ELSEVIER

Robotics and Autonomous Systems 941 (2001) 1–16

Robotics and Autonomous Systems

www.elsevier.com/locate/robot

Learning human arm movements by imitation: evaluation of a biologically inspired connectionist architecture

Aude Billard [a,b,*], Maja J. Matarić [a]

[a] *Computer Science Department, University of Southern California, USA*

[b] *Kawato Dynamic Brain Project (ERATO/JST), ATR Human Information Processing Research Labs, 2-2-Hikaridai, Seika-cho, Soraku-gun, Kyoto 619-0288, Japan*

Abstract

This paper evaluates a model of human imitation of abstract, two-arm movements. The model consists of a hierarchy of artificial neural networks, which are abstractions of brain regions involved in visuo-motor control. The model is validated in a biomechanical simulation of a 37 degrees of freedom (DOF) humanoid. Input to the model are data from human arm movements recorded using video and marker-based tracking systems. Results show a high qualitative and quantitative agreement with human data. The model's reproduction is better or comparable to that of human subjects imitating the same movements.

Keywords: Imitation; Learning; Artificial neural networks

1. Introduction

A goal of robotics is to have robots become a part of human everyday lives. A key challenge to make this possible is developing flexible motor skills in order to give robots the ability to be programmed and interacted with more easily and naturally, and to assist humans in various tasks. A very exciting area of current research is concerned with developing human-like robots (humanoids) for assisting humans in medical surgery [34,36] and rehabilitation [6], for providing help in everyday tasks to the elderly and the disabled [58], and for replacing humans in low-level industrial tasks and unsafe areas [25,30] (including space, nuclear, and waste management industries).

Providing robots with human-like capabilities, and in particular, with sophisticated motor skills for flexible and precise motions is a very difficult task, requiring important low-level programming (with high cost) for fine tuning of the motor parameters and re-calibration of sensor processing [18,47]. An alternative is to provide the robot with *learning* or *adaptive* capabilities, which can be used for on- and/or off-line optimization of predefined motor control parameters [13,28,55]. Particularly challenging is the problem of how to teach a robot new motor skills through demonstration rather than through reprogramming. In such a scenario, the robot learns novel motor sequences by replicating those demonstrated by a human instructor and by tuning its motor program descriptions so as to successfully achieve the task. The method is interesting because it allows the robot to be programmed and interacted with merely by human demonstration, a natural and simple means of human–machine interface. Furthermore, it makes the robot flexible

* Corresponding author. Present address: 3614 Wyatt Way, Los Angeles 90089-2520, CA, USA. Tel.: +1-213-740-92-23; fax: +1-213-740-56-87.
E-mail address: billard@usc.edu (A. Billard).

0921-8890/01/$ – see front matter © 2001 Published by Elsevier Science B.V.
PII: S0921-8890(01)00155-5

with respect to the tasks it can be taught and, thus, facilitates the end-use of robotic systems.

1.1. Related work

The first robotics work to address imitation was focused on assembly task-learning from observation. Typically, a series of arm trajectories of a human, performing object moving/stacking tasks, were recorded either using a manipulandum, with the advantage of measuring directly the joint torques [4,14,27], or using video images [23,32,53]. Data were analyzed to remove inconsistencies and extract key features of movement. An industrial non-human-like robotic arm would then be trained to reproduce the trajectory which maximizes the data key features. These efforts constitute a significant body of research in robotics, and contribute to data segmentation and understanding. However, they provide highly task-specific solutions, with little flexibility for applying the same algorithm to imitation after different types of movements and tasks.

More recent efforts, including our own [3,8,10,37], have been oriented toward analyzing the underlying mechanisms of imitation in natural systems and modeling those on artificial ones. Atkeson and Schaal [5,49] developed a control strategy in which the robot learns a reward function from the demonstration and a task model from repeated attempts to perform the task. The algorithm has proven to be robust, fast and applicable to different tasks, such as juggling and pole balancing. In a more biological approach, Demiris and co-workers [15,16] performed experiments in which a robotic head equipped with a pair of cameras observes and imitates the head movements of a human demonstrator. These approaches use visual feature detectors, which inform a built-in system that directly mapped a set of possible observed head movements to the robot's own head movements. The inspiration for the visual feature detectors comes from evidence in monkey of neurons specialized to particular orientation of motion [44] and the observed-performed mapping is based on Meltzoff's proposed innate visuo-motor map [39]. Following a similar research line, Kuniyoshi and co-workers achieved fine oculo-motor control of a robot head for on-line tracking [7,32] and reproduction [12] of human torso motion by a humanoid robot. Schaal and Sternad [50,51] explored the idea of creating complex human-like movements from biologically motivated movement primitives. Each degree of freedom (DOF) of a robot's limb is assumed to have two independent abilities to create movement, one through a discrete dynamic system (for point-to-point movements), and one through a rhythmic system (type of *central pattern generator* (CPG) [54]). The model was implemented on a humanoid robot for a drumming task. Jenkins et al. [26] described an imitation model based on a set of perceptuo-motor primitives. A simple version of the model was validated on a 20 DOF humanoid simulation with dynamics, using real vision data (same as those used in this work) to imitate a movements from athletics and dance. Fod et al. [19] contributed to this model by providing a method for automatically extracting a set of primitives from human movement data.

1.2. Our approach

Our work aims to complement the above approaches, by investigating a connectionist-based model validated on a biomechanical simulation of a humanoid. The endeavor is to, on the one hand, build biologically plausible models of animal imitative abilities, and, on the other hand, to develop architectures for visuo-motor control and learning in robots which would show some of the flexibility of natural systems. We follow neuroscience studies of primate motion recognition and motor control. Specifically, our work is driven by the observation that: (1) visual recognition of movements is done in both extrinsic and intrinsic frames of reference [42,56]; (2) a neural system, possibly the *mirror neuron system*, encapsulates a high-level representation of movements, the link between visual and motor representation [17,46]; (3) motor control and learning are hierarchical and modulate (evolutionary) primitive motor programs (e.g. CPGs, located in primate spinal cord [54]).

Our model is composed of a hierarchy of artificial neural networks and gives an abstract and high-level representation of the neurological structure underlying primates brain's visuo-motor pathways. These are the spinal cord, the primary and pre-motor cortices (M1 and PM), the cerebellum and the temporal cortex (TC). The model has first been evaluated in a pair of demonstrator–imitator humanoid avatars with 65 DOFs [8] for learning by imitation gestures and complex movements involving all the avatar's limbs. In

this paper, we evaluate the model's performance at reproducing human arm movements. A biomechanical simulation is developed which models the muscles and the complete dynamics of a 37 DOF humanoid. [1] The aim of these experiments is to evaluate the realism of the model and the dynamic simulation at modeling human imitation.

In the experiments presented here, only 11 DOFs are actively commanded to match the observed performance (4 DOFs per arm and 3 for the torso), while the rest of the joints are kept immobile. In the experiments reported in [8,10], we demonstrated the validity of the architecture for controlling the 65 DOFs of our avatar for imitating complex movements requiring all limbs. There, data for the imitation were simulated, produced by a demonstrator avatar, and we could generate data for the whole body. In this paper, we use human data. However, because of the limitation of our tracking system, we could not record motion of the whole body and were constrained to using movements of the upper torso only. In future work, we will use a full body tracking system which will allow us to further validate the model for controlling the whole 37 DOFs on real data (as opposed to simulation data we have used previously). Preliminary work in this direction is reported in [11].

The rest of the paper is organized as follows. In Section 2, we describe in detail the model, and, in particular, the visual processing of the data and the learning algorithm. In Section 3, we evaluate the model's performance on a series of experiments for reproducing human arm motion, namely reaching movements and oscillatory movements of the two arms. We compare the model's performance to that of humans in the same imitation task. Section 4 concludes this paper with a short summary of the presented work.

2. The model

We have developed a highly simplified model of primate imitative ability [8] (see Fig. 1). This model is biologically inspired in its function, as its composite modules have functionalities similar to that of specific brain regions, and in its structure, as the modules are composed of artificial neural architectures (see Fig. 2). It is loosely based on neurological findings in primates and incorporates abstract models of some brain areas involved in visuo-motor control, namely the TC, the spinal cord, the primary motor cortex (M1), the premotor (PM) area and the cerebellum.

