

# Goal-Directed Imitation in a Humanoid Robot

Sylvain Calinon, Florent Guenter and Aude Billard

*Autonomous Systems Lab (ASL3)*

*Swiss Federal Institute of Technology Lausanne (EPFL)*

*CH-1015 Lausanne, Switzerland*

*{sylvain.calinon, florent.guenter, aude.billard}@epfl.ch*

**Abstract**—Our work aims at developing a robust discriminant controller for robot programming by demonstration. It addresses two core issues of imitation learning, namely “*what to imitate*” and “*how to imitate*”. This paper presents a method by which a robot extracts the goals of a demonstrated task and determines the imitation strategy that satisfies best these goals. The method is validated in a humanoid platform, taking inspiration of an influential experiment from developmental psychology.

## I. INTRODUCTION

Robots programming by demonstration has become a key topic of research in robotics (see [1] for a recent overview of core approaches in the domain). Work in that area tackles the development of robust algorithms for motor control, motor learning, gestures recognition and visuo-motor integration. While the field exists for more than twenty years, recent developments, taking inspiration in biological mechanisms of imitation, have brought a new perspective to the domain [2]. Programming by demonstration, now, encompasses more of the learning components of traditional approaches and is often referred to as *learning by imitation*.

Two core issues of imitation learning are known as “*what to imitate*” and “*how to imitate*” [3]. *What to imitate* refers to the problem of determining which of the features of the demonstration are relevant for the achievement of the task [4]. *How to imitate*, also referred to as the *correspondence problem* [5], is the problem of transferring an observed motion into one’s own capabilities. Works tackling this issue have followed either an approach in which the correspondence is unique and the imitation must produce an exact, but parameterizable, reproduction of the trajectories [6]–[8], or an approach in which only a subset of predefined goals must be reproduced (e.g. [9]–[12]).

While prior work has concentrated on either of these issues separately, we propose a system that combines a method for solving the *what to imitate* problem by extracting the task constraints, with a method for solving the *how to imitate* problem given a set of task constraints. This paper extends our theoretical framework for solving the *what to imitate* problem [4], in incorporating the notion of goal preference and including a method for optimizing the reproduction (*how to imitate*). The complete system is validated in a humanoid platform, reproducing an influential experiment from developmental psychology.

The experimental set-up and the methods for data pre-processing are detailed in Section II. In Section III, we present a statistical method for extracting the constraints and inferring their relative importance. In Section IV, we describe a method to optimize the trajectory, using the task constraints extracted in the first phase of the analysis. Results and discussion are presented in Section V and VI.

## II. EXPERIMENTAL SET-UP

### A. Experimental scenario

The experiment starts with the (human) demonstrator and the (robot) imitator standing in front of a table, facing each other (see Figure 4). On both sides of the table, two colored dots (red and green) have been stamped at equal distance to the demonstrator and imitator’s starting positions. In a first set of demonstrations, the demonstrator reaches for each dot alternatively with left and right arm. If the demonstrator reaches for the dot on the left handside of the table with his left arm, it is said to perform an ipsilateral motion. If conversely the demonstrator reaches the dot on the right handside of the table with his left arm, it is said to perform a contralateral motion. Then the demonstrator produces the same ipsilateral and contralateral motions, but without the presence of dots.

Each of these motions are demonstrated five times consecutively. In each case, the demonstrator starts from the same starting position. While observing the demonstration, the robot tries to make sense of the experiment by extracting the demonstrator’s intention underlying the task. I.e. it determines a set of constraints for the task, by extracting relevant features in a statistical manner. When the demonstration ends, the robot computes the trajectory that satisfies best the constraints extracted during the demonstration and generates a motion that follows this trajectory.

