fig. 3. Throw robot to the lions: endow the robot with a set of learning techniques and let it explore the world alone.

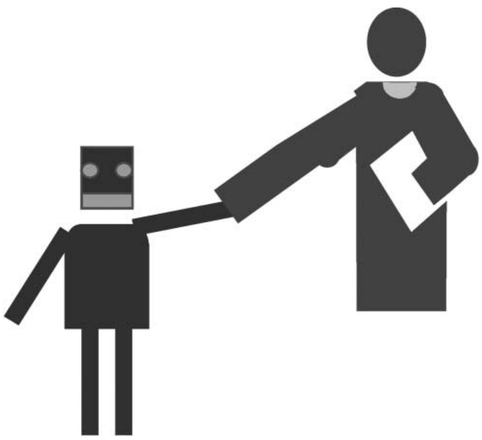

Fig. 4. Robot fostering: giving the robot a helping hand.

their shape because of wearing flippers, skates or skis, etc.

After learning the sensory-motor coordination humans can learn motor skills by imitation. Interestingly, it appears that human children are more inclined to use imitation than young chimpanzees [13]! When learning by imitation the learner observes a solution to the motion control problem, which he converts to a reference system associated with his own body, creating thus a solution that can be used directly or as a reference model. In the latter case, he tries to minimize the difference between his behavior and that of the model.

Learning can also be cognitive, in which case a trainer describes or explains the movement. This information can be used to build a controlling cognitive model or simply for guidance of the body during the movement. Once the correct movement has been achieved it is usually repeated with increasing performance, until it becomes reflex.

Robots could learn in the same ways humans do. Several authors [9,11] describe systems in which the robots learned sensory-motor control by exploration, following the circular reaction mechanism. Exploration is one way to generate examples of associations between actions and perceptions; another way to generate examples is to have a teacher guide the robot through the movement. The guidance can be done by analogic teaching [18], which is particularly useful when the precise coordinates where the robot should go are not known exactly, but the operator can see where he wants the robot to move. In most cases this is done using a teaching pendant; in other cases the human drives the robot directly. This is the case with NAVLAB, a vehicle that learned to drive on the freeway from recorded example pairs of visual scenes and the associated wheel steering commands used by the human while driving during a training session [15]. The most popular recent technique for learning sensory-motor control from examples employs neural networks (NNs); this was used in the eye–hand coordination examples mentioned [9,11] and in vision-guided mobile robots such as NAVLAB.

A simulation in which the robot starts with a cognitive phase (a descriptive knowledge of the movement) is presented in [10]. The knowledge is initially stored in a knowledge base, and the robot moves according to the description in the knowledge base. An NN gradually takes over the movement control, learning to produce the same motor control sequence, the movement becoming reflexive. But obtaining the knowledge about the movement may be non-trivial: knowledge acquisition is the acknowledged bottleneck of knowledge-based systems. Movements may be difficult to describe in words, even if simple to perform. (Once learned, *the knowledge of how to move* could be captured from the first learning robot in a generation and directly transferred/downloaded to new, similarly looking robots.)

3.2. Imitation

Humans prefer to demonstrate movements, rather then describe them linguistically. By demonstration they offer a visual model, which can be used for learning by imitation. Thus, from the perspective of learning motor skills humanoid robots have an unmatched advantage on other robots: they have a body shape that allows them to imitate humans.

The most straightforward way to force a robot to imitate human movements is to completely take control over its actions, moving it by telemanipulation. For example, NASA Johnson Space Center (JSC) has a full immersion telepresence testbed (Fig. 5), which allows operators to be virtually immersed in the environment where a two-arm dexterous anthropomorphic robot operates [6]. The operator headset allows the human to see through robot's eyes — the cameras mounted on the robot head, and special gloves allow the operator to move the robot arms, while also getting force feedback.

An extension forcing overall body imitation is possible if the body is covered with appropriately placed sensors. Imitation and capturing of elements of human movement is of great interest not only to robotics engineers but also to computer-assisted movie and game makers. For such users, Sarcos (Salt Lake City, Utah) has developed the SenSuit^{TM} (illustrated in Fig. 6), that enables real-time teleoperator control of robotic figures and computer-generated icons [16].