2.1. Brief description of the modules

Visual information is processed in TC for recognition of the direction and orientation of movement of the demonstrator's limbs relative to a frame of reference located on the demonstrator's body, i.e., the TC module takes as input the Cartesian coordinates of each joint of the demonstrator's limbs in an excentric frame of reference (whose origin is fixed relative to the visual tracking system). It then transforms these coordinates to a new set of coordinates relative to an egocentric frame of reference. Our assumption of the existence of orientation-sensitive cells in an egocentric frame of reference in TC is supported by neurological evidence in monkeys [42,43] and humans [2,31,56]. The vision system also incorporates a simplified attentional mechanism which is triggered whenever a significant change of position (relative to the position at the previous time step) in one of the limbs is observed. At this stage of the modeling and given the simplicity of this module, the attentional module does not relate to any specific brain area. The attentional mechanism creates an inhibition, preventing information flow from M1 to PM and further to the cerebellum, therefore, allowing learning of new movements only when a change in the limb position is observed. In Section 2.2, we describe the motion tracking system we used in the experiments and explain in more detail the stages of visual processing in the TC module.

Motor control in our model is hierarchical with, at the lowest level, the spinal cord module, composed of primary neural circuits (CPGs [54]), made of *motor neurons* and *interneurons* [2] (see Section 2.3). The motor neurons in our simulation activate the muscles of the humanoid avatar, see Section 2.5. The M1 module monitors the activation of the spinal networks. Nodes

[1] The previous implementation of the model used a partial dynamic simulation of a 65 DOF humanoid avatar, where we did not compute the internal torques of the humanoid.

[2] Inter- and motor neurons are spinal cord neurons with no direct and direct input to the muscles, respectively.

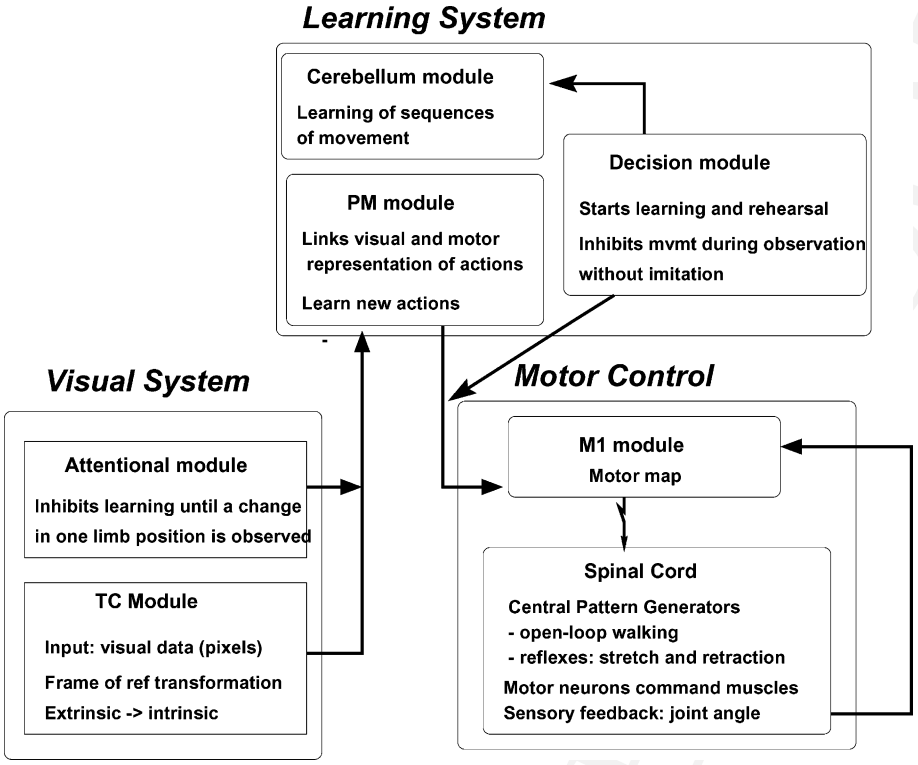

Fig. 1. The architecture consists of seven modules which give an abstract and high-level representation of corresponding brain areas involved in visuo-motor processing. The seven modules are: the attentional and TC modules, the primary motor cortex and spinal cord modules, the PM cortex and cerebellum module, and the decision module.

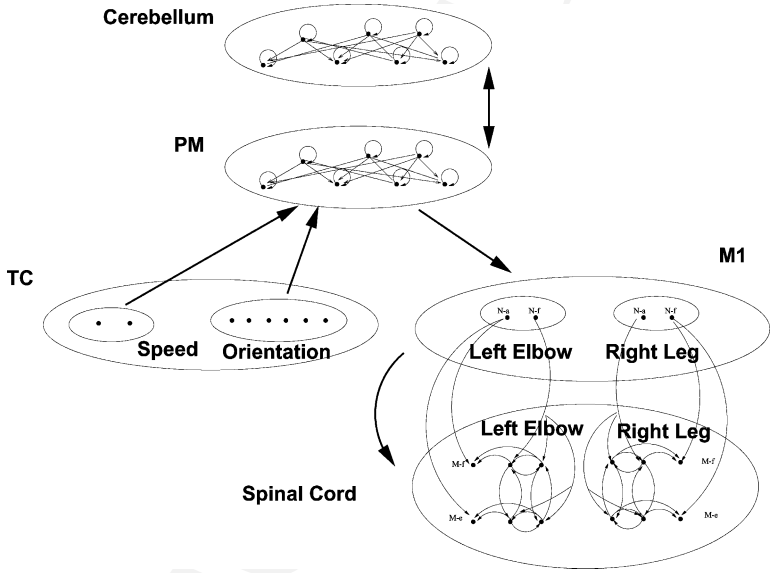

Fig. 2. A schematic of the interconnections between the modules, and the neural structure within each module.

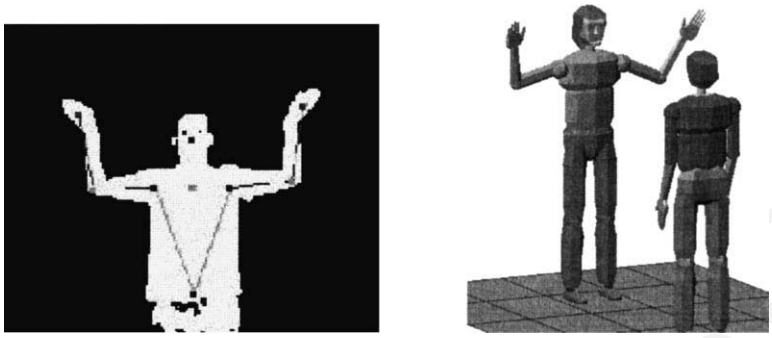

Fig. 3. Motion tracking system of human movement (left); the Cosimir simulator (right).

in M1 are distributed following a topographic map of the body.

Learning of movements is done in the PM and cerebellum modules. These modules are implemented using the Dynamical Recurrent Associative Memory Architecture (DRAMA) [9] which allows learning of time series and of spatio-temporal invariance in multi-modal inputs (see Section 2.4 for details). Finally, the decision module controls the transition between observing and reproducing the motor sequences, i.e., it inhibits PM neural activity due to TC (visual) input to flow downwards to M1 (for motor activation). It is implemented as a set of if–then rules and has no direct biological inspiration.

Neurons in the PM module respond to both visual information (from the TC) and to corresponding motor commands produced by the cerebellum. As such, they give an abstract representation of *mirror neurons*. Mirror neurons refer to neurons located in the rostral part of inferior PM area 6 in monkey [17,46], which have been shown to fire both when the monkey grasps an object and when it observes another monkey or a human performing a similar grasp.

In the next section, we describe in more detail the visual, motor, and learning parts of our model.

2.2. Visual segmentation

Data for our experiments (see Section 3) are recordings of human motion. The first set of data was recorded using a vision-based motion-tracking system. The system we used is capable of selecting a collection of features from the moving image, based on a constrained (unoccluded and unambiguous) initial position and kinematic model of a generic adult human (see [57] for a detailed description). Tracking is done off-line and based on image frequency of 15 Hz. The system allows tracking of the upper body in the vertical plane, where the body features correspond to those of a stick figure (see Fig. 3). It calculates the positions (relative to a fixed, excentric frame of reference) of nine points on the body: two located on the wrists, two on the elbows, two on the shoulders, one on the lower torso, one on the neck and one on the head.