The scenario of our experiment is a replication of a set of psychological experiments conducted with young children and adults [13]. In these experiments, Bekkering and colleagues have shown that children have a tendency to substitute ipsilateral for contralateral gestures, when the dots are present. In contrast, when the dots are absent from the demonstration, the number of substitutions drop significantly. Thus, despite the fact that the gesture is the same in both conditions, the presence or absence of a physical object (the dot) affects importantly the reproduction. When the object is present, object selection takes the highest

priority. Children, then, nearly always direct their imitation to the appropriate target object, at the cost of selecting the “wrong” hand. When removing the dots, the complexity of the task (i.e. the number of constraints to satisfy) is decreased, and, hence, constraints of lower importance can be fulfilled (such as producing the same gesture or using the same hand). Similar experiments conducted with adults have corroborated these results, by showing that the presence of a physical object affects the reproduction<sup>1</sup>.

These experiments are informative to robotics, in helping us determine how to prioritize constraints (that we will also name goals throughout this paper) in a given task (and as such help us solve the “*correspondence problem*”). For instance, in the particular scenario, knowing the trajectory of the demonstrator’s arm and hand path might not allow us to determine unequivocally the angular trajectories of the robot’s arm. Indeed, depending on where the target is located, several constraints (goals) might compete and satisfying all of those would no always lead to a solution. For instance, in the case of contralateral motions, the robot’s arm is too small to both reach the target and perform the same gesture. In that case, it must find a trade-off between satisfying each of the constraints. This amounts to determining the importance of each constraint with respect to one another.

### B. Hardware

The demonstrator’s motions are recorded by 5 x-sens motion sensors, attached to the torso and the upper- and lower-arms. Each sensor provides the 3D absolute orientation of each segment, by integrating the 3D rate-of-turn, acceleration and earth-magnetic field, at a rate of 100Hz. The angular trajectories of the shoulder joint (3 degrees of freedom) and the elbow (1 degree of freedom) are reconstructed by taking the torso as referential, with an accuracy of 1.5 degrees.

A color-based stereoscopic vision system tracks the 3D-position of the dots, the demonstrator’s hands, and the robot’s hands at a rate of 15Hz, with an accuracy of 10 mm. The system uses two Phillips webcams with a resolution of 320x240 pixels. The tracking is based on color segmentation of the skin and the objects in the YCbCr color space (only Cb and Cr are used, to be robust to changes in luminosity).

The humanoid robot is a Fujitsu HOAP-2. In this experiment, trajectory control affects only the two arms (4 DOFs each). The torso and legs are set to a constant position to support the robot’s standing-up posture.

### C. Encoding of the data into Hidden Markov Models

In order to reduce the dimensionality of the dataset to a subset of critical features, we pre-segment the joint angle

<sup>1</sup>In that case, the response latency is used instead of the proportion of errors

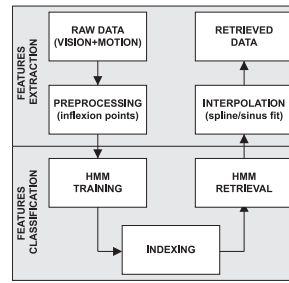


Fig. 1. Data processing loop. The trajectories are segmented into a set of keypoints. The keypoints sequences are classified using HMMs. The trajectories for the reproduction are generated by interpolation through the keypoints sequence regenerated by the HMMs.




Fig. 2. Example of a left-right continuous HMM with 5 hidden states and 2 output variables  $y_t$  and  $y'_t$ .

trajectories and the hand path into a set of keypoints, corresponding to the inflexion points (see Figure V).

The trajectories are subsequently encoded into Hidden Markov Models (HMMs), following our earlier work [14]. We here only briefly summarizes the method.

Each of the 4 joint angle trajectories is encoded in a left-right continuous HMM [15] (see Figure 1 and 2). The number of states is determined by the sequence with the highest number of keypoints in the training set. Each hidden state represents a keypoint  $j$  in the trajectory, and is associated with a stochastic representation, encoding two variables  $\{y_j, y'_j\}$ , namely the time lag between two keypoints and the absolute angle. The hand path is encoded in the same way, with 3 output distributions to encode the Cartesian components.