Early references on the use of imitation for anthropomorphic/humanoid robots (1992–1995) include [20,22–24,29]. The topic was not much in the attention of researchers, partly because the whole field of humanoids was largely non-existent outside Japan. The only notable exception was the COG project [4], which in its early phase focused more on using ideas

Fig. 5. NASA JSC full immersion telepresence testbed (after [6]).

related to the subsumption architecture and behavioral robotics — that has changed a few years later to emphasize the interaction with human users. Learning by watching was a precursor of the learning by imitation, yet the focus was on task learning and not how to move. More recently, imitation learning has received a much larger attention, the role of learning by imitation for humanoid robots being well argued in the work of Schaal and Vijayakumar [17,27], and also Mataric's group, e.g. [2].

Researchers worldwide are working on different aspects of imitation. For example, researchers at Tokyo University developed a human skull shaped robot imitating the facial expression of a human teacher. (The analysis and understanding of human gestures has recently received considerable interest from the perspective of developing a next generation of human-friendly computer interfaces.) Schaal's work uses the SARCOS humanoid robot and is likely the most advanced current work on imitation by humanoids.

Real imitation is when the robot itself watches and moves freely to reproduce human movements. This paper presents the first work in the area of learning by imitation using a real anthropomorphic/humanoid robot arm, first describing how an anthropomorphic robot arm learns to imitate the movements of a teacher (see [22] for a more detailed presentation of the subject). The robot learner and the teacher (a human or an identical anthropomorphic robot) stay next to each other and the eye of the learning robot watches to its side the movement of the teacher. Fig. 7 shows the set-up of a 2D imitation experiment detailed later in the paper.

Fig. 6. Operator in Sarcos' SenSuit controls a virtual anthropomorphic creature (copyright Sarcos, reprinted with permission).

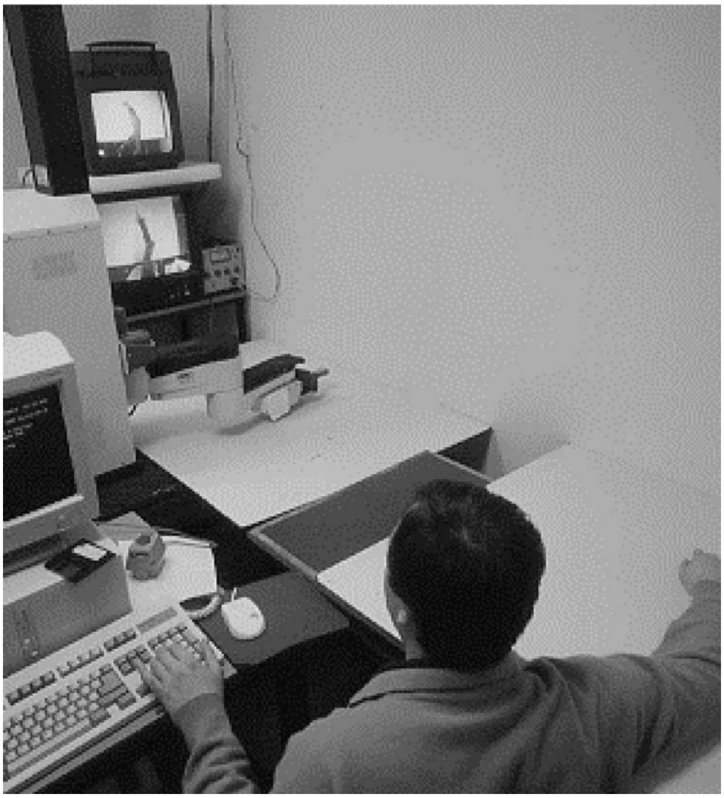

Fig. 7. Anthropomorphic robot arm imitates the movements of master arm.

Learning by imitation appears promising for making humanoid robots move like the humans, however, many other important aspects (e.g. correlating the movement with the task) need to be addressed when aiming for such an endeavor.

3.3. Collaboration/aid

Fostering expands greatly beyond providing examples for imitation. Very importantly, teachers can set-up learning experiments and provide reinforcement. They can also get directly involved in interacting physically with the robot while it is learning. They can help the robot by providing learning aids. Consider biped motion control, which must contend with the problem of maintaining stability. One can, however, alleviate the need for a very stable design by initially supporting the robot on a walker, such as a circular ring at waist level used to maintain stability, on which the robot's hands lean on, and which is pushed along when the robot walks. Force sensors would then provide feedback, and the control of the robot could be adaptively changed. Thus the robot could learn to walk while trying to minimize the force applied to the walker (an optimization problem); when finally no force is put on the walker, the robot will be able to maintain by itself a stable biped motion. This approach would lessen the possibility of an expensive robot accidentally losing stability and falling, possibly damaging itself.