A second set of human arm data, used in the experiments, was gathered by Matarić and Pomplun in a joint interdisciplinary project conducted at the National Institutes of Health Resource for the Study of Neural Models of Behavior, at the University of Rochester [38,45]. Subjects watched and imitated short videos of arm movements, while wearing the FastTrak marker mechanism for recording the positions of four markers on the arm: at the upper arm, near the elbow, the wrist, and the hand.

In the experiments, these Cartesian coordinates are input to the TC module of our model, in which they are processed in four stages. Data are first transferred into a frame of reference relative to the demonstrator's body, by calculating the joint angles of the elbows and shoulders. In a second stage, a low-pass filter is applied to the calculation of the angular velocity for each of the four joints. This stage corresponds to the attentional mechanism of Fig. 1. This allows us to eliminate small arm movements which we consider noise for these experiments. These small motions are due to two factors: (1) the locations of the nine points of reference of the tracking are imprecise; the coordinates

Table 1
Thresholds (°) for visual filtering[a]

Experiment	θ_0	T_0
LS_x	PI/16	15
LS_y	PI/16	15
RS_x	PI/16	15
RS_y	PI/16	15
LE	PI/8	10
RE	PI/8	10

[a] LS_x is the DOF_x of the left shoulder. LE is the left elbow. θ_0 (in radians) is the minimum displacement for detecting a motion. T_0 (in recording cycles) is the minimum time delay during which no displacement greater than θ_0 has been observed.

are extrapolated across three time steps of recording; (2) because of the interaction torques across the body, movement of one limb results in small motions of the rest of the body. These small movements are noise to us, as we wish to recognize only voluntary movements (as opposed to movements made to compensate for the interaction torques). Since shoulders and elbows have different dynamics, due to their different lengths and muscular composition, we applied different filter parameters to each. The filtering process depends on a set of two parameters per DOF. They are: (1) the minimum displacement θ_0 (in joint angle) for detecting a motion; (2) the minimum time delay T_0 during which no displacement greater than θ_0 has been observed. The latter is then considered as a stop of the motion or small, noisy movements. Table 1 shows the values we used for the experiments reported in Section 3. Note that in the experiments, we used at most 2 (abduction and flexion) of the 3 DOFs of the shoulders, as the third DOF, humeral rotation, was not recorded by either of the two tracking systems. Fig. 6 shows the results of the visual segmentation for three oscillatory movements of the two arms. Only the large movements are segmented.

In the third stage, we calculate the direction of movement of each limb relative to the limb to which it is attached (elbow relative to shoulder and shoulder relative to the torso). The direction of movement is positive or negative depending on whether the limb moves upwards or downwards, respectively. In the fourth stage, the TC module activates a series of cells coding for the possible joint angle distributions. There are two cells per DOF per joint, coding for positive and negative direction of movement, respectively. The output of the cells encodes both the direction and speed of the movement. The faster the speed, the greater the output excitation of the cell. Only one cell of the pair is active at a time. If both cells are inactive, the limb is not moving. The decomposition of the limb motion can easily be mapped to the muscular structure of the imitator; each DOF of a limb is directed by a pair of flexor–extensor muscles. Upward and downward directions of movement correspond to the activation of the extensor and flexor muscles, respectively.

In summary, the visual module performs four levels of processing on the data: (1) a transformation from extrinsic to intrinsic frame of reference; (2) filtering of small and noisy motions; (3) a parameterization of the movements in terms of speed and direction; (4) segmentation of the motion, based on changes in velocity and movement direction.

2.3. Motor control

2.3.1. Spinal cord module

In our model, motor control is hierarchical. On the lowest level of motor control is the *spinal cord* module. It is composed of primary neural circuits made of *motor neurons* (afferent to the muscles and responsible for the muscle activation or inhibition) and *interneurons*.

In our experiments, the spinal circuits are built-in and encode extending and retracting arm movements, as well as rhythmic movements of legs and arms involved in locomotion, following a biological model of the walking neural circuits in vertebrates [24]. The neurons of the spinal cord module are modeled as leaky-integrators, which compute the average firing frequency [22]. According to this model, the mean membrane potential m_i of a neuron N_i is governed by the equation

$$\tau_i \frac{\mathrm{d}m_i}{\mathrm{d}t} = -m_i + \sum w_{i,j} x_j \quad (1)$$

where $x_j = (1 + \mathrm{e}^{(m_j + b_j)})^{-1}$ represents the neuron's short-term average firing frequency, b_j the neuron's bias, τ_i a time constant associated with the passive properties of the neuron's membrane, and $w_{i,j}$ the synaptic weight of a connection from neuron N_j to neuron N_i.

2.3.2. Motor cortex module: M1

The primary motor cortex (M1) module contains a *motor map* of the body (similar to the corresponding brain area [41]). It is divided into layers of three neuron networks, each activating distinct (extensor–flexor) muscle pairs (see Fig. 2). The three-neuron network allows for independently regulating the amplitude (two nodes, one for each muscle) and the frequency (one node) of the oscillation of the corresponding flexor–extensor pair, similar to [24]. An oscillation of a limb segment is generated by activating all three neurons, allowing a small time delay between activation of the first and second neuron, thus creating an asymmetry between the two motor neurons' activity and the corresponding muscle contraction. Motion of a single muscle (flexor or extensor) is obtained by activating only one of the two amplitude nodes, while keeping the frequency node at zero. The speed of the movement, i.e., the speed of contraction of the muscle, is controlled by increasing the output value of the amplitude neuron and consequently that of the corresponding motor neuron in the spinal cord. The amplitude of the movement (in the case of one-muscle activation) is controlled by the duration of the neuron activation. The longer the activation of the amplitude neuron (and subsequently of the motor neuron), the longer the duration of muscle contraction, the larger the movement.

M1 receives sensory feedback, in the form of joint angle position, from the spinal cord module. Each motor area of M1 receives sensory feedback from its related sensory area (arm area receives feedback on joint positions of the shoulder joints). This is used to modulate the amplitude or speed of the movement, by increasing or decreasing (for smaller or larger speed) the output of the M1 nodes. The sensory feedback provides inhibition; the larger the feedback, the slower the movement. In the experiments of Section 3.1, this is used to modulate reaching movements. When the movement starts, the sensory feedback is at its minimum and consequently the tonic input (i.e., the amplitude of the M1 nodes' output) is at its maximum. When the arm has reached half the required distance, the sensory feedback is at its maximum and, consequently, the tonic input is decreased to 10% of its maximum. The arm stops shortly afterwards when the torque produced by the muscle (proportional to the motor neuron's output, see Section 2.5) equals that of gravity.

2.3.3. PM cortex module

The PM module creates a direct mapping between the parameterization of the observed movement in TC, following visual segmentation, and that used for motor control in M1. In TC, the observed motion is segmented in terms of speed, direction and duration of movement (the delay between two changes in velocity and motion direction) of each limb (see Section 2.2). In M1, speed and direction of movement of each limb CPG (in the spinal cord) are controlled by the amplitude of the nodes which project to the relevant interneurons. PM nodes transfer the activity of the TC nodes (observation of a specific movement) into an activity pattern of M1 nodes (motor command for the corresponding movement). A large output activity in TC cells (comprised between 0 and 1) will lead to an important output from PC nodes, and further from M1 nodes which further the activation of the corresponding amplitude node. Duration of movement is proportional to the duration of activation of the amplitude node. Learning of the movements consists, then, of storing the sequential activation (recording the amplitude and the time delay) of each of the TC nodes, and mapping these to the corresponding M1 nodes. This will be further explained in Section 2.4.

2.3.4. Decision module

Finally, the execution of a movement (as during rehearsal of the motion in the experiments, see Section 3) is started by the decision module, by activating one of the cerebellum nodes (the node which encodes the corresponding sequence of muscle activation, described in Section 2.4). The activity of the cerebellum node is passed down to the nodes of the PM cortex, which encode co-activation of the muscle in a specific step of the sequence (described in Section 2.4), and, further, down to the nodes of the second layer of primary motor cortex (M1). Finally, the activity of the nodes in the second layer of M1 activates the nodes in the spinal cord module, which further activates the motor neurons. These in turn activate the simulated muscles of the avatar.