The transition probabilities  $P(q_t=j|q_{t-1}=i)$  and the emission distribution  $p(y_t|q_t=i)$  are estimated by the *Baum-Welch* iterative method. The *forward-algorithm* is used to estimate a log-likelihood value that an observed sequence could have been generated by one of the model. The *Viterbi algorithm* is used to generate a generalization of a trajectory over the demonstrations, by retrieving the best sequence of hidden states and the associated keypoint components. The corresponding trajectory is then reconstructed by applying a 3rd-order spline fit when using the Cartesian trajectory, and by applying a cosine fit when using the joint angle trajectory (see Figure V). The cosine fit corresponds to a cycloidal velocity profile, and keeps the keypoints as local maxima or minima during the reproduction.

## III. DETERMINING THE TASK CONSTRAINTS

In [4] and [14], we have developed a general formalism for determining the metric of imitation performance. The




Fig. 3. Function used to transform a standard deviation  $\sigma$  to a weight factor  $w \in [0, 1]$ .  $\sigma_{min}$  corresponds to the accuracy of the sensors.  $\sigma_{max}$  represents the maximal standard deviation measured during a set of demonstrations generated by moving randomly the arms around the setup during one minute.

metric (or cost function) measures the quality of the reproduction, and, as such, drives the selection of an appropriate controller for the reproduction of the task.

One way to compare the relative importance of each set of variables is to look at their variability. We take the perspective that the relevant features of the movement, i.e. those to imitate, are the features that appear most frequently, i.e. the invariants in time. We apply a cost function to determine the relevance of the different goals (or constraints) to reproduce a gesture. Following this framework, we model the task's cost function as a weighted linear combination of cost functions applied to the collected data.

#### A. Unidimensional case

Let  $D = \{x_1, x_2, \dots, x_T\}$  and  $D' = \{x'_1, x'_2, \dots, x'_T\}$  be the demonstration and the reproduction datasets of a variable  $x$ . The cost function  $J$  is defined by:

$$\begin{aligned} J(D, D') &= 1 - f(e) \\ e &= \frac{1}{T} \sum_{t=1}^T |x'_t - \mu_t| \end{aligned} \quad (1)$$

$J \in [0, 1]$  gives an estimate of the quality of the reproduction. Optimizing the imitation consists of minimizing  $J$ .  $J=0$  corresponds to a perfect reproduction.  $e$  is a measure of deviation between the observed data  $D'$  and the training data  $D$ , using the HMM representation of the data. The *Viterbi algorithm* is first used to retrieve the best sequence of states  $\{q_1, q_2, \dots, q_T\}$ , given the observation data  $D' = \{x'_1, x'_2, \dots, x'_T\}$  of length  $T$ , where  $\{\mu_1, \mu_2, \dots, \mu_T\}$  are the sequence of means associated with the sequence of states. A transformation function  $f()$  normalizes and bounds each variable within minimal and maximal values (see Figure 3). This eliminates the effect of the noise, intrinsic to each variable, so that the relative importance of each variable can be compared.

#### B. Multidimensional case

Let us consider a dataset of  $K$  variables. The complete cost function  $J_{tot}$  is given by:

$$J_{tot} = \frac{1}{K} \sum_{i=1}^K w_i J(D_i, D'_i) \quad (2)$$

where  $w_i \in [0, 1]$  is a measure for the relative importance of each cost function. These weights reflect the variance of the data during the demonstration. To evaluate this variability, we use the statistical representation provided by the HMM. If  $\{q_1, q_2, \dots, q_T\}$  is the best sequence of states retrieved by the model, and if  $\{\sigma_1^i, \sigma_2^i, \dots, \sigma_T^i\}$  is the associated sequence of standard deviations of variable  $i$ , we define:

$$w_i = f\left(\frac{1}{T} \sum_{t=1}^T \sigma_t^i\right) \quad (3)$$

If the variance of a given variable is high, i.e. showing no consistency across demonstrations, this suggests that satisfying some particular constraints on this variable will have little bearing on the task. The factors  $w_i$  in the cost function equation reflect this assumption: if the standard deviation of a given variable is low, the value taken by the corresponding  $w_i$  are close to 1. This way, the corresponding variable will have a strong influence in the reproduction of the task.