A human can give a helping hand too, providing the required balance for the first steps. This can be done when teaching other movements require maintaining balance such as learning to use a bicycle. In fact this can be taken gradually too as for most children, a tricycle is the first step in learning to ride. "A tricycle has only two things to teach a child: steering and pedaling. The steering usually comes first, because the child can stand on the back step with one foot and push along with the other. Once the basic concept of steering has been learned, the child can start to use the pedals." (From "Teaching Kids to Ride", by Sheldon "Two Wheeler" Brown, http://www.sheldonbrown.com/teachride.html). Preparing the right set of learning experiences is an important part of fostering.

In unsupervised (robot) learning the teacher's primary role would be to prepare the environment. It would provide structure, order the tasks in increasing degree of difficulty, and provide tests and playground. In reinforcement learning the humans provide either direct feedback or interaction. Robots can also teach each other. *It is possible that the capacity of teaching, and not that of learning, is the decisive factor that ensures human species' superiority*.

4. Fostering by imitation

Humanoids could learn by imitation. One way would be to initially rely on teleoperation by a human. For example, for walking a teleoperator in an appropriate suite, say a "Cyber Suite" (an extension of "cyber gloves") could walk (or in another context ride the bicycle, or manipulate objects for certain task). However, the approach that is perceived as the most promising for humanoid learning is imitation. This includes learning from a present instructor or from recorded/broadcasted images of humans moving. In the following approach, fostering starts with human imitation of the robot; once the eye–arm coordination is learned the robot will imitate the human and learn from his examples.

4.1. A model for eye–arm coordination

This section describes an approach to the transfer of motor skills to such robots, in which the robot's capability to control its limbs starts with the learning of motor coordination using self-directed exploration. Once it has control over its limbs, the robot could imitate the movements of an instructor, or execute movements described verbally or as a succession of coordinates in Cartesian or joint spaces. The approach described in this section uses learning by imitation (or teaching by demonstration, as seen from human's perspective), considered promising because it is human-friendly and efficient in illustrating postures hard to capture in linguistic descriptions or quantified in programming instructions. The model of eye–arm sensory-motor coordination proposed here is characterized by a system of equations, solved numerically using NNs.

Motor skills can be broadly divided into two large categories: planning skills, i.e. the know-how expertise, and motor control skills, i.e. the ability acquired after performing a movement many times.

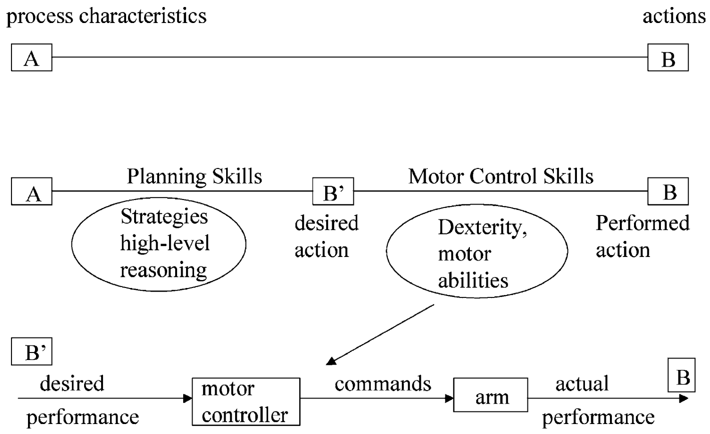

Fig. 8. Planning skills and motor control skills.

Accordingly, the mapping between process characteristics and actions can be divided in a mapping between process characteristics and desired actions (determining the planning skills), and the mapping between desired and performed actions (corresponding to motor control skills), as in Fig. 8.

The former are related to strategies at a higher level, while the latter refer to dexterity and motor abilities. In the case of arm coordination, the mapping between desired and actual performance is subject to a representation in which a motor controller maps the desired performance into commands, and the arm plays the role of the controlled plant, mapping commands to actual performance. The motor controller also performs a transformation from a sensory coordinate system to a motor coordinate system.

In order to be able to place its arm in a desired position, the robot needs to have a model of its eye–arm coordination. Traditionally, visuo-motor coordination in robotics addressed eye–hand coordination. For redundant manipulators (including here the human arm and anthropomorphic robot arms), the associated inverse kinematics problem is under constrained, admitting more than one solution. In the context of acquiring motor skills by imitation, when the task requires specific postures, or imposed by obstacles in the environment, eye–hand coordination is insufficient. This is illustrated in the 3D situation in Fig. 9, where posture (1) given by an eye–hand model is unacceptable due to an obstacle, while a posture (2) shown by an instructor provides a feasible alternative.

