Learning Robot Grasping and Manipulation

TECHNISCHE UNIVERSITÄT DARMSTADT

Learning to Interact with Objects

A key long-term goal of robotics is to create autonomous robots that can perform a wide range of tasks to help humans in daily life. One of the main requirements of such domestic and service robots is the ability to manipulate a wide range of objects in their surroundings.

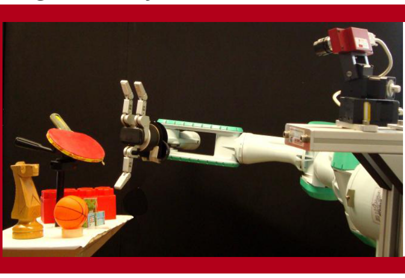

A robot preparing to manipulate objects in a cluttered scene. The robot consists of a three-fingered hand, an anthropomorphic arm, and a stereo camera.

Modern industrial robots usually only have to manipulate a single set of identical objects, using preprogrammed actions. However, future service robots will need to operate in unstructured environments and perform tasks with novel objects. These robots will therefore need to learn to optimize their actions to specific objects, as well as generalize their actions between objects.

Improved Grasps through Experience

The ability to grasp objects is an important prerequisite for performing various manipulation tasks. In order to autonomously learn better grasps, the grasp selection process can be framed as a *continuum-armed bandit* reinforcement learning problem. In this manner, the robot can actively balance executing grasps that are known to be good and exploring new grasps that may be better.

A dynamic grasp of a table-tennis racket, with preshaping (left), that the robot learned using a reinforcment learning approach.

Affordance Learning

An object's affordances are the actions that the robot can perform using the object. The affordances of basic objects, such as tools, are usually defined by their surface structures. By finding similar surface structures in different objects, the robot can transfer its knowledge of afforded actions between objects. In this manner, the robot can predict whether a novel object affords a specific action, as well as adapt the action to this object.

As shown in the figure below, the robot can learn an initial action from a human demonstration. Adapting this action to new objects is achieved autonomously by the robot, using a trial-and-error learning approach.

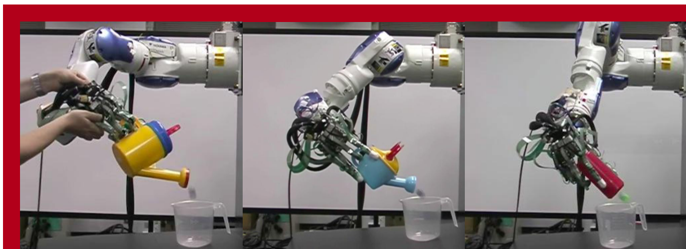

Generalization of a demonstrated pouring action to novel objects.

Task Planning

The individual grasping and manipulation actions learned by the robot form the basis for performing more complex tasks. The robot performs actions in a continuous state space. However, these actions are represented as discrete elements to allow higher-level planning. Hence, the structure of tasks can be represented using hybrid system methods, which allow for planning in both discrete and continuous states.

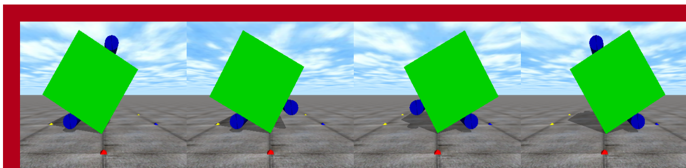

Four stages of a simulated box tumbling task using two (blue) hands.

H. Ben Amor - amor@ias.tu-darmstadt.de
A. Boularias - boularias@tuebingen.mpg.de
H. van Hoof - hoof@ias.tu-darmstadt.de
O. Krömer - kroemer@ias.tu-darmstadt.de
Prof. J. Peters - mail@jan-peters.net

Intelligent Autonomous Systems