#### C. Goal-directed case

Extracting the relative importance of each set of variables statistically is sometimes too slow, and requires to have observed enough data to estimate their distribution. In order to address this problem, we have added a set of *goal-preference parameters*  $\alpha_i$ , extending our cost function to express explicitly how a constraint can be prioritized over another.

The cost function is thus applied to 4 sets of variables, namely the joint angle trajectories, the hand path, the location of the objects at which actions are directed (the dots), and the laterality of the motion (which hand is used).

Let  $D = \{\Theta, X, d, h\}$  and  $D' = \{\Theta', X', d', h'\}$  be the datasets generated by the demonstrator and imitator respectively.  $\{\theta_1, \theta_2, \theta_3, \theta_4\}$  are the generalized joint angle trajectories over the demonstrations,  $\{\vec{x}_1, \vec{x}_2, \vec{x}_3\}$  the generalized Cartesian trajectory of the hand,  $\{o_{11}, o_{12}, o_{13}\}$  and  $\{o_{21}, o_{22}, o_{23}\}$  the 3D location of the first and second dot respectively. We compute  $d_{kj} = x_j - o_{kj}$  a difference measure for component  $j$  between the hand and a dot  $k$ , at the end of a trajectory.  $h \in \{1, 2\}$  corresponds to the usage of the left and right arm respectively.

With  $N = 4$  joint angles for each arm and  $O = 2$  objects, and given the position of the hand and the objects defined by  $P = 3$  variables in Cartesian space (we make the assumption that only one hand is used at the same time), we define the general cost function  $J_{tot}$  as:

$$\begin{aligned}
J_{tot} &= \alpha_1 \frac{1}{N} \sum_{i=1}^N w_1^i J_1(\vec{\theta}_i, \vec{\theta}'_i) \\
&+ \alpha_2 \frac{1}{P} \sum_{j=1}^P w_2^j J_2(\vec{x}_j, \vec{x}'_j) \\
&+ \alpha_3 \frac{1}{OP} \sum_{k=1}^O \sum_{j=1}^P w_3^{kj} J_3(d_{kj}, d'_{kj}) \\
&+ \alpha_4 w_4 J_4(h, h')
\end{aligned} \tag{4}$$

The new factors  $\alpha_i$ , with  $\sum_i \alpha_i = 1$  and  $0 \leq \alpha_i \leq 1$ , are fixed by the experimenter, and determine the relative importance of each subset of data, i.e. the importance of each constraint (or goal) in the overall task ( $\alpha_1$  for the joint angle trajectories,  $\alpha_2$  for the hand path,  $\alpha_3$  for the hand-object distance, and  $\alpha_4$  for the laterality<sup>2</sup>).

The cost function  $J_{1,2,3}$  are given by Equation 2<sup>3</sup>.  $J_4$  is given by  $J_4(h, h') = |p(h=1) - p(h'=2)|$ , where  $p(h=1)$  is the probability of using the left arm during the demonstrations.

$w_{1,2,3}^j$  follow Equation 3, and are set to 0 if the corresponding component is missing (e.g. if the dot is not detected by the vision system).

$w_4$  is given by  $w_4 = 2 |p(h=1) - 0.5|$ . It represents the importance of using either the left or right hand (laterality of the imitation) and is based on a measure of the probability with which either hand has been used over the whole set of demonstrations ( $w_4=0$  if there is no preference).