Eye–arm coordination adds to other models of coordination as shown in Fig. 10.

The following presents a model of eye–arm coordination model schematically illustrated in Fig. 11. During the learning of the visuo-motor model the visual inputs could be from the robot's own arm, from the arm to follow, or from another teaching arm. During the imitation of human arm movements, the visual inputs are images of the human arm. The model (W) can be considered to reflect the mapping between visual inputs (X) and joint motor commands (Y).

The inputs (X) to the model are low-resolution images originating in the images obtained from video cameras. The first experiment uses a 2-link arm. The

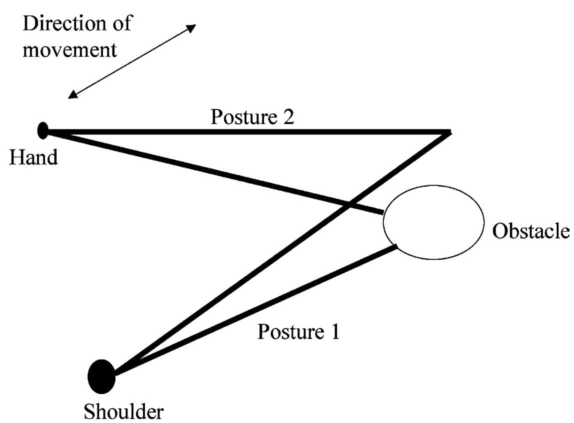

Fig. 9. Two-arm postures for the same hand position.

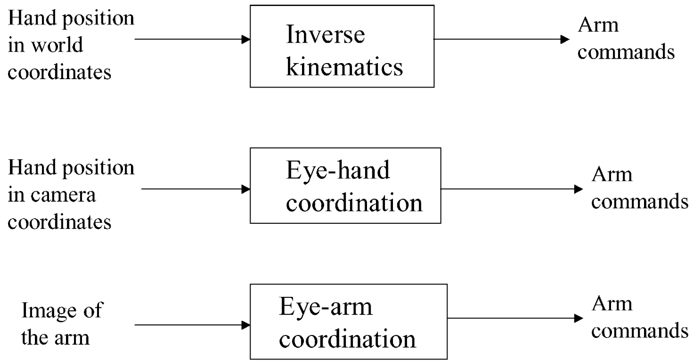

Fig. 10. Models of arm coordination.

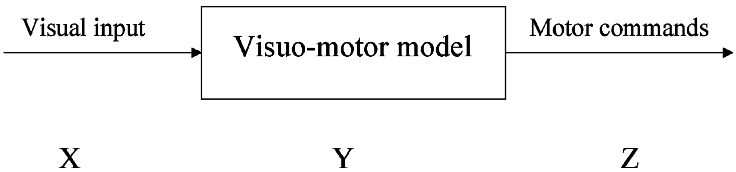

Fig. 11. A model of visuo-motor coordination.

output (Y) is associated with shoulder and elbow joint angles as in Fig. 12.

Model identification from training examples consists in finding W, for given X and Y pairs. In order to identify a model it is necessary to obtain input–output data characterizing it. In this case, one needs to associate visual inputs to motor control outputs, which would position the arm in a posture similar to the visual input (showing the teacher arm). In most problems the associations are between actions and determined perceptions through the same system. Here the robot must give controls to the own arm to place it like what it sees for the teacher's arm. To surpass this problem, in the technique adopted here for collecting training examples, *the human (teacher) imitates the robot*. The robot randomly flails its arm, and for each position of the arm, the human places his arm in a similar posture giving also a validation signal. Thus, the robot receives the information how the human arm looks like when it is in a posture similar to that of his arm resulted as an effect of controls Y. Whenever will need to achieve a posture like X the robot will have to provide the commands Y.

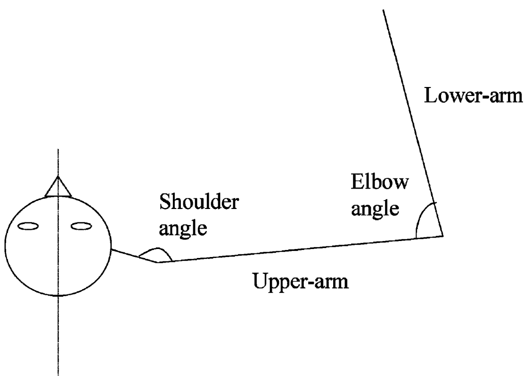

Fig. 12. Arm skeleton showing shoulder and elbow angles.