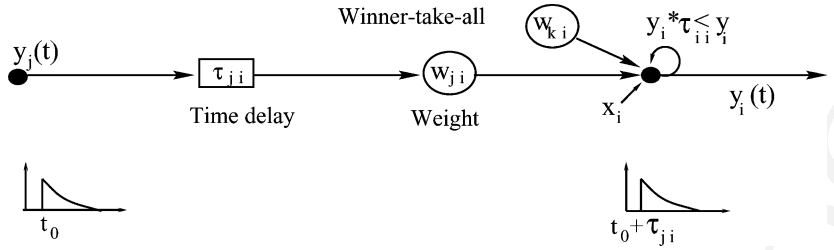

Fig. 4. A schematic of the node connection from unit i and unit j. Each connection of the DRAMA network is associated with two parameters, a weight w_{ij} and a time parameter τ_{ij}. Weights correspond to the synaptic strength, while the time parameter specifies a synaptic delay. Each unit has a self connection. Retrieval follows a winner-take-all rule on the weights.

2.4. The learning modules

Learning of motor sequences is done by updating the connectivity between the primary cortex (M1), the PM cortex, and the cerebellum modules. PM and cerebellum modules consist of a DRAMA [9], a fully recurrent neural network without hidden units. Similarly to time delay networks [35], each connection is associated with two parameters, a weight w_{ij} and a time parameter τ_{ij} (see Fig. 4). Weights correspond to the synaptic strength, while the time parameter specifies a synaptic delay, i.e., a delay on the time required to propagate the activity from one neuron to the other. Both parameters are modulated by learning in order to represent the spatial (w) and temporal (τ) regularity of the input to a node. The parameters are updated following Hebbian rules, given by Eqs. (2) and (3). Learning starts with all weights and time parameters set to zero, unless specified differently to represent predefined connection (as between PM and M1 modules, see Section 2.3).

$$\delta w_{ji}(t) = a y_i(t) y_j(t) \tag{2}$$

$$\tau_{ji}(t) = \left(\frac{(\tau_{ji}(t-1))(w_{ji}/a) + (y_j(t)/y_i(t))}{(w_{ji}/a) + 1} \right) \times y_i(t) y_j(t) \tag{3}$$

where a is a constant factor by which the weights are incremented.

In the present experiment, learning across TC–PM, PM–M1 and PM–cerebellum consists of building up the connectivity of nodes across these modules so as to represent spatio-temporal patterns of activation in the TC and PM modules, respectively. The connectivity PM–M1 is constructed simultaneously to that of TC–PM to represent the isomorphism between visual and motor representation.

In DRAMA, the neuron activation function follows a linear first order differential equation givenas follows.

$$y_i(t) = F\left(x_i(t) + \tau_{ii} y_i(t-1) + \sum_{j \neq i} G(\tau_{ji}, w_{ji}, y_j(t-1)) \right) \tag{4}$$

where F is the identity function for input values less than 1 and saturates to 1 for input values greater than 1 ($F(x) = x$ if $x \leq 1$ and $F(x) = 1$ otherwise) and G the retrieving function is given as follows.

$$G(\tau_{ji}, w_{ji}, y_j(t-1)) = A(\tau_{ji}) B(w_{ji}),$$
$$A(\tau_{ji}) = 1 - \Theta(|y_j(t-1) - \tau_{ji}|, \varepsilon(\tau_{ij})),$$
$$B(w_{ji}) = \theta(w_{ji}, \delta(w_{ij})) \tag{5}$$

The function $\Theta(x,H)$ is a threshold function that outputs 1 when $x \geq H$ and 0 otherwise. The factor ε is an error margin on the time parameter. It is equal to $0.1\tau_{ij}$ in the simulations, allowing a 10% imprecision in the record of the time delay of units co-activation. The term $\delta(w_{ij})$ is a threshold on the weight. It is equal to $((\max_{y_j} > 0(w_{ji}))/\theta(w_{ij}))\theta(w_{ij}) = 2$ in the experiments. $\max_{y_j} > 0(w_{ji})$ is the maximum value of the weight of all the connections between activated units j and unit i, which satisfy the temporal condition encoded in $A(\tau_{ji})$.

Each unit in the network has a self-connection, associated with a time parameters τ_{ii}. This provides a

short-term memory of unit activation, whose rate is specified by the value of $\tau_{ii} < 1$. This decay is represented by the term $\mathrm{d}y_i/\mathrm{d}t = (\tau_{ii} - 1)y_i$, obtained from Eq. (4), when putting to zero all other terms.

Eq. (4) can be paraphrased as follows: the output y_i of a unit i in the network takes values between 0 and 1: $y_i(t) = 1$, when (i) an input unit x_i (TC nodes input to the PM and PM nodes input to the cerebellum) has just been activated (new movement) or (ii) when the sum of activation provided by the other network units is sufficient to pass the two thresholds of time and weight, represented by the function G (see Eq. (5)). A value less than 1 represents the memory of a past full activation (value 1).

2.5. 3D biomechanical simulation of a humanoid

We added dynamics to the 3D Cosimir graphical humanoid simulation [48] of a 37 DOF avatar. Shoulders, hips, wrists, ankles and head have 3 DOFs. Elbows and knees have 1. The trunk is made of three segments with 2 DOFs each. All limbs are attached by hinge joints. The external force applied to each joint is gravity. Balance is handled by supporting the hips; ground contact is not modeled. There is no collision avoidance module.

The acceleration $\ddot{\boldsymbol{X}}_i$ and angular acceleration $\ddot{\boldsymbol{\theta}}_i$ of each link i depends on $\boldsymbol{E}_i$, the forces exerted by the environment, on $\boldsymbol{T}_i^j$, the torques due to the paired muscles of joint(s) j, and on $\boldsymbol{C}_i^j$, the inner forces due to the constraints of joint(s) j:

$$m_i\ddot{\boldsymbol{X}}_i = \boldsymbol{E}_i + \sum_j \boldsymbol{C}_i^j \tag{6}$$

$$[\mathbf{I}]_i\ddot{\boldsymbol{\theta}}_i = \sum_j \boldsymbol{T}_i^j + \sum_j \boldsymbol{C}_i^j \times \boldsymbol{r}_i^j \tag{7}$$

where m_i and $[\mathbf{I}]_i$ are the mass and the moment of inertia of link i. $\boldsymbol{r}_i^j$ is the position vector of joint j compared to the center of mass of link i.

These dynamic equations are solved using MathEngine's Fastdynamics[3] which computes the internal forces keeping the links connected, as well as the forces due to contacts, while the external forces such as the torques of the muscles, the forces due to gravity and to the air damping are given by the user.

[3] See www.mathengine.com.

2.5.1. Muscle torques

A muscle is simulated as a combination of a spring and a damper [33]. The torque exerted on each joint is determined by a pair of opposed flexor and extensor muscles. These muscles can be contracted by input signals from motor neurons, which increase their spring constant, and, therefore, reduce their resting length. The torque acting at a particular joint is therefore determined by the motoneuron activities (M_f and M_e) of the opposed flexor and extensor muscles:

$$T = \alpha(M_\mathrm{f} - M_\mathrm{e}) + \beta(M_\mathrm{f} + M_\mathrm{e} + \gamma)\Delta\varphi + \delta\Delta\dot{\varphi} \tag{8}$$

where $\Delta\varphi$ is the difference between the actual angle of the joint and the default angle. The different coefficients α, β, γ, and δ determine, respectively, the gain, the stiffness gain, the tonic stiffness, and the damping coefficient of the muscles.

3. Experiments

We present a series of experiments in which we measured the performance of the model at reproducing well-known features of human arm movement during reaching and the precision with which the model reproduced sequences of oscillatory arm movements. We also compared the performance of the model to human subjects imitating the same arm movements.

The model was implemented on eight sets of human arm motions. The first three sets were recorded using the video tracking system described in [57], and consisted of 2D oscillatory movements of the two arms in the vertical plane (lifting the shoulders up and down and bending the elbows). The other five sets were recorded using a FastTrak marker-based system (see [45] for a complete report) and consisted of 3D oscillatory movements of the left arm.

