

MT: Learning robotic juggling

Motivation:

Dexterous manipulation is an emerging field where more and more capabilities have been shown. These capabilities are usually fairly controlled and slow, and interact with static, or quasi-static objects.

In this thesis, we want to learn fast reactive skills, particularly robotic juggling with balls or juggling clubs (as shown on the right side) and a single hand mounted on a UR5. We start out in simulation, and then deploy it on the real robot. Particularly for dealing with juggling clubs, the dexterity of the hand is crucial, compared to juggling with balls, which can be done with simpler end effectors [1].

**Goal:**

The goal of this project is the development and eventually the deployment of such a juggling and/or catching/throwing policy on the real robot.

In order to learn robotic juggling, we will start from a simpler task: catching, and throwing using a robot hand. As the task has a very sparse reward, we intend to explore either curriculum learning, or learning from data, using an imitation reward. To achieve that, we need pose estimation for the object that is being thrown, and reactive and reliable (and safe) controllers for both the arm and the hand.

Interested?

Ideally, you already have strong experience with Python and C++ and know both theory and practice of robotic reinforcement learning.

For further information or to apply for this project, please contact Valentin Noah Hartmann (valentin.hartmann@inf.ethz.ch) with your CV and transcripts.

[1] <https://www.youtube.com/watch?v=pKJEbs64Y2o>