Project no.:	**027657**
Project full title:	**Perception, Action & Cognition through Learning of Object-Action Complexes**
Project Acronym:	**PACO-PLUS**
Deliverable no.:	**D8.2.2**
Title of the deliverable:	**Demonstration of an artificial cognitive system that imitates human actions**

Contractual Date of Delivery to the CEC:	January 31st, 2008	
Actual Date of Delivery to the CEC:	March 25th, 2008	
Organisation name of lead contractor for this deliverable:	JSI	
Author(s):	Tamim Asfour, Pedram Azad, Aleš Ude, Andrej Kos, Volker Krüger, Dennis Herzog	
Participants(s):	UniKarl, JSI, AAU	
Work package contributing to the deliverable:	WP1, WP2, WP3, WP8	
Nature:	D	
Version:	1.0	
Total number of pages:	7	
Start date of project:	1st Feb. 2006	**Duration:** 48 months

Project co-funded by the European Commission within the Sixth Framework Programme (2002-2006)		
Dissemination Level		
PU	Public	**X**
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	

Abstract:

In this deliverable we present our work on imitating human actions on a humanoid robot. The purpose of this work is to provide means for the acquisition of action knowledge on the cognitive system. The deliverable consists of three demonstrations, for which we provide the videos:

- Action recognition and synthesis with parametric hidden Markov models, where training and synthesis uses the marker-based data while recognition uses real vision data. The demonstration shows recognition of pointing actions and synthesis of reaching movements.
- Coaching and goal-directed action synthesis from a library of example movements. We demonstrate how to train the robot to learn tasks such as ball throwing without providing any prior physical task models.
- Generation of discrete reaching and grasping movements using dynamic movement primitives.
- Reproduction of captured human motion on ARMAR-III.

Keyword list: Imitation, coaching, human motion capture, hidden Markov models, locally weighted regression

Table of Contents

Introduction

This deliverable shows the demonstrations realized on the PACO+ robotic systems in WP8.2. Here the robot needs to learn actions and associations between the perceived objects and possible actions through observation of a human demonstrator. The focus is on structuring actions by learning, imitating and understanding actions of humans, in contrast to the emphasis on exploration in WP8.1. In the first two years we explored a number of representations that can be used to achieve these goal: parametric hidden Markov models, memory-based representations and locally-weighted regression combined with the coaching paradigm, and the dynamic movement primitives. In the next three sections we present the current state of implementation on our hardware.

Action Recognition and Synthesis with parametric hidden Markov models

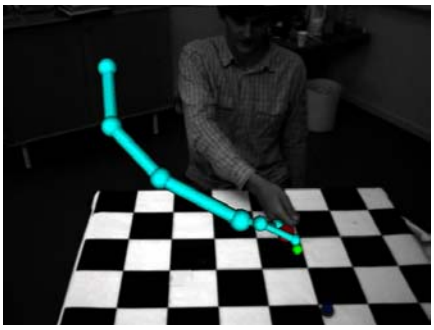

The video **pointing1.mpg** shows our system for recognition and synthesis of movements in a table-top scenario. A human agent is sitting in front of a table, on which two objects are located. The person is advising a virtual robot arm how to relocate these objects by first pointing on the object to be moved and by then pointing at the new position. Here, pointing and reaching is the matter of movement recognition from video.

A sliding ball in the video indicates the likelihood for observing a pointing action performed by the agent and its most likely respective parameterization. The green colour of the ball indicates a high likelihood for pointing, red indicates a low likelihood. Both, action and position likelihoods are estimated based on a sliding window of the recent history of the wrist and elbow positions. This is done in our exemplar-based parametric HMM framework developed by AAU (WP3.1.2). The wrist and elbow positions (**pointing2.mpg**) are given by our model-based motion capturing approach (WP3.1.1, WP3.1.3), which relies on depth images generated by a stereo camera (**depth.mpg**). For recognition and also for synthesis of reaching movements, parametric Hidden Markov Models (HMMs) are trained based on human demonstrations (not shown in video). The recognition part is based on the noisy vision data.

Video **pointing2.mpg** gives better insights into the whole setup of the scene. The calibration of the table-top and the observing camera are verified by a re-projection of the object positions onto table-top (rendered in green). The object positions, and the pointed at positions in the training session (balls at the corners of the table) are indicated with tiny balls.

