Robotics

Figure 1.12: In the past, we have studied a single robot playing table tennis against a human or ball gun. In 2018, we migrated to a new lab and a setup with two robots. We aim to enable the unique high-speed Barrett WAM robot arms PING and PONG to jointly play table tennis in a collaborative game using methods from robotic skill learning.

Creating autonomous robots that can learn to assist humans in situations of daily life is a fascinating challenge for machine learning. While this aim has been a long-standing vision of artificial intelligence, we have yet to create robots that can learn to accomplish many different tasks triggered by environmental context or higher-level instruction. Our goal is the investigation of the ingredients for such a general approach to motor skill learning, in order to get closer towards human-like performance in robotics. We thus focus on the solution of basic problems in robotics by developing domain-appropriate machine-learning methods.

While many machine learning methods work in theory, in simplified simulations and textbook control plants, it is essential to study real robot systems to understand the learning of high-performance motor skills. We focus on the problem of learning robot table tennis as our "Drosophila" (i.e., a model system used in experiments to gain insight that we hope will be generalizable) of robot learning. This task has a number of components that are representative of tasks encountered by natural intelligent systems, including perception and action, as well as various aspects of social interaction (opponent modeling, competition, collaboration).

In our studies, we focus on using machine learning approaches for improved tracking, imitation of demonstrated behavior, and self-improvement by robot reinforcement learning with a strong focus on high-speed skill learning.

Learning approaches have to generalize a complex hitting behavior from relatively few demonstrated trajectories, which neither cover all ball trajectories nor all desired hitting directions. Therefore, past approaches that only modeled a deterministic mean behavior without capturing the variability of the movement have been fairly limited. Recent work on capturing trajectory distributions using probabilistic movement representations [5] opens new possibilities for robot table tennis. We have presented several methods to adapt probabilistic movement primitives, e.g., for adapting hitting movements learned in joint space to have a desired end effector position, velocity and orientation [229], as well as to find the initial time and duration of the movement primitive in order to intercept a moving object like the table tennis ball [5]. The resulting methods rely on straightforward operations from probability theory and provide a more principled approach to solve some of the challenges of robot table tennis compared to previous approaches. This also enabled us to learn several other tasks.

We have worked on various other questions of robot motion control with the context of robot table tennis in real robot experiments on the Barrett WAM. Among these approaches, we have studied the properties of optimal trajectory generation in robot table tennis strikes [231], learning striking controllers [6]. We have also recently demonstrated how a table tennis serve can be captured and successfully reproduced [4].

More information: https://ei.is.mpg.de/project/robotics