A first type of experiment involved a robot imitating human arm movements performed in a horizontal plane. In its horizontal performance, the robot is anthropomorphic (see Fig. 7). The second type of experiments targeted the extension of this approach to 3D performance.

Two identical looking robots (RTX) were used, one learning to imitate the other (shown before in Fig. 4). The human operator controlled the teacher robot via a computer. This time the camera was placed at the

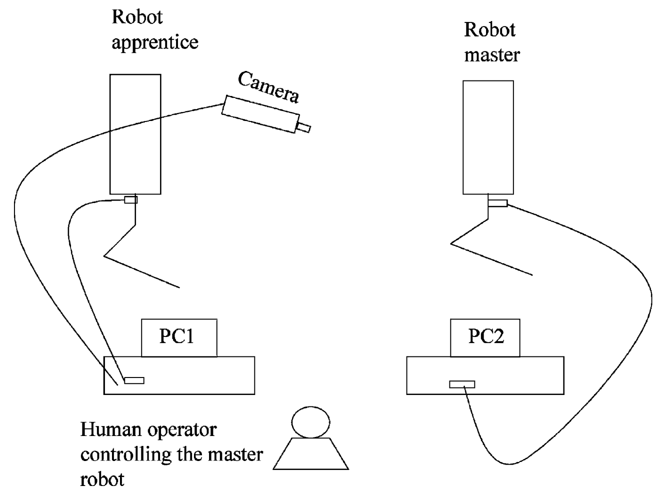

Fig. 13. Set-up for 3D learning.

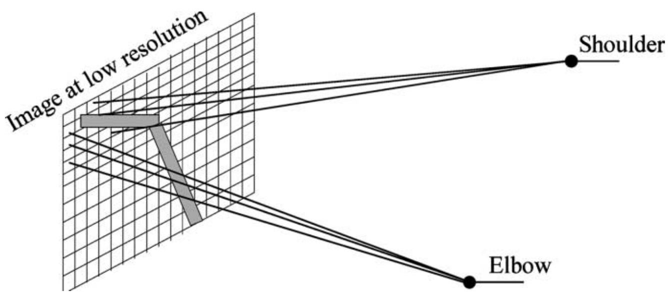

Fig. 14. Shoulder and elbow neurons that map images to joint commands.

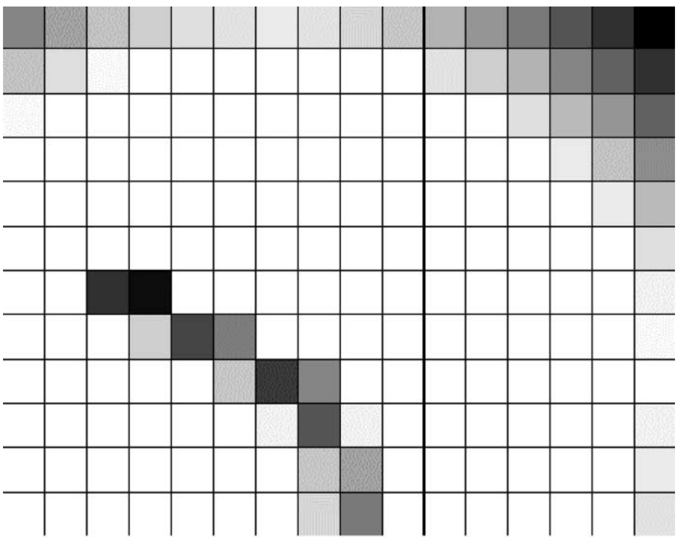

Fig. 15. Low-resolution image.

approximate position of the human eye, gazing at an oblique angle to the teacher arm, as illustrated in the drawing in Fig. 13. The image-command pairs were selected to (approximately) uniformly cover the workspace. A total of 97 image-command pairs was collected, and separated in a training set (88 pairs) and a test set (nine pairs), the number of pairs in the test being about 10% of the training set. Both addressed models (a neural and fuzzy-neural one) used only one neuron per joint (one neuron for the shoulder and one neuron for the elbow, for the 2D case in Fig. 14).