#### IV. DETERMINING AN OPTIMAL IMITATION

Once the cost function and the relative influence of each constraint have been determined, we generate an optimal (with respect to the cost function  $J_{tot}$ ) trajectory. In order to do this, we first generate a set of candidate trajectories for the hand path, using the HMMs. Because of the difference in size between the demonstrator and the robot, we rescale the trajectories so that the starting and final positions of the demonstrator's hand correspond to the rest position of the robot's hand and corresponding location in the robot's space, respectively. In order to allow the robot to do by default a mirror-imitation, we transform the demonstrated trajectory following a vertical symmetry. The trajectories are then interpolated from the set of keypoints.

The new trajectories are considered as candidates for the hand path. To generate the corresponding joint angle trajectories, we have to solve the inverse kinematics equation given by  $\vec{x} = \mathbf{T}\vec{\theta}$ , where  $\mathbf{T}$  is the Jacobian. A solution to this equation can be found using the *pseudo-inverse with optimization* numerical solution [16]:

<sup>2</sup>We make the assumption that a mirror imitation is used in preference by the robot

<sup>3</sup> $J_3$  is similar to  $J$ , with a HMM of one state

$w_1^1 = 0.96$	$w_2^2 = 0.98$	$w_3^3 = 0.88$	$w_4^4 = 0.73$
$w_1^2 = 0.89$	$w_2^3 = 0.93$	$w_3^2 = 0.82$	
$w_3^{11} = 1.00$	$w_3^{12} = 1.00$	$w_3^{13} = 1.00$	
$w_3^{21} = 1.00$	$w_3^{22} = 1.00$	$w_3^{23} = 1.00$	
$w_4 = 1.00$			

TABLE I

MEAN WEIGHT VALUES FOUND BY THE SYSTEM, WITH THE PRESENCE OF DOTS. BY REMOVING THE DOTS, THE VALUES STAY THE SAME, EXCEPT THAT THE  $w_3^j$  VALUES ARE NULL.

	$\alpha_1=\alpha_2=\alpha_3=\alpha_4$		$\alpha_1=\frac{1}{2}\alpha_2=\frac{1}{4}\alpha_3, \alpha_4=0$	
	Dots	No dots	Dots	No dots
Left contralateral	<b>0.16</b>	<b>0.14</b>	0.22	<b>0.11</b>
Right ipsilateral	0.36	0.47	<b>0.08</b>	0.16

TABLE II

VALUES OF THE COST FUNCTION  $J_{tot}$ . IN BOLD: OPTIMAL REPRODUCTION, WITH A DEMONSTRATION OF A CONTRALATERAL MOTION WITH RIGHT HAND. WITH NO GOAL-PREFERENCE (LEFT COLUMN), THE ROBOT IMITATES IN MIRROR-FASHION. WITH A GOAL-DIRECTED COST FUNCTION (RIGHT COLUMN), THE ROBOT USES ITS CLOSEST HAND WHEN THE DOTS ARE PRESENT.

$$\begin{aligned}
\vec{\theta}_c &= \mathbf{T}^+ \vec{x} + \alpha(\mathbf{I} - \mathbf{T}^+ \mathbf{T})g(\vec{\theta}) \\
\mathbf{T}^+ &= \mathbf{T}^T (\mathbf{T} \mathbf{T}^T)^{-1} \\
g(\vec{\theta}) &= \vec{\theta}_r - \vec{\theta}
\end{aligned} \tag{5}$$

$g(\vec{\theta})$  is an optimization term which tends to minimize the distance between the arm position and a rest position given by the middle range of each joint angle. To avoid the singularities, we give some bounds to the joint angles.  $\vec{\theta}_c$  gives a solution for the Cartesian trajectories. In order to account for the influence of the demonstrator's joint angle trajectories, we add a second constraint to the pseudo-inverse:

$$\vec{\theta} = \gamma \vec{\theta}_c + (1 - \gamma) \vec{\theta}_d \tag{6}$$

where  $\vec{\theta}_d$  is the derivative of the joint angle trajectory generated by the HMM after training, and  $\gamma$  is a factor used to tune the influence of the two different terms (reproduction of hand path or joint angle trajectories). For each candidate path and associated set of joint angle trajectories, we compute the value of the cost function  $J_{tot}$ . We, then, determine a local optimum for  $J_{tot}$ , by gradient-descent on  $\gamma$ , starting with  $\gamma = 0$ . The corresponding (locally) optimal trajectory is, then, run on the robot to reproduce the demonstration.