3.1. Reaching movements

We evaluated the model's performance in reproducing reaching movement of the left arm based on the data recorded using the FastTrak system (see Section

2.2). In this experiment, the model was given the target of the trajectory (i.e., the desired angle for each DOF of the shoulder and elbow) as input for the reproduction. These values were used by the spinal cord module of the model to modulate the sensory feedback. There is no learning in this example. The model's predefined connectivity for reaching (in the PC module) is exploited to generate the motions. We tested the correctness of the model in reproducing two main features associated with human arm movements, namely the bell-shaped velocity profiles and the quasi-straight hand trajectory in space [1,40,52].

Rows 4–6 of Fig. 5 show the trajectory (row 4), velocity profile (row 5), and the projected path (row 6) of the avatar's hand during a reaching movement directed towards a point at 25° in the *x*-direction and 30° in the *z*-direction. Rows 1–3 of Fig. 5 show the same quantity for the human hand in a similar reach (aimed at the same target). In both avatar and human movements, the velocity profiles for the largest directions of movements (*x* and *z*) follow a bell-shape curve. In the direction of small movements (*y*-axis), which result from internal torques caused by movement in the two other DOFs, the velocity profile is made of small oscillatory movements in both the avatar and the human. Similarly to the human data, the avatar's hand trajectory is smooth, reaching its sharpest slope at middle distance (a fact reflected by the bell-shape velocity pro-

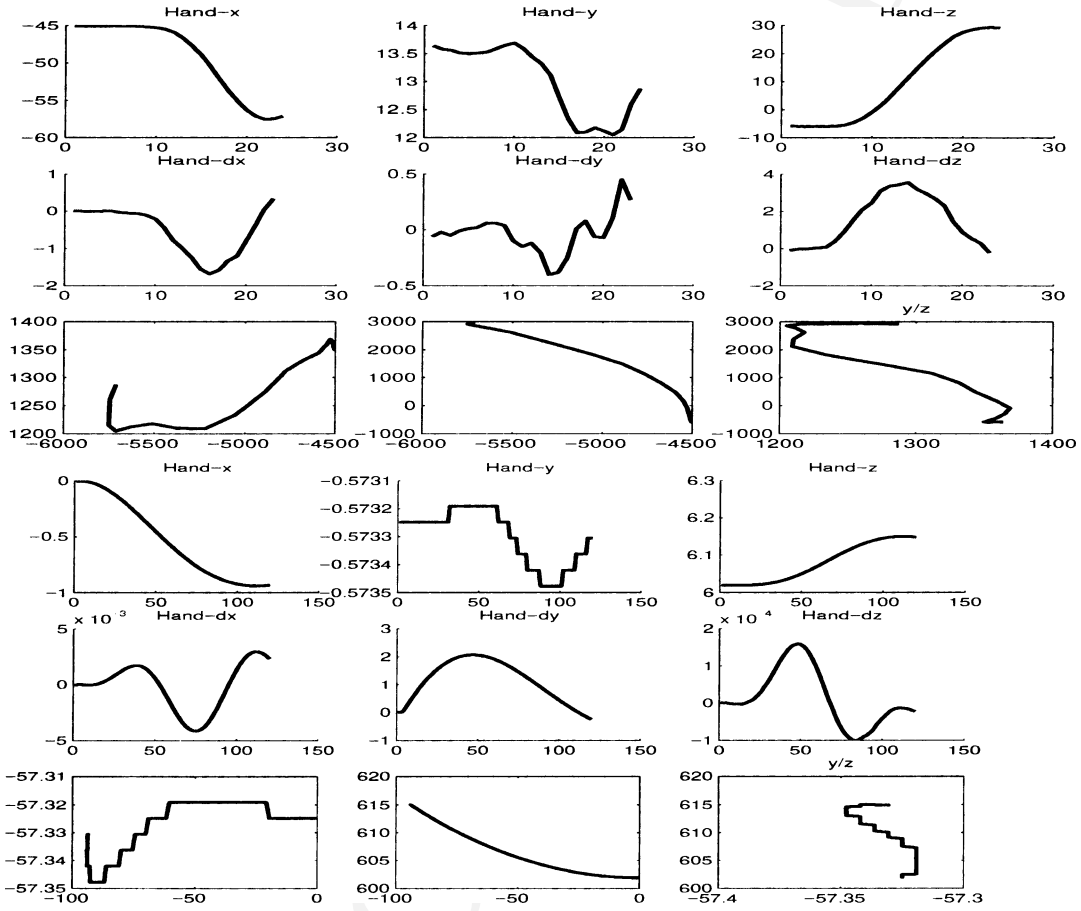

Fig. 5. Rows 1–3 from the top: human data, rows 4–6: simulation data. Trajectory (rows 1 and 4), velocity profile (rows 2 and 5) and path (rows 3 and 6) of the hand in *x*-, *y*-, *z*-directions during a reaching movement directed towards a point at 25° in the *x*-direction and 30° in the *z*-direction.

file). In our model, the slow increase of velocity for the first half of the distance is due to the smooth increase of neural activation of the motor neuron (the motor neuron's output is directly proportional to the elasticity constraint of the modeled muscles, see Eq. (8)), which follows a sigmoid (see Eq. (1)). The plateau and decrease of the velocity starting at mid-distance is due to two factors. The first is the damping of the muscle model (see Eq. (8)). The second is a mechanism in the controller which decreases the tonic input (from PM and M1 nodes) sent to the motor neurons, when receiving feedback (from the spinal cord module) indicating that the joints are at about half the desired angle.

3.2. Oscillatory arm movements

This section describes results using the three motion sets recorded with the video tracking system, which consisted of lifting up and lowering left and right upper arms (vertical rotation around the shoulders), while bending and extending the lower arms (rotation around the elbows), respectively. For each set, the motion was repeated twice.

For these experiments, the reproduction of the movement was not driven by a joint angle target as in the previous section. Here, observed motions of each limb were fed continuously to the TC module. Each change of movement triggered the TC cells. Their activity, which encoded the new orientation and speed of the movement, was passed further to the PC and cerebellum module to learn the sequence of movement. At the end of the observation, the cerebellum and PC were activated by the decision module to trigger rehearsal of the learned sequence.

Fig. 6 shows superimposed trajectories of the left and right shoulders and elbows of the avatar and the human for the three sets of motions. The black vertical lines show the instants during the movement at which the visual segmentation triggered (detecting a start or end of the motion based on velocity and direction changes). The avatar's reproduction shows a qualitative and quantitative agreement with the human movement. It reproduces all the large movements of the shoulders and elbows, with a similar amplitude. A good reproduction of the amplitude of the movement is obtained in the model by keeping a good measure of the speed of the observed movement. The speed of the movement is transmitted by the amplitude of the output of the TC cells (see Section 2.2), which is then recorded in the PM weights and further transmitted to motor neurons (in the spinal cord) as the amplitude of PM and M1 nodes' output. In the above example, we chose a 1% precision in the speed recording.

3.3. Comparison with human imitative performance

Using the data gathered in [45] on human imitation of arm movements, we evaluated the precision within which the subjects reproduce arm movements. Fig. 7 shows the trajectories of the left hand of each of four human imitators, that of the human demonstrator and that of the avatar's reproduction of the same trajectory.

The imitation by the human subjects is qualitatively similar to the demonstration, as they correctly reproduced the two oscillations in the z-direction. However, some subjects produced movements in the x- and y-directions as well. The amplitude and timing of the movement is not reproduced very well. In these two respects, the avatar's reproduction is as good as that of the human. Note that the imprecise reproduction of the avatar results from the imprecise sensory information which is given to the simulation. The avatar is given the position of each of the subject's joints, as well as that of its own joints, within 20° of precision. It is also given the speed of the subject's movement with 20% error. These values were fixed to reproduce somewhat similar imprecision as that displayed by the proprioceptive and visual sensing in humans. Had perfect sensory information been given to the avatar, the reproduction would have been perfect. However, the aim here was to make the input of the system sufficiently imprecise so as to get an output which will show patterns of imprecision similar to that of humans in their first imitation trial.