The video **reaching.mpg** shows the robot HOAP-3 at JSI relocating objects on a table. The robot movements are generated based on synthesized reaching movements. The synthesized reaching movements are based on trajectories that were previously trained on movements observed on human agents (acquired using a marker-based system). Before synthesis the original human arm movements are mapped onto the robot body based on the proportions of the robot and the human arm.

The video starts with a calibration sequence where the robot places the object at for different but specific table-top positions. The object is tracked using a color tracker. After that, the robot grasps the object, wherever it is, and puts it back to a given position.

The aim of this demo is to a) test our parametric action representation that is able to take into account the state of the object (here, location and size) and thus gives rise to early OACs and to b) use movements that were previously performed by a human demonstrator.

Goal-directed action synthesis from a library of example movements

The aim of this demonstration is to show how the robot can extract full action knowledge from initial human movements, which were demonstrated by a human instructor to accomplish the task in a given situation. Motion capture has been used successfully to generate fairly complex movements on a humanoid robot. However, direct reproduction of movements, even if it includes the physical constraints of a robot, rarely results in a successful execution of the task that involves external goals. The initial movements need to be adapted to the body of the robot. Video **coaching.mpg** shows this process on the example of training throwing movements. It was created on the site of our Japanese partner ATR. As can be seen in the video, the direct reproduction ended up in a throw that missed the basket. The initial performance was suboptimal in many other ways such as for example timing of the ball release and smoothness. It was therefore necessary to develop a methodology to adapt the initial robot motion.

We explored the coaching paradigm to solve these problems [Riley et al. 2006]. Coaching provides a familiar setting to most people for interacting with and directing the behavior of a complex humanoid robot where human-robot communication takes the form of coach's demonstrations and high-level qualitative instructions. In this way it is possible to generate throwing movements that result in successful throws with good dynamical properties, which are suitable for generalization.

Based on these results we explored in simulation how to generate actions for any given configuration of the external world. Video **training.avi** shows the simulated training examples. By using a suitable representation of the trajectories and locally weighted regression [Ude at al. 2007], we have shown in simulation that based on such training data we can accurately generate movements to throw the ball anywhere in space (**generated.avi**).

Videos **PA10throws.mpg** and **HOAPthrows.mpg** shows the performance of the algorithm on the real systems. In particular, we used the Mitsubishi PA10 robot arm and the HOAP-3 humanoid robot, which are both available at JSI, to perform the experiments. On PA-10 we achieved the average precision of 3.4 cm for throwing of the ball to the distance 1.4 – 2.1 meters. The videos of HOAP-3 show the training throws, the action synthesis part on HOAP-3 is currently under evaluation. See [Ude et al., 2007] for more details.

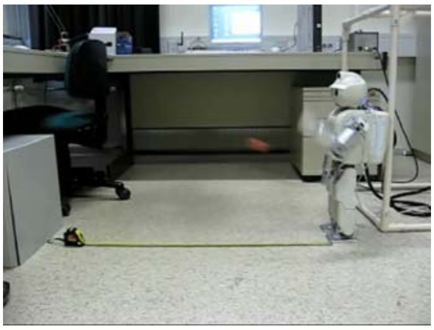

Generation of discrete movements with dynamic movement primitives

In other work we studied the applicability of dynamic movement primitives (Schaal, 2003; Ijspeert et al., 2003) in the context of the realization of an imitation learning system for object grasping and manipulation tasks. In particular, movement segmentation algorithm was developed which decomposes recorded demonstrations such that the system is able to acquire a set of DMPs from a single but complex demonstration. Semantic information has been added to the segmented movement primitives to allow the system to recall movement sequences and apply them in object manipulation tasks. The feasibility of this approach has been demonstrated in simulation as well as on a real robot. The robot learned a pick-and-place operation and a pouring water task and could generalize these tasks to novel situations.