#### V. EXPERIMENTAL RESULTS

Table I shows the values of the weights  $w_i^j$  found for the different conditions, in the two sets of experiment (with or without the dots). As expected, we find little variation ( $w_i^j$  close to 1) in either the joint angle trajectories, the

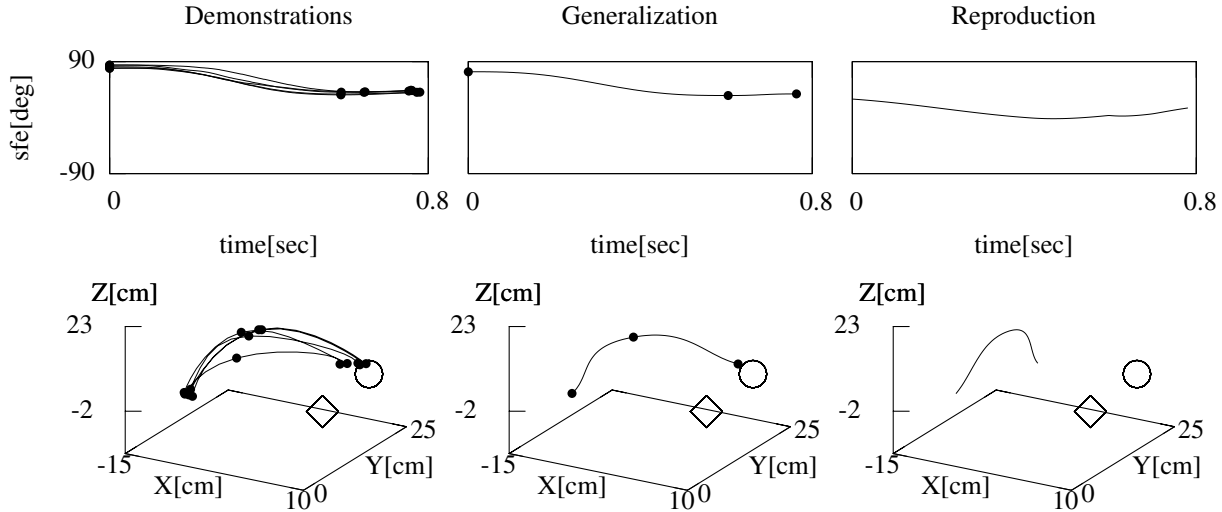


Fig. 6. Joint angle and hand path plots of 5 demonstrations of a contralateral motion with right hand (left column), the trajectory retrieved by the HMM model of the 5 demonstrations (middle column), and reproduction of a new motion by the robot to minimize the general cost function  $J_{tot}$ , with  $\alpha_i$  set to the same value (right column). The points in the graphs represent the keypoints segmented and retrieved by the HMMs. The square and the circle show the position of the two dots on the table. Only the shoulder flexion-extension is represented for the joint angles. The trajectory is retrieved with a cosine fit for the joint angles, and with a 3rd-order spline fit for the hand path.




Fig. 4. *Top*: demonstration of an ipsi- (*top-left*) and contralateral (*top-right*) motion of the right arm. *Bottom*: reproduction of the motion candidate with lowest cost function  $J_{tot}$ , by using only statistics to extract the goals ( $\alpha_1=\alpha_2=\alpha_3=\alpha_4$ ).




