Hierarchical Reinforcement Learning for Bimanual Mobile Manipulation

In order to perform complex contact-rich mobile manipulation tasks it is necessary to coordinate the movements of platform and the end-effectors. Reactively planning such a coordinated behavior is challenging due to unforeseen events such as humans entering the scene or imprecise localization of objects. There exist different phases during mobile manipulation tasks, such as navigation, grasp selection and execution. These are highly interleaved and dependent on each other.

To grasp an object that is located at an arbitrary position in the room, an agent must therefore (1) control the platform, (2) dynamically choose suitable grasp candidates and (3) control the end-effectors to reach and grasp objects with both hands.

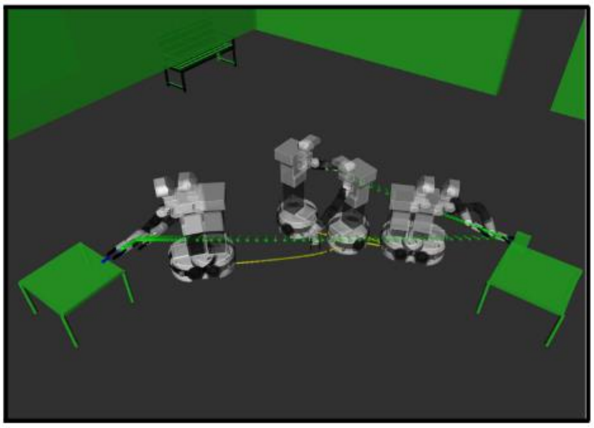

Fig. 1: Simulated mobile manipulation task [1]

Hierarchical reinforcement learning is a promising approach to learn such composite behaviors. Given a global plan including an initially collision-free path for the platform, a set of grasp candidates and a grasping strategy, the agent must adapt according to to the environment model (humans, obstacles, perceived object pose) obtained by the perception system.

In this work, you will investigate how to apply hierarchical reinforcement learning for the coordination of mobile manipulation tasks. The policy should be able to grasp objects from a cluttered scene that may consist of stacked objects. To this end, the learned policy in simulation should be transferred to the real robot ARMAR-6.

Relevant research questions include:

- How to model this problem as a reinforcement learning problem? How can a suitable reward signal be defined ?
- How can the learned policy be transferred from simulation to reality?
- How can objectives / constraints e.g. collision-free motion, time-optimal behavior, high manipulability be incorporated into the system?

This work will use the humanoid robot ARMAR-6 and its digital twin:

- ArmarX (C++): armarx.humanoids.kit.edu
- Python

Contact: Patrick Hegemann (patrick.hegemann@kit.edu), Fabian Reister (fabian.reister@kit.edu)

[1] Honerkamp, Daniel, Tim Welschehold, and Abhinav Valada. "Learning Kinematic Feasibility for Mobile Manipulation through Deep Reinforcement Learning." *arXiv preprint arXiv:2101.05325* (2021).