The inputs (X) came from a 192 pixels (12×16) low-resolution image (Fig. 15), obtained by averaging regions of 16×16 neighboring pixels of a higher resolution image obtained from the frame-grabber. Their intensity values were in the [0,1] interval, with 256 gray levels. Similarly, the outputs Y were normalized to [0,1], which was the required definition domain for the fuzzy-neural model.

Details of the implementation can be found in [21,22]. The performance of the neural model (evaluated on the test set and on the quality of imitation in a performance illustrated in the following) was considered good.

4.2. Robot imitates the human and learns arm movement

The robot used the neural models determined by training to imitate (track) the movements of the teacher arm. The qualitative evaluation consisted of subjective assessments of the closeness of the posture of the robot arm to the posture of the human master arm. A series of images during imitation is shown in Fig. 16. When the training set included data from several different looking human arms the model generalized and become robust to variation in the appearance of the teacher arm.

In another set of experiments the robot apprentice as illustrated in Fig. 17 imitated 3D movements of a robot master arm.

4.3. Future work

The work described in this section is only a first attempt to learn motor skills by imitation. The technique proposed here for obtaining training examples is general and can be applied to learning other types

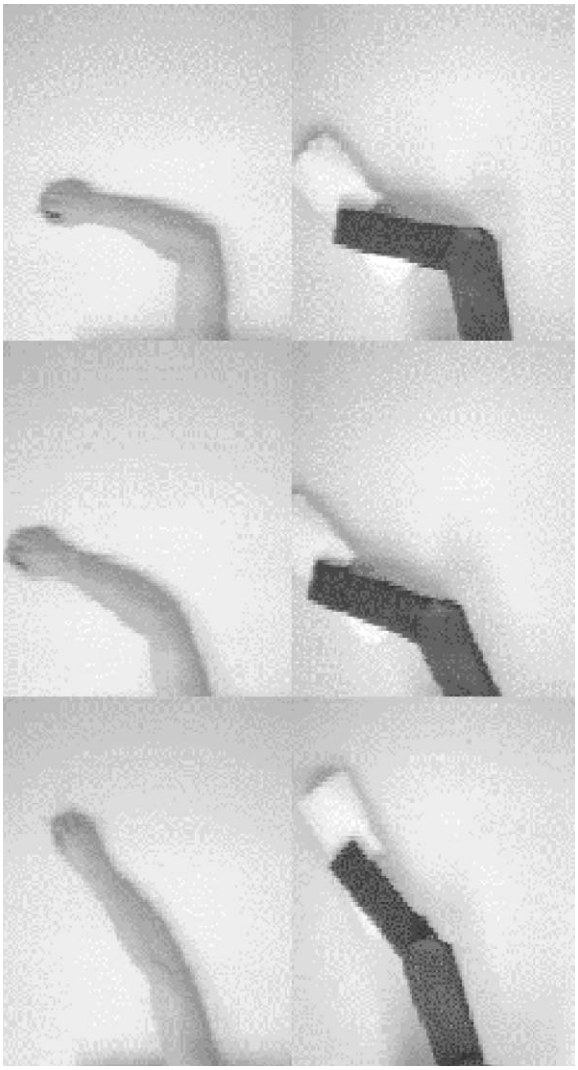

Fig. 16. Images showing human arm and imitation by robot arm.

of movement. The models employed here are simple and of limited power. The neural models used in this section require that the vision system always see the teacher's shoulder at the bottom of the image. To obtain a robust system tolerant to position and rotation variations, one could expand the described system by introducing pre-processing models that perform appropriate compensating image transformation. A similar pre-processing is also needed to insure scale invariance, etc. When imitation becomes possible in real-world environments, the issue of correlating the