Fig. 5. Same experiment as in Figure 4, by adding a notion of goal-dominance to the general imitation cost function ( $\alpha_1=\frac{1}{2}\alpha_2=\frac{1}{4}\alpha_3$ ,  $\alpha_4=0$ ). This time, the contralateral motion is reproduced by doing an ipsilateral motion with the other arm, closer to the dot to touch.

hand paths, the hand-object distances or the laterality in any of the conditions, forcing the satisfaction of all constraints during the reproduction of goal-directed motion. In such a situation, since the robot does not share the same embodiment as that of the user, it will not be able to satisfy every constraint. The factors  $\alpha_i$  of the cost function  $J_{tot}$  must, thus, prioritize the constraints to fulfill.

When the target dots are not present, their associated weights  $w_3^j$  values are null. Since there are no more object in the scene, the hand path and joint angle trajectories become the sole relevant features to reproduce.

In order to test the influence of setting the preference for each set of variables (i.e. goals) on the performance of imitation, we have computed  $J_{tot}$  with two sets of values: 1)  $\alpha_1=\alpha_2=\alpha_3=\alpha_4$  (general cost function), i.e. no preference among the goals; 2)  $\alpha_1=\frac{1}{2}\alpha_2=\frac{1}{4}\alpha_3$ ,  $\alpha_4=0$  (goal-directed cost function).

For each set, we have selected the optimal controller. Table II gives the values of the cost functions  $J_{tot}$  in each case. Figures 4 and 5 show the resulting trajectories. The trajectories displayed on the images are the ones tracked by the vision system. Due to the limited range of motion

and short arm, the robot can not reach the goal with the contralateral motion. By adding a notion of goal-dominance to the general imitation cost function, the robot reproduces the contralateral motion of the demonstrator by doing an ipsilateral motion with the other arm, that is closer to the dot to touch (see Figure 5).

## VI. DISCUSSION AND FUTURE WORK

This paper has presented a method to extract the constraints of a task and to use these to determine the optimal imitation strategy. The method has been validated on a humanoid platform to perform selective reproduction of a simple set of reaching tasks, and to illustrate an example of the goal-directed imitation. The method is generic and will be applied to more complex tasks in future work. However, the implementation has made a number of assumptions that will need to be further investigated. For instance, while the weights  $w_i^j$  can be extracted easily by statistics, the factors  $\alpha_i$  (preference for the different goals), can not be determined solely through statistics. By fixing different values for  $\alpha_i$  manually, we have observed different imitative behaviors. This notion of goal dominance is more difficult to learn only through passive observation: it requires feedback from the user or explicit teaching.

This experiment have shown that using a simple goal-directed cost function can be advantageous, when a statistical analysis of the data does not permit to determine the priorities of the constraints and/or when multiple constraints can not all be fulfilled.

In our everyday life, there is a rich variety of goals, and we select only a subset of goals to reproduce a task. Using a goal-directed approach in a robotic application has also the advantage to let the robot focus on only a subset of goals, if required. If a goal has been determined, the robot can then allow itself to loose track of the other events with lower priorities, and, thus, reduces the memory capacity used to compute these sensory information.

In future work, we plan to set up experiments to learn the parameters  $\alpha_i$  of the cost function, so that the robot can determine to which extent it is more important to reproduce the same hand-object relationships, compared to the reproduction of the same hand paths, or the same joint angle trajectories. We have observed the effect that such factor can have in determining the laterality of the imitation. The choice of laterality in imitation depends on the type of actions to perform (as well as on handedness in humans). For example, the choice of hand used for pointing to an object might be less influenced by the demonstration than that of reproducing a grasping action. The factor  $\alpha_4$  could, then, be adjusted depending on the type of action being recognized in the demonstration.

In the experiments presented in this paper, the objects are static (dots on the table), and the relationships are the hand-object relative distances. Our next work will extend this notion to object-object and hand-object relationships,

taking into account relative/absolute translation and rotation.

