Karlsruhe Institute of Technology

Autonome Lernende Roboter (ALR)

Prof. Gerhard Neumann

Project Type

- [x] Master Thesis
- [ ] Bachelor Thesis
- [ ] Research Project

Supervisors

Ge Li

ge.li@kit.edu

Hongyi Zhou

hongyi.zhou@kit.edu

Difficulty

Algorithmic

Math

Application

Requirements

- [x] Python
- [x] Pytorch
- [x] Knowledge about RL and ML

Deep Reinforcement Learning with Sequenced Movement Primitives in Robot Manipulation tasks

Description

Movement Primitives (MP) [1] serve as fundamental building blocks for representing and generating smooth trajectories in robot manipulation tasks. They have gained significant popularity in conjunction with deep reinforcement learning algorithms, enabling the solution of challenging tasks characterized by non-linear causalities or high-order sensory inputs. However, the majority of MP-based tasks have focused solely on single operations, such as basic movements or striking, thus limiting the complexity of tasks that can be effectively addressed. When confronted with tasks involving multiple operations, numerous challenges arise. Firstly, the value function must accurately predict the outcome of each individual operation. Secondly, the transition from one sub-task to another must be seamlessly executed to ensure smooth movement. Moreover, the agent must operate at a sufficiently fast pace to ensure the timely completion of the entire task. Additionally, the sub-tasks may require a specific order or sequence to be solved effectively.

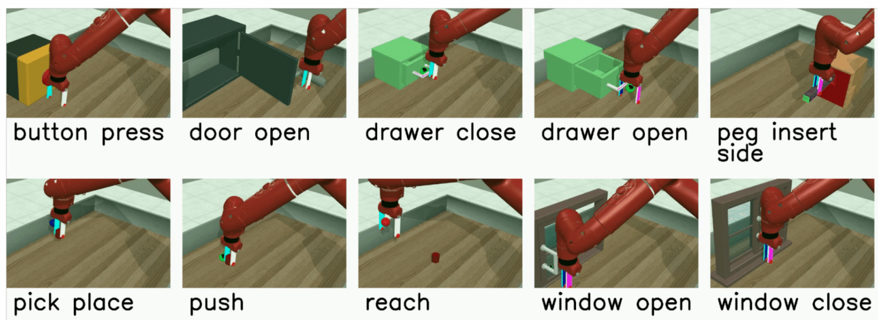

Figure 1: Metaworld tasks [2]. Most of the tasks require only single operation, but you can easily build multi-operation tasks from it. For example, first open a drawer, then pick up a ball and put it in the drawer, and close the drawer in the end.

In this thesis, we want to use deep reinforcement learning methods to solve complex robot manipulation tasks that require multi-phases operations, including:

Tasks

- Literature research and get familiar with the latest deep RL and MP methods.
- Build up or reuse or build up simulation tasks with multiple operations.
- Design the algorithm for multi-phases learning, such as event-driven trajectory replanning and switching.
- Evaluate your algorithm and compare it with baseline [3].

References

[1] Ge Li, Zeqi Jin, Michael Volpp, Fabian Otto, Rudolf Lioutikov, and Gerhard Neumann. Prodmp: A unified perspective on dynamic and probabilistic movement primitives. *IEEE Robotics and Automation Letters*, 8(4):2325–2332, 2023.

[2] Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on robot learning*, pages 1094–1100. PMLR, 2020.

[3] Haichao Zhang, Wei Xu, and Haonan Yu. Generative planning for temporally coordinated exploration in reinforcement learning. In *International Conference on Learning Representations*, 2021.