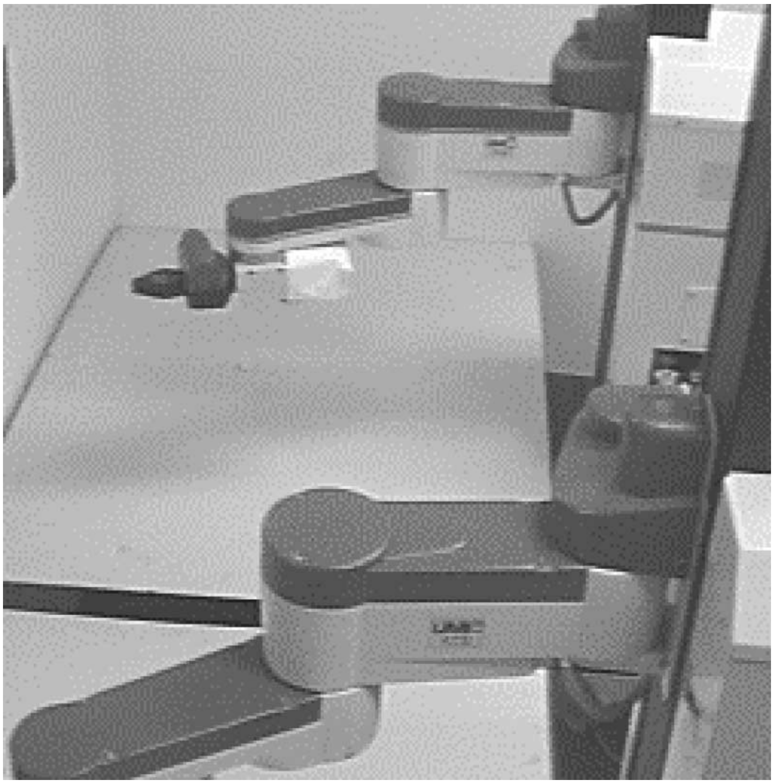

Fig. 17. Master arm and slave arm: imitation in 3D.

Fig. 18. Toward 3D imitation of arm movement.

motor behavior to the task to which it relates, needs to be addressed. For example, a robot may imitate quite well human arm movement in hammering a nail, with the small exception of hitting one millimeter away from the nail, or hitting the nail at a low speed: it is necessary to have some *understanding* of the purpose of the movement. Imitation has the important role of providing a rough example of a movement, however, to enable task-related motor skill learning one needs more sophisticated models than those for simple perceptual skill addressed here.

Future experiments will benefit from a 7-DOF anthropomorphic arm and a stereovision system, illustrated in Fig. 18. The experiments will target learning human-like 3D movement from imitation of an arm performing unconstrained in the environment.

5. Summary and conclusion

The paper argued in favor of developmental robotics: perfecting techniques inspired from human motor and cognitive development for forming/educating robots.

The focus technique explored here is fostering: a role that the teacher assumes to facilitate robot learning. The teacher can prepare experimental environments conditioning data for unsupervised learning, provide reinforcement depending on how the robot does during learning, interact (including physically) with the robot helping the robot perform the task, showing the robot how to do the task.

An example was given to illustrate a fostering role: the teacher initially imitated a robot in its action; later it acted as a model and the robots imitated him, thus getting examples of how to do things. The example showed the learning of eye–arm coordination and then the learning of arm movements by imitation. In another example, the teacher's role was played by another robot: it appears useful for the future to have the robots teach and foster themselves.

Acknowledgements

This work was performed at Jet Propulsion Laboratory, California Institute of Technology, and was sponsored by the National Aeronautics and Space Administration. The learning by imitation experiment was performed at Victoria University, Melbourne, Australia.

References

[1] M. Asada, K.F. MacDorman, H. Ishiguro, Y. Kuniyoshi, Cognitive developmental robotics as a new paradigm for the design of humanoid robots, in: Proceedings of the First IEEE-RAS Conference on Humanoid Robotics, Cambridge, MA, September 7–8, 2000.
[2] A. Billard, M. Mataric, Learning human arm movements by imitation: Evaluation of biologically inspired connectionist architecture, in: Proceedings of the First IEEE-RAS Conference on Humanoid Robotics, Cambridge, MA, September 7–8, 2000.
[3] J.S. Bruner, Nature and uses of immaturity, American Psychologist 27 (8) (1972) 687–709.
[4] COG (MIT), www.ai.mit.edu/projects/cog.
[5] P. Dario, E. Guglielmelli, V. Genovese, M. Toro, Robot assistants: Applications and evolution, Robotics and Autonomous Systems 18 (1996) 225–234.
[6] DART, http://tommy.jsc.nasa.gov/~li/dartfitt.html.
[7] J. Engelberger, Service robots arise, International Federation of Robotics Newsletter 18 (1) 1995.
[8] First IEEE Conference on Humanoid Robots, Cambridge, MA, September 7–8, 2000.
[9] M. Kuperstein, INFANT neural controller for adaptive sensory-motor coordination, Neural Networks 4 (1991) 131–145.
[10] S.H. Lane, D.A. Handelman, J.J. Gelfand, Can robots learn like people do? in: S.K. Rogers (Ed.), Applications of Artificial Neural Networks, Proc. SPIE, Vol. 1294, SPIE, Bellingham, WA, 1990, pp. 296–309.
[11] B. Mel, A connectionist model may shed light on neural mechanisms for visually guided reaching, Journal of Cognitive Neuroscience 3 (3) (1991) 274–292.
[12] P.R. Merlyn, Toward a Humanoid Robot: Artificial Intelligence and the Confluence of Technologies, SRI Consulting, Business Intelligence Program, Report D96-2031, October 1996.
[13] K. Nagell, R.S. Olguin, M. Tomasello, Processes of social learning in the tool use of chimpanzees (pan troglodytes) and human children (homo sapiens), Journal of Comparative Psychology 107 (2) (1993) 174–186.
[14] J. Piaget, Play, Dreams and Imitation in Childhood, International Library of Psychology, 1999.
[15] D.A. Pomerleau, Neural Network Perception for Mobile Robot Navigation, Kluwer Academic Publishers, Boston, MA, 1993.
[16] SARCOS, http://www.sarcos.com.
[17] S. Schaal, Is imitation learning the route to humanoid robots? Trends in Cognitive Science 3 (6) (1999) 233–242.
[18] T.B. Sheridan, Telerobotics, Automation, and Human Supervisory Control, MIT Press, Cambridge, MA, 1992.