The video **DMPs.avi** shows experimental results performed on the SARCOS robot arm at USC. The video shows two different experiments for testing the developed framework for learning by imitation and generalization to new goal positions. The first sequence in the video demonstrates the ability to reproduce the observed movement and the second sequence demonstrates the ability to generalize to new goal positions. The two experiments were:

- Pick-and-place operation, which was realized as a sequence of three DMPs: grasping, placing and releasing. The cup's initial position as well as its target location were hard-coded and used to estimate the goal position of the end effector. Furthermore, the end effector orientation after placing the cup as well as the final end effector position and orientation after releasing the cup were also hard-coded. However, the goal orientation for the grasping movement was estimated automatically.
- Pouring task, which was incorporated to show the practicability of movement primitives in quaternion space, i.e. that generalization to new goals does result in reasonable trajectories not only in Cartesian space but also in quaternion space. Therefore, the experiment required the robot to grasp a bottle, pour water into three cups, and release it on the table. This was achieved using four DMPs: grasping, pouring, repositioning, and releasing. The setup for this experiment was similar to the first in that the goal state for each of the four DMPs (except for the grasping movement) were hard-coded. Adaptation to new goal positions was obtained by setting the two appropriate goal postures such that the water was flowing into the white and black cup.

This work has been performed in cooperation with Prof. Stefan Schaal at the University of Southern California (USC), Computational Learning and Motor Control Lab, USA. Mr. Peter Pastor from UniKarl spent a semester (April – October 2007) at USC developing and implementing an imitation learning framework using DMPs under the supervision of Prof. Stefan Schaal (Scientific Advisory Board in PACO-PLUS) and Dr. Asfour.

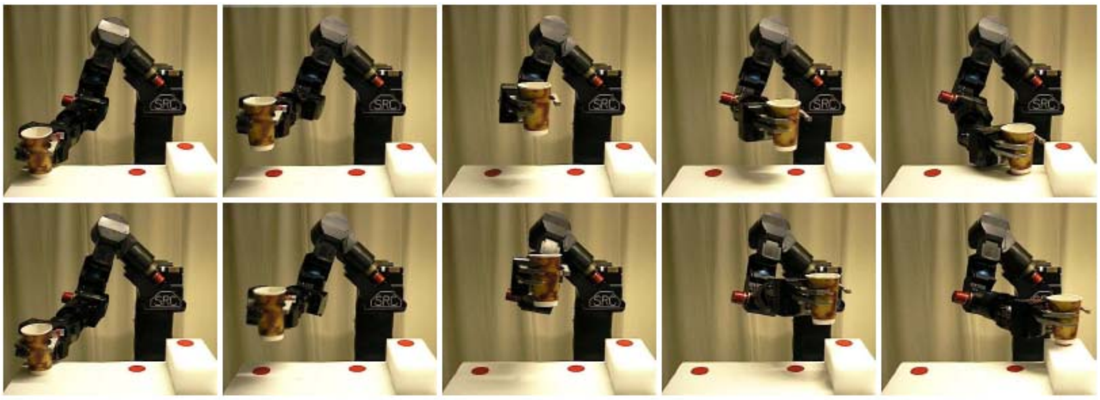

Figure 2: Pick-and-place operation performed with the Sarcos slave robot arm. Reproduction of the placing movement (upper row) and adaptation to new situation by changing the goal variable (bottom row).

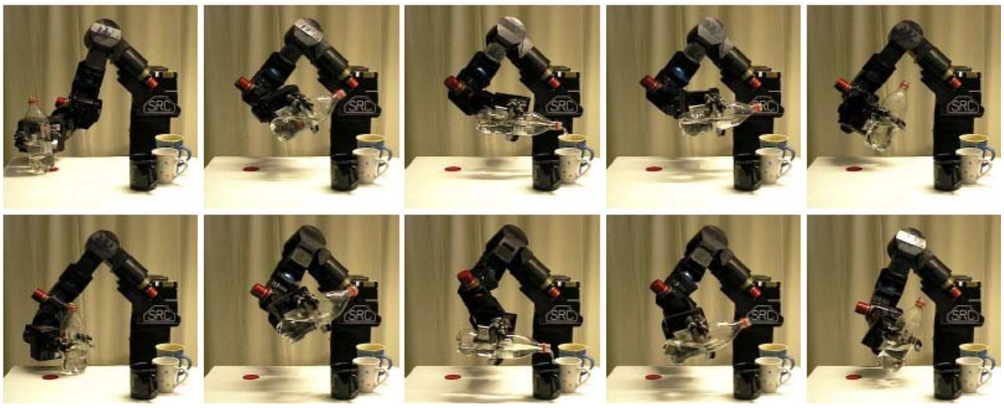

Figure 3: Pouring task perform with the Sarcos slave robot arm. Reproduction of the pouring and repositioning movement (upper row) for pouring water into the blue cup and adaptation to new situation by changing the goal variable (bottom row) for pouring water into the black cup..