## ACKNOWLEDGMENT

Thanks to the undergraduate students C. Dubach and V. Hentsch who helped to develop the motion tracking with x-sens and vision. This work was partially conducted within the EU Integrated Project COGNIRON funded by the European Commission Division FP6-IST Future and Emerging Technologies under Contract FP6-002020, and was supported by the Swiss National Science Foundation, through grant 620-066127 of the SNF Professorships program.

## REFERENCES

- [1] A. Billard and R. Siegwart, "Robot learning from demonstration," *Robotics and Autonomous Systems*, vol. 47:2-3, pp. 65–67, 2004.
- [2] S. Schaal, A. Ijspeert, and A. Billard, "Computational approaches to motor learning by imitation," *Philosophical Transaction of the Royal Society of London: Series B, Biological Sciences*, vol. 358, no. 1431, pp. 537–547, 2003.
- [3] C. Nehaniv and K. Dautenhahn, "Of hummingbirds and helicopters: An algebraic framework for interdisciplinary studies of imitation and its applications," in *Interdisciplinary Approaches to Robot Learning*, J. Demiris and A. Birk, Eds. World Scientific Press, 2000, vol. 24, pp. 136–161.
- [4] A. Billard, Y. Epars, S. Calinon, G. Cheng, and S. Schaal, "Discovering optimal imitation strategies," *Robotics and Autonomous Systems*, vol. 47:2-3, pp. 69–77, 2004.
- [5] A. Alissandrakis, C. Nehaniv, and K. Dautenhahn, "Imitating with alice: Learning to imitate corresponding actions across dissimilar embodiments," *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans*, vol. 32, pp. 482–496, 2002.
- [6] A. Ijspeert, J. Nakanishi, and S. Schaal, "Movement imitation with nonlinear dynamical systems in humanoid robots," in *Proceedings of the IEEE Intl Conference on Robotics and Automation*, 2002.
- [7] A. Ude, C. Atkeson, and M. Riley, "Programming full-body movements for humanoid robots by observation," *Robotics and Autonomous Systems*, vol. 47:2-3, pp. 93–108, 2004.
- [8] J. Demiris and G. Hayes, "Imitation as a dual-route process featuring predictive and learning components: A biologically-plausible computational model," in *Imitation in Animals and Artifacts*, C. Nehaniv and K. Dautenhahn, Eds. MIT Press, 2001.
- [9] R. Dillmann, "Teaching and learning of robot tasks via observation of human performance," *Robotics and Autonomous Systems*, vol. 47:2-3, pp. 109–116, 2004.
- [10] J. Zhang and B. Rössler, "Self-valuing learning and generalization with application in visually guided grasping of complex objects," *Robotics and Autonomous Systems*, vol. 47:2-3, pp. 117–127, 2004.
- [11] M. Skubic and R. Volz, "Acquiring robust, force-based assembly skills from human demonstration," in *IEEE Transactions on Robotics and Automation*, vol. 16:6, 2000, pp. 772–781.
- [12] A. Billard and S. Schaal, "A connectionist model for on-line learning by imitation," in *Proceedings of the IEEE/RSJ Int. Conference on Intelligent Robots and Systems*, hawaii, 2001.
- [13] H. Bekkering, A. Wohlschlger, and M. Gattis, "Imitation of gestures in children is goal-directed," *Quarterly Journal of Experimental Psychology*, vol. 53A (1), pp. 153–164, 2000.
- [14] S. Calinon and A. Billard, "Stochastic gesture production and recognition model for a humanoid robot," in *Proceedings of the IEEE/RSJ Intl Conference on Intelligent Robots and Systems (IROS)*, Sendai, Japan, September 28 - October 2 2004, pp. 2769–2774.
- [15] L. Rabiner, "A tutorial on hidden markov models and selected applications in speech recognition," *Proceedings of the IEEE*, vol. 77:2, pp. 257–285, February 1989.
- [16] A. Liegeois, "Automatic supervisory control of the configuration and behavior of multibody mechanisms," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 7, 1977.