[19] A. Stoica, Anthropomorphic Systems, JPL Report D-14842, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 1997.
[20] A. Stoica, Evolving creatures that can learn by imitation: Apprentice behavior and its role in robot motor learning, in: Proceedings of the International Conference on Autonomous Robots and Artificial Life (PERAC'94): From Perception to Action, EPFL, Lausanne, Switzerland, September 7–9, 1994, pp. 440–443.
[21] A. Stoica, Learning eye–arm coordination using neural and fuzzy neural techniques, in: H.N. Teodorescu, A. Kandel, L. Jain (Eds.), Soft Computing in Human-related Sciences, CRC Press, Boca Raton, LA, 1999, pp. 31–61.
[22] A. Stoica, Motion learning by robot apprentices: A fuzzy neural approach, Ph.D. Thesis, Victoria University, Melbourne, Australia, 1995.
[23] A. Stoica, Neurocontrol of visually guided anthropomorphic manipulators, in: Proceedings of the First International US–Mexico Conference on Neural Networks and Neurocontrol, Quintana Roo, Mexico, September 5–15, 1995, pp. 392–406.
[24] A. Stoica, On learning arm movements in behavior-based robots: Towards the integration of visual and linguistic information in learning from humans, in: Proceedings of the Machine Learning and Hybrid Systems Workshop, Sixth Australian Joint Conference on Artificial Intelligence, Melbourne, Australia, November 16 (1993) 20–24.
[25] A. Stoica, Robot fostering techniques for sensory-motor development of humanoid robots, in: Proceedings of the First IEEE-RAS Conference on Humanoid Robots, Cambridge, MA, September 7–8, 2000.
[26] J.A. Adam, Technology combats disabilities, IEEE Spectrum 31 (10) (1994) 24–26.
[27] S. Vijayakumar, S. Schaal, Real-time learning in humanoids: A challenge for scalability of online algorithms, in: Proceedings of the First IEEE-RAS Conference on Humanoid Robotics, Cambridge, MA, September 7–8, 2000.
[28] J.J. Wagner, H.F.M. van der Loos, L.J. Leifer, Construction of social relationships between users and robot, Robotics and Autonomous Systems 31 (2000) 185–191.
[29] M. Wingate, A. Stoica, A step towards robot apprentices — visual learning, in: Proceedings of the IASTED International Conference on Robotics and Manufacturing, Honolulu, I, August 19–22, 1996.

Adrian Stoica is a Senior Researcher in the Advanced Computing Technologies Group at the Jet Propulsion Laboratory (JPL), California Institute of Technology, Pasadena, California. He received his M.S. degree in Electrical Engineering from the Technical University of Iasi, Romania in 1986, and his Ph.D. in Electrical Engineering and Computer Science from Victoria University of Technology, Melbourne, Australia in 1996. His research is directed along two themes: adaptive hardware for autonomous space systems, and next-generation robots. Current projects are in the areas of evolvable hardware, adaptive and configurable computing, smart sensing, soft computing, spacecraft survivability, robot learning, and education of anthropomorphic robots. He is the initiator of the NASA/DOD Workshops on Evolvable Hardware. He received the 1999 Lew Allen Award, the highest distinction for excellence in research at JPL.