We measure the precision of the imitation based on two criteria: we measure α_i and β_i the ratio of amplitude and speed of the hand trajectory relative to axis i of demonstrator and imitator. Let $D_i(t)$ and $I_i(t)$ be the angular displacements of joint i at time t of demonstrator and imitator, respectively. $T^{\mathrm{D}}_{1,\ldots,n}$ and $T^{\mathrm{I}}_{1,\ldots,n}$ are the time steps (for each oscillation) at which $D_i(t)$ and $I_i(t)$ are maximal and T^{D}_1, T^{I}_1 are the maxima

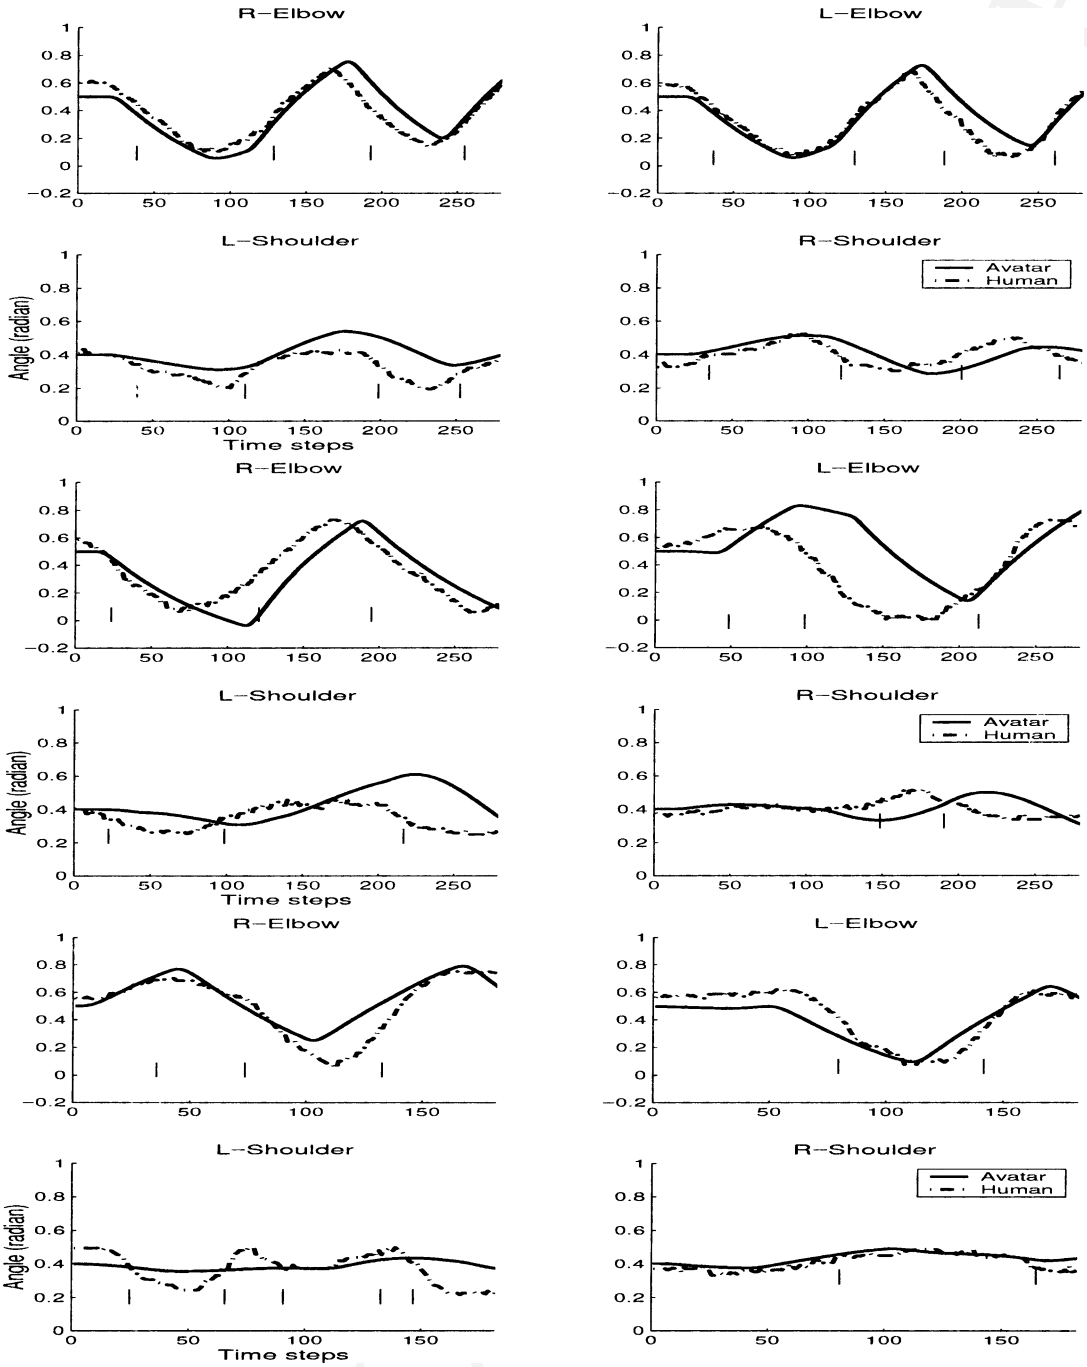

Fig. 6. Superimposed trajectories of left/right shoulder/elbow of the avatar and the human during the three movement sequences (from top to bottom).

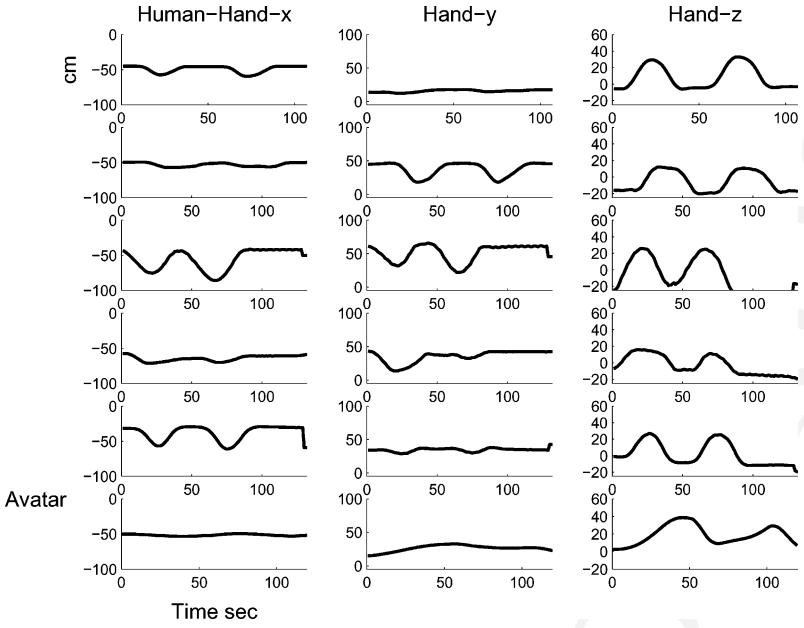

Fig. 7. Trajectories of hand motion of four human subjects and the avatar imitating an oscillatory movement of the left arm, demonstrated by another human subject. Top row: human demonstration; rows 2–5: imitation by four human subjects; sixth row: imitation by the humanoid avatar.

of each series. Then

$$\alpha = \frac{D_i(T_1^{\mathrm{D}})}{I_i(T_1^{\mathrm{I}})}$$

$$\beta = \frac{1}{n}\left(\sum_{j=2}^{n}\left\|\frac{(D_i(T_j^{\mathrm{D}}) - D_i(T_{j-1}^{\mathrm{D}}))/(T_i^{\mathrm{D}} - T_j^{\mathrm{D}})}{(I_i(T_j^{\mathrm{I}}) - I_i(T_{j-1}^{\mathrm{I}}))/(T_i^{\mathrm{I}} - T_j^{\mathrm{I}})}\right\|\right)$$

This is a straightforward measure of the observable dissimilarities between the two trajectories. α is a direct measure of the amplitude difference between the movements, while β is an indirect measure of the speed difference. In [45], other measures of similarity between the trajectories for the same reaching tasks are presented and evaluated.

Table 2 shows the mean values of these measures across imitation of the eight data sets for human imitation and avatar replication. Avatar and human performance following these measures are quantitatively similar. Both show an imprecision of over 20% on average for reproducing the amplitude and the speed of the movement.