Reproduction of captured human motion on ARMAR-III

In D2.1.3, it has been shown that human motion can now be captured in real-time using our stereo-based markerless human motion capture system. One important step towards *online* imitation learning is the ability to reproduce perceived trajectories on the robot.

For this purpose, we have defined the so-called Master Motor Map (MMM) [Azad et al., 2007], which serves as an exchange data format by specifying a reference kinematics model. In the presented video, the captured trajectory from D2.1.3 was reproduced in simulation on the humanoid robot ARMAR III. For this, the trajectory was post-processed by a smoothing filter, and in particular, the rotations given in terms of Euler angles were mapped to a continuous trajectory for the target shoulder joint.

Due to the mechanical implementation of the shoulder joint not all perceived poses can be mapped to ARMAR III. As illustrated in Fig.4 (Center), such infeasible configurations can only be approached to the extent to which it is allowed by the joint constraints. By solving an optimization problem, a similar configuration within the robot's joint space is found, which shows small deviation to the perceived pose while moving the tool center point close to its designated position.

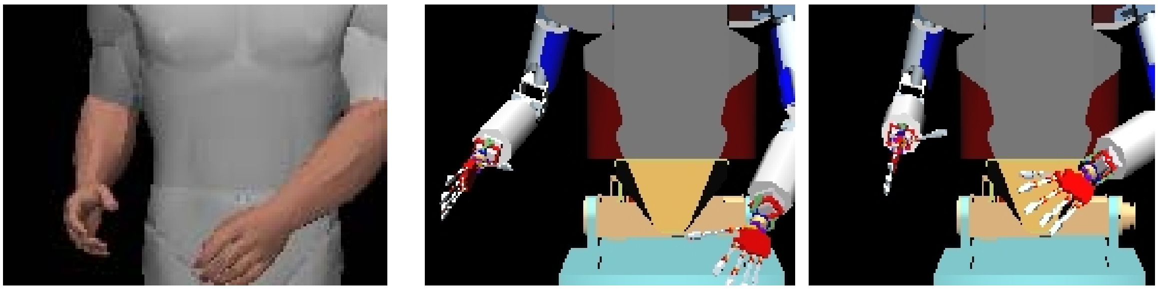

Figure 4 Left: Infeasible configuration visualized by a human model. Center: Infeasible configuration on ARMAR III. Right: Approximated configuration on ARMAR III.

The video **MovementReproductionARMAR.avi** shows the stages from the source trajectory, as it is acquired by the human motion capture system, to the final trajectory which is reproduced on ARMAR III. The video consists of the following four parts:

1. The left camera image of the stereo sequence is shown which is used as input to the markerless human motion capture system.
2. Simplified 3D visualization of the trajectory acquired by the stereo-based markerless human motion capture system presented in D2.1.3.
3. Visualization of trajectory after mapping to the MMM representation. The 3D human model is used for visualizing any trajectory given in the MMM format.
4. Visualization of the simulation result of the trajectory after mapping from the MMM representation to ARMAR III. Joint contstraints have been handles by solving the optimization problem, as explained above.

Attached Papers

P. Azad, T. Asfour, and R. Dillmann (2007) Toward an Unified Representation for Imitation of Human Motion on Humanoids. *IEEE International Conference on Robotics and Automation (ICRA)*, Rome, Italy.

A. Ude, M. Riley, A. Kos, B. Nemec, T. Asfour, and G. Cheng (2007) Synthesizing Goal-Directed Actions from a Library of Example Movements. In *Proc. IEEE-RAS International Conference on Humanoid Robots (Humanoids 2007)*, Pittsburgh, USA.

References

A. J. Ijspeert, J. Nakanishi, and S. Schaal. Learning Attractor Landscapes for Learning Motor Primitives. In Advances in Neural Information Processing Systems 15 (NIPS), 2003.

M. Riley, A. Ude, C. G. Atkeson, and G. Cheng (2006) Coaching: An approach to efficiently and intuitively create humanoid robot behaviors. In *Proc. IEEE-RAS International Conference on Humanoid Robots (Humanoids 2006)*, Genoa, Italy.

S. Schaal. Dynamic Movement Primitives - A Framework for Motor Control in Humans and Humanoid Robotics. In The International Symposium on Adaptive Motion of Animals and Machines (AMAM), 2003.