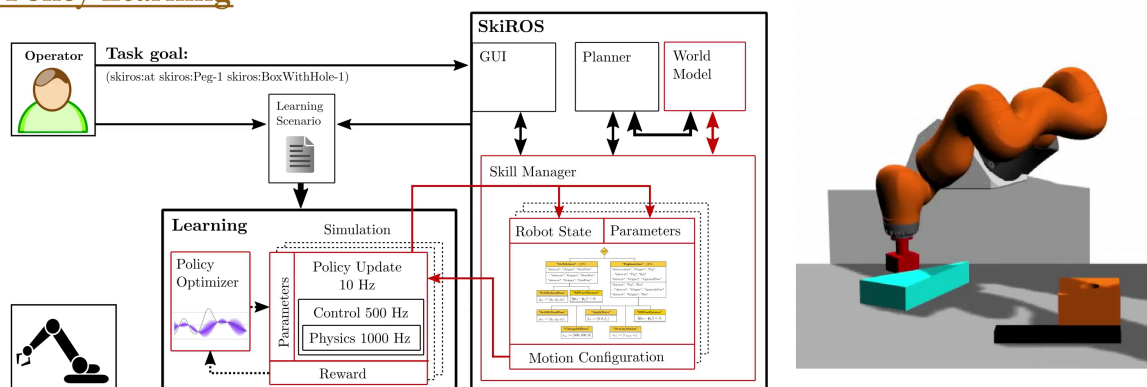


Master Project: Knowledge-Supported Reasoning for Reinforcement Learning Scenarios

Policy Learning



Knowledge- and skill-based approaches are powerful as long as all relevant aspects of a task are sufficiently well described. Many real-world tasks have *tacit* (implicit, unmodeled) aspects that are usually tuned by an operator. An alternative to this manual tuning is to use reinforcement learning (RL). However for every task to be learned, a variety of hyperparameters, such as learning iterations, need to be set. Furthermore, reward functions need to be specified for the task at hand. This is often manually done by an operator. The research question of this thesis is whether this can be done automatically with a reasoning system based on the knowledge in the world model. For example adding a negative reward for being close to a fragile object in the environment.

Your Tasks

At first you search and discuss related work in the field and get familiar with the skill-based system SkiROS¹ and its world model. A structure for both RL scenario hyperparameters and reward functions is to be developed. Such a structure should be easily reconfigurable since we learn and execute frequently changing tasks. The approach would be evaluated with a set of learning problems in an industrial setting.

Requirements

- Independent, diligent and structured way of working
- Knowledge in Python
- (Optional) A course that covered AI or RL methods
- (Optional) Experience with Linux & ROS

What we offer

- State-of-the-art research
- Option to publish the results
- Regular supervision
- Freedom to explore interesting topics
- Freedom to schedule your time

Start Date

- as soon as possible

Key Words

- Knowledge Representation, Reasoning
- AI, Reinforcement Learning
- Robotics

¹<https://github.com/RVMI/skiros2>