This similarity between human and avatar data is encouraging, as the long-term goal of this study is to design a model of human ability to learn movements by imitation. Further work will focus on developing precise measures of trajectory similarities and on determining the influence of each parameter of the model and of the biomechanical simulation on the model's performance.

Table 2
Qualitative comparisons of human and avatar imitative performance[a]

	Avatar	Human
α	0.22 ± 13	0.27 ± 16
β	0.23 ± 0.21	0.25 ± 0.19

[a] α is the ratio of maxima of amplitude and β is the ratio of speed for each oscillation (mean value along the whole trial) for human and avatar hand trajectories. Data are mean values and standard deviation across imitation of eight data sets.

4. Conclusion

This paper presented a series of experiments to evaluate the performance of a connectionist model for imitating human arm movements. The model is com-

posed of a hierarchy of artificial neural network models, each of which gives an abstract representation of the functionality of some brain area involved in motor control. These are the spinal cord, the primary and PM cortices (M1 and PM), the cerebellum, and the TC.

The model was implemented in a biomechanical simulation of a humanoid avatar with 37 DOFs. Data for the imitation were recordings of human arm motions for reaching and oscillatory movements. To validate the model using real data, as opposed to simulation, and using a complete biomechanical simulation was very important to us, as our goal is to implement the system on a real robotic platform.

Results showed that the model could reliably reproduce all motions, in spite of the highly noisy input data. We measured a good quantitative agreement between simulated and real data, based on an error measure of the amplitude and speed of the movement. Moreover, the measured error in the model's reproduction was comprised within the range of error made by humans engaged in the same imitation task. These results suggest that the connectionist model, coupled to the biomechanical simulation, could be a good first approximation of human imitation. Future work will aim at evaluating further the model's performance on more data and at comparing its performance in tasks covered by other models of human motor control, such as [20,21,29,51].

Acknowledgements

Thanks to Stefan Weber for providing the data and the vision-based motion tracking software. Many thanks to the Robotics Institute at the University of Dortmund for providing the Cosimir simulator. This work was supported in part by the National Science Foundation, CAREER Award IRI-9624237 to M.J. Matarić and in part by the Office of Naval Research. A. Billard was also supported in part (1999) by a fellowship from the Medicus Foundation, New York, and by a fellowship from the Swiss National Science Foundation (2000).

References

[1] W. Abend, E. Bizzi, P. Morasso, Human arm trajectory formation, Brain 105 (1981) 331–348.

[2] R.A. Andersen, H.S. Lawrence, D.C. Bradley, J. Xing, Multimodal representation of space in the posterior parietal cortex and its use in planning movements, Annual Review of Neuroscience 20 (1997) 303–330.

[3] M. Arbib, A. Billard, M. Iacobonni, E. Oztop, Mirror neurons, imitation and (synthetic) brain imaging, Neural Networks (2000), in press.

[4] H. Asada, H. Izumi, Automatic program generation from teaching data for the hybrid control of robots, IEEE Transactions on Robotics and Automation 5 (2) (1989) 166–173.

[5] C.G. Atkeson, S. Schaal, Learning tasks from a single demonstration, in: Proceedings of IEEE International Conference on Robotics and Automation, Vol. 2, 1997.

[6] A.K. Bejczy, Towards development of robotic aid for rehabilitation of locomotion-impaired subjects, in: Proceedings of the First Workshop on Robot Motion and Control, 1999, pp. 9–16.

[7] L. Berthouze, P. Bakker, Y. Kuniyoshi, Learning of oculo-motor control: a prelude to robotic imitation, in: Proceedings of the 1996 IEEE/RSJ International Conference on Intelligent Robots and Systems'96, 1996, pp. 376–381.

[8] A. Billard, Learning motor skills by imitation: a biologically inspired robotic model, Special Issue on Imitation in Animals and Artifacts, Cybernetics and Systems Journal 32 (1–2) (2001) 155–193.

[9] A. Billard, G. Hayes, Drama, a connectionist architecture for control and learning in autonomous robots, Adaptive Behavior 7 (1) (1999) 35–64.

[10] A. Billard, M.J. Matarić, Learning motor skills by imitation: a biologically inspired robotic model, in: Proceedings of the Fourth International Conference on Autonomous Agents (Agents 2000), Barcelona, Catalonia, Spain, June 3–7, 2000.

[11] A. Billard, S. Schaal, A connectionist model for on-line learning by imitation, in: Proceedings of the International Conference on Intelligent Robots and Systems, IROS'01, Hawaii, USA, November 2001.

[12] G. Cheng, Y. Kuniyoshi, Complex continuous meaningful humanoid interaction: a multi sensory-cue based approach, in: Proceedings of IEEE International Conference on Robotics and Automation (ICRA 2000), 2000, pp. 2235–2242.

[13] V.R. de Angulo, C. Torras, Self-calibration of a space robot, IEEE Transactions on Neural Networks 8 (4) (1997) 951–963.

[14] N. Delson, H. West, Robot programming by human demonstration: adaptation and inconsistency in constrained motion, in: Proceedings of IEEE International Conference on Robotics and Automation, Vol. 1, 1996, pp. 30–36.

[15] J. Demiris, Movement imitation mechanisms in robots and humans, Ph.D. Thesis, Department of Artificial Intelligence, University of Edinburgh, May 1999.

[16] J. Demiris, S. Rougeaux, G.M. Hayes, L. Berthouze, Y. Kuniyoshi, Deferred imitation of human head movements by an active stereo vision head, in: Proceedings of the Sixth IEEE International Workshop on Robot Human Communication, IEEE Press, Sendai, Japan, September 1997, pp. 88–93.

