Project no.:	**027657**
Project full title:	**Perception, Action & Cognition through learning of Object-Action Complexes**
Project Acronym:	**PACO-PLUS**
Deliverable no.:	**D8.1.4**
Title of the deliverable:	**Demonstration Application I: Robot-vision system realising a three level hierarchy able to (1) compute and perform sophisticated plans (including plan monitoring and replanning) and (2) to perform learning on all three levels**

Contractual Date of Delivery to the CEC:	31.3.2009
Actual Date of Delivery to the CEC:	16.4.2009
Organisation name of lead contractor for this deliverable:	SDU

Author(s): Dirk Kraft, Norbert Krüger, Florentin Wörgötter, Nicolas Pugeault, Christopher Geib, Ron Petrick, Mark Steedman, Renaud Detry, Justus Piater, Ales Ude, Tamim Asfour and Rüdiger Dillmann
Participant(s): JSI, UEDIN, SDU, ULg, BCCN

Work package contributing to the deliverable:	WP1,WP2,WP3,WP4,WP5
Nature:	R
Version:	Final
Total number of pages:	4
Start date of project:	1st Feb. 2006 **Duration:** 48 month

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:

We present 4 videos showing our progress on developping a cognitive system which performs learning, categorisations, and planning on different levels of the processing hierarchy.

Keyword list: Cognitive Architecture, Grounding, Learning, Object-Action Complexes

Table of Contents

1. Learning and Execution of Grasp Affordance Densities

(grasp_densities.avi)

The video shows the learning of a grasp empirical density [1], the execution of grasps from the resulting empirical density, a set of examples of failed grasps, and a set of examples of successful grasps on two different objects.

The learning of an empirical density involves the execution of a large number of samples from a hypothesis density. The video presents the methods involved in the computation and execution of one grasp on a toy basket: extraction of 3D scene descriptors (multi-modal primitives), alignment of the basket model (pose estimation), repetitive sampling from the hypothesis density, execution of the first sample for which a robot motion can be planned, and assessment and storage of the outcome of the grasp.

The next sequence shows the execution of a set of samples from the empirical density of the basket, followed by a set of examples of failed grasps. The last sequence shows a set of examples of successful grasps for the basket and a toy pan.

2. Goal-directed object manipulation with ARMAR III and PKS

(CupStacking.mp4)

This video demonstrates the current state of ongoing integration efforts between UEDIN and UniKarl, to integrate the PKS planning system with the ARMAR humanoid robot platform. We consider a simple object stacking scenario where the robot is given the task of stacking a blue cup into a green cup. The planner is responsible for generating a sequence of high-level actions that achieves this goal, and for sending those actions to the low-level robot control system for execution, using the multi-level communication architecture developed as part of WP4 and the robot API developed in WP1. Four scenarios are considered, illustrating some of the complications that can arise during plan generation and execution. In the first scenario, the two cups are stacked in a straightforward manner. In the second scenario, the green cup is blocking the access to the blue cup and the planner must guide the robot to first relocate the green cup before grasping the blue cup. In the third scenario, the green cup is lying in a toppled state and the robot must first upright this object before manipulating it further. In the final scenario, both cups must be uprighted to successfully complete the task. This work provides the basis for the next stage of integration which involves more complex high-level planning and execution monitoring, in conjunction with more sophisticated robot-level interactions with objects in the kitchen domain.

3. Learning to pour by reinforcement learning

(RL-Hoap.mpg)

In this experiment we show that correct pouring position can be learned by a robot using reinforcement learning (RL). Pouring movement is hard to pre-calculate, because the stream of liquid running out of a container may have complicated physical properties. Humans would perform the task using vision to observe the stream. Robot vision is not yet so advanced to execute pouring through observing the stream. But the task of liquid pouring fits well with reinforcement learning on a robot, where the movement is performed in feed-forward fashion (without using vision for feed-back). In the first part of the movie we show that for a human (without specific training) is problematic to make a correct pouring movement in the case the eyes are closed. In the second part of the movie we show the HOAP3 robot learning to perform the same pouring task. Reinforcement learning with function approximation is used. Reward is calculated as the amount of water correctly poured into the lower glass. In this experiment the reward was obtained using scales, but measuring amount of liquid in a glass is in principle solvable task for robot vision. Actions in our RL framework are defined through gradients of the value function. Robot starts from the initial pouring position with 8 to 10 attempts to pour along the trajectory. After 3-5 such trajectories robot learns to pour correctly. This is a reasonable result for robot learning, where hundreds of unsuccessful trials would not be acceptable, but 20-30 not fully successful trials are acceptable. The movie shows the setup of the experiment, as well as the sequences of robot movements. One can see how the pouring task at first is performed unsuccessfully, but then is performed correctly after learning. Intermediate learning steps are not shown in the movie.

4. Plan execution with monitoring, resensing, and replanning

(PlanningWithMonitor.avi)

This video demonstrates the current state of ongoing integration efforts between SDU and UEDIN, to integrate the PKS planning system with SDU's robot/vision system. We consider a simple object manipulation scenario where the robot is given the task of grasping particular domain objects, in order to stack and remove those objects from a surface. The planner is responsible for generating a sequence of high-level actions that achieves this goal, and for sending those actions to the low-level robot/vision system for execution, using the multi-level communication architecture developed as part of WP4. This demo highlights the operation of a high-level plan execution monitor that communicates with the planner to help facilitate resensing and replanning activities, based on state reports provided by the robot/vision system. More details about this work can be found in Deliverable 4.3.5.

References

[1] Renaud Detry, Emre Başeski, Norbert Krüger, Mila Popović, Younes Touati, Oliver Kroemer, Jan Peters, and Justus Piater. Learning object-specific grasp affordance densities. In *International Conference on Development and Learning*, (submitted).