[17] G. di Pellegrino, L. Fadiga, L. Fogassi, V. Gallese, G. Rizzolatti, Understanding motor events: a neurophysiological study, Experimental Brain Research 91 (1992) 176–180.
[18] J. Dias, A. de Almeida, H. Araujo, J. Batista, Camera recalibration with hand-eye robotic system, in: Proceedings of IEEE International Conference on Industrial Electronics, Control and Instrumentation, Vol. 1, 1991, pp. 1923–1928.
[19] A. Fod, M.J. Matarić, O.C. Jenkins, Automated derivation of primitives for movement classification, in: Proceedings of the First IEEE-RAS International Conference on Humanoid Robotics, MIT Press, Cambridge, MA, 2000.
[20] H. Gomi, M. Kawato, Equilibrium-point control hypothesis examined by measured arm stiffness during multijoint movement, Science 272 (1996) 117–120.
[21] S.R. Goodman, G.L. Gottlieb, Analysis of kinematic invariances of multijoint reaching movement, Biological Cybernetics 73 (4) (1995) 311–322.
[22] J.J. Hopfield, Neurons with graded response properties have collective computational properties like those of two-state neurons, in: Proceedings of the National Academy of Sciences, Vol. 81, The Academy, Washington, 1984, pp. 3088–3092.
[23] G.E. Hovland, P. Sikka, B.J. McCarragher, Skill acquisition from human demonstration using a hidden Markov model, in: Proceedings of IEEE International Conference on Robotics and Automation, Minneapolis, MN, 1996, pp. 2706–2711.
[24] A.J. Ijspeert, J. Hallam, D. Willshaw, Evolving swimming controllers for a simulated lamprey with inspiration from neurobiology, Adaptive Behavior 7 (2) (1999) 151–172.
[25] N. Ishikawa, K. Suzuki, Development of a human and robot collaborative system for inspecting patrol of nuclear power plants, in: Proceedings of the Sixth IEEE International Workshop on Robot and Human Communication, 1997, pp. 118–123.
[26] O.C. Jenkins, M.J. Matarić, S. Weber, Primitive-based movement classification for humanoid imitation, in: Proceedings of the First IEEE-RAS International Conference on Humanoid Robotics, MIT Press, Cambridge, MA, 2000.
[27] M. Kaiser, R. Dillmann, Building elementary robot skills from human demonstration, in: Proceedings of the International Conference on Robotics and Automation, Vol. 3, 1996.
[28] F. Kanehiro, M. Inaba, H. Inoue, Action acquisition framework for humanoid robots based on kinematics and dynamics adaptation, in: Proceedings of IEEE International Conference on Robotics and Automation, Vol. 2, 1999, pp. 1038–1043.
[29] A. Karniel, G.F. Inbar, A model for learning human reaching movements, Biological Cybernetics 77 (3) (1997) 173–183.
[30] K. Kawamura, D.M. Wilkes, T. Pack, M. Bishay, J. Barile, Humanoids: future robots for home and factory, in: Proceedings of the First International Symposium on Humanoid Robots, Waseda University, Tokyo, Japan, 1996, pp. 53–62.
[31] C. Kertzman, U. Schwarz, T.A. Zeffiro, M. Hallett, The role of posterior parietal cortex in visually guided reaching movements in humans, Experimental Brain Research 114 (1) (1997) 170–183.
[32] M.I. Kuniyoshi, I. Inoue, Learning by watching: extracting reusable task knowledge from visual observation of human performance, in: Proceedings of IEEE Transactions on Robotics and Automation, Vol. 10, No. 6, 1994, pp. 799–822.
[33] F. Lacquaniti, J.F. Soechting, Simulation studies on the control of posture and movement in a multi-jointed limb, Biological Cybernetics 54 (1986) 367–378.
[34] M.A. Lewis, G.A. Bekey, Automation and robotics in neurosurgery: prospects and problems, Neurosurgery for the Third Millenium (1992) 65–79 (Chapter 6).
[35] D.T. Lin, P.A. Ligomenides, J.E. Dayhoff, Learning spatio-temporal topology using an adaptive time-delay neural network, in: Proceedings of World congress on Neural Networks, Vol. 1, Portland, OR, 1993, pp. 291–294.
[36] Y. Louhisalmi, I. Leinonen, On research of directly programmable surgical robot, in: Proceedings of the 18th Annual International Conference of the IEEE, Engineering in Medicine and Biology Society, Bridging Disciplines for Biomedicine, Vol. 1, 1997, pp. 229–230.
[37] M.J. Matarić, Sensory-motor primitives as a basis for imitation: linking perception to action and biology to robotics, in: C. Nehaniv, K. Dautenhahn (Eds.), Imitation in Animals and Artifacts, MIT Press, Cambridge, MA, 2001, in press.
[38] M.J. Matarić, M. Pomplun, Fixation behavior in observation and imitation of human movement, Cognitive Brain Research 7 (2) (1998) 191–202.
[39] A. Meltzoff, The human infant as imitative generalist: a 20-year progress report on infant imitation with implications for comparative psychology, in: C.M. Heyes, B.G. Galef (Eds.), Social Learning in Animals: The Roots of Culture, Academic Press, New York, 1990.
[40] P. Morasso, Spatial control of arm movements, Experimental Brain Research 42 (1981) 223–428.
[41] W. Penfield, I. Rassmussen, The Cerebral Cortex of Man: A clinical Study of Localisation of Function, Macmillan, New York, 1950.
[42] D.I. Perret, M. Harries, R. Bevan, S. Thomas, P.J. Benson, A.J. Mistlin, A.J. Chitty, J.K. Hietanene, J.E. Ortega, Frameworks of analysis for the neural representation of animate objects and actions, Journal of Experimental Biology 146 (1989) 87–113.
[43] D.I. Perret, M. Harries, A.J. Mistlin, A.J. Chitty, Three stages in the classification of body movements by visual neurons, in: H.B. Barlow, et al. (Eds.), Images and Understanding, Cambridge University Press, Cambridge, 1989, pp. 94–107.
[44] D.I. Perret, P.A.J. Smith, A.J. Mistlin, A.J. Chitty, A.S. Head, D.D. Potter, R. Broennimann, A.D. Milner, M.A. Jeeves, Visual analysis of body movements by neurons in the temporal cortex of the macaque monkey: a preliminary report, Behavioral Brain Research 16 (1985) 153–170.
[45] M. Pomplun, M.J. Matarić, Evaluation metrics and results of human arm movement imitation, in: Proceedings of the First IEEE-RAS International Conference on Humanoid Robotics, MIT Press, Cambridge, MA, September 7–8, 2000.
[46] G. Rizzolatti, L. Fadiga, V. Gallese, L. Fogassi, Premotor cortex and the recognition of motor actions, Cognitive Brain Research 3 (1996) 131–141.

[47] J. Roning, A. Korzun, A method for industrial robot calibration, in: Proceedings of IEEE International Conference on Robotics and Automation, Vol. 4, 1997, pp. 3184–3190.
[48] E. Freund, J. Rossmann, Projective virtual reality: bridging the gap between virtual reality and robotics, in: Proceedings of IEEE Transaction on Robotics and Automation Special Section on Virtual Reality in Robotics and Automation, Vol. 15, No. 3, June 1999, pp. 411–422. www.irf.de/cosimir.eng/.
[49] S. Schaal, Learning from demonstration, Advances in Neural Information Processing Systems 9 (1997) 1040–1046.
[50] S. Schaal, Is imitation learning the route to humanoid robots? Trends in Cognitive Sciences 3 (6) (1999) 233–242.
[51] S. Schaal, D. Sternad, Programmable pattern generators, in: Proceedings of the Third International Conference on Computational Intelligence in Neuroscience, Research Triangle Park, NC, 1998, pp. 48–51.
[52] L.E. Sergio, S.H. Scott, Hand and joint paths during reaching movements with and without vision, Biological Cybernetics 122 (2) (1998) 157–164.
[53] K. Sing Bing Kang, Ikeuchi, A robot system that observes and replicates grasping tasks, in: Proceedings of the Fifth International Conference on Computer Vision, 1995.
[54] P.S.G. Stein, S. Grillner, A.I. Selverston, D.G. Stuart, Neurons, Networks and Motor Behavior, A Bradford Book, MIT Press, Cambridge, MA, 1997.
[55] S. Thrun, Explanation-based Neural Network Learning — A Lifelong Learning Approach, Kluwer Academic Publishers, Boston, MA, 1996.
[56] G. Vallar, E. Lobel, G. Galati, A. Berthoz, L. Pizzamiglio, D. Le Bihan, A fronto-parietal system for computing the egocentric spatial frame of references in humans, Experimental Brain Research 124 (1999) 281–286.
[57] S. Weber, Simple human torso tracking from video, Technical Report IRIS-00-380, Institute for Robotics and Intelligent Systems, University of Southern California, 2000.
[58] W.-K. Song, H.-Y. Lee, J.-S. Kim, Y.-S. Yoon, B. Zeungnam, Kares: intelligent rehabilitation robotic system for the disabled and the elderly, in: Proceedings of IEEE, 20th Annual International Conference on Engineering in Medicine and Biology Society, Vol. 5, 1998, pp. 2682–2685.

Aude Billard is a research assistant professor at the Department of Computer Science at the University of Southern California. She received a B.Sc. (1993) and an M.Sc. (1995) in Physics from the Swiss Federal Institute of Technology in Lausanne (EPFL), with specialization in Particle Physics, done at the CERN (European Center for Nuclear Research). She received a Ph.D. from the Department of Artificial Intelligence at the University of Edinburgh in 1998. She has worked as a postdoctoral fellow at IDSIA and LAMI-EPFL (1998–1999) and then joined the University of Southern California as research associate. Her work has been on designing Artificial Neural Network algorithms with application in robotics and computational neuroscience. She does research on learning by imitation and language acquisition, and seeks to understand the neural and cognitive mechanisms underlying these skills in animals. She takes inspiration in the biology of imitation learning to develop adaptive robot controllers, offering flexible means of human–robot interaction.

Maja J. Matarić is an associate professor in the Computer Science Department and the Neuroscience Program at the University of Southern California, the Director of the USC Robotics Research Lab and an Associate Director of IRIS (Institute for Robotics and Intelligent Systems). She received her Ph.D. in Computer Science and Artificial Intelligence from MIT in 1994, her MS in Computer Science from MIT in 1990, and her BS in Computer Science from the University of Kansas in 1987. She is a recipient of the NSF Career Award, the IEEE Robotics and Automation Society Early Career Award, and the MIT TR100 Innovation Award and is featured in the documentary “Me and Isaac Newton”. She has worked at NASA’s Jet Propulsion Lab, the Free University of Brussels AI Lab, LEGO Cambridge Research Labs, GTE Research Labs, the Swedish Institute of Computer Science, and ATR Human Information Processing Labs. Her research is in the areas of control and learning in behavior-based multi-robot systems and skill learning by imitation based on sensory-motor primitives.