Learning a Meta-Controller for Dynamic Grasping

Yinsen Jia[1*], Jingxi Xu[1*], Dinesh Jayaraman[2], Shuran Song[1]
[1]Columbia University [2]University of Pennsylvania

SCAN ME

Motivation

Grasping moving objects is a challenging task that combines multiple modules such as *object pose predictor, arm motion planner*, etc. Each module operates under its own set of meta-parameters.

Many previous works assign fixed values to these parameters; however, at different time steps within a single episode of dynamic grasping, there should be different optimal values for each parameter, depending on the current scene.

Two Key Meta-Parameters

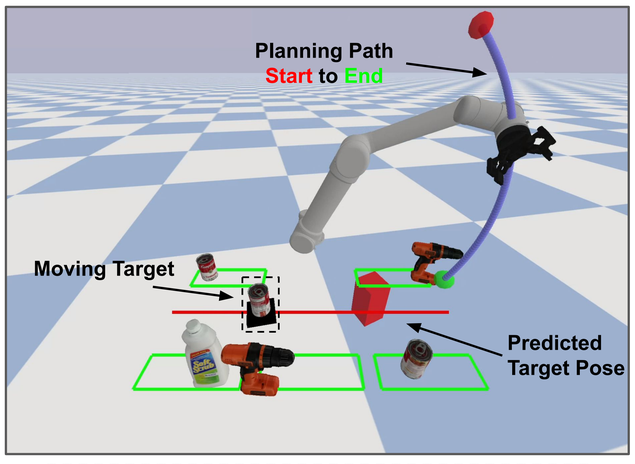

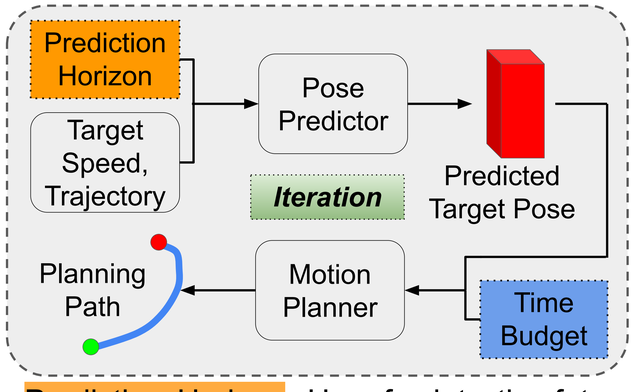

- Prediction Horizon, How far into the future the predicted object pose should be
- Time Budget, How much time should be assigned to plan a path

Dynamic Grasping with a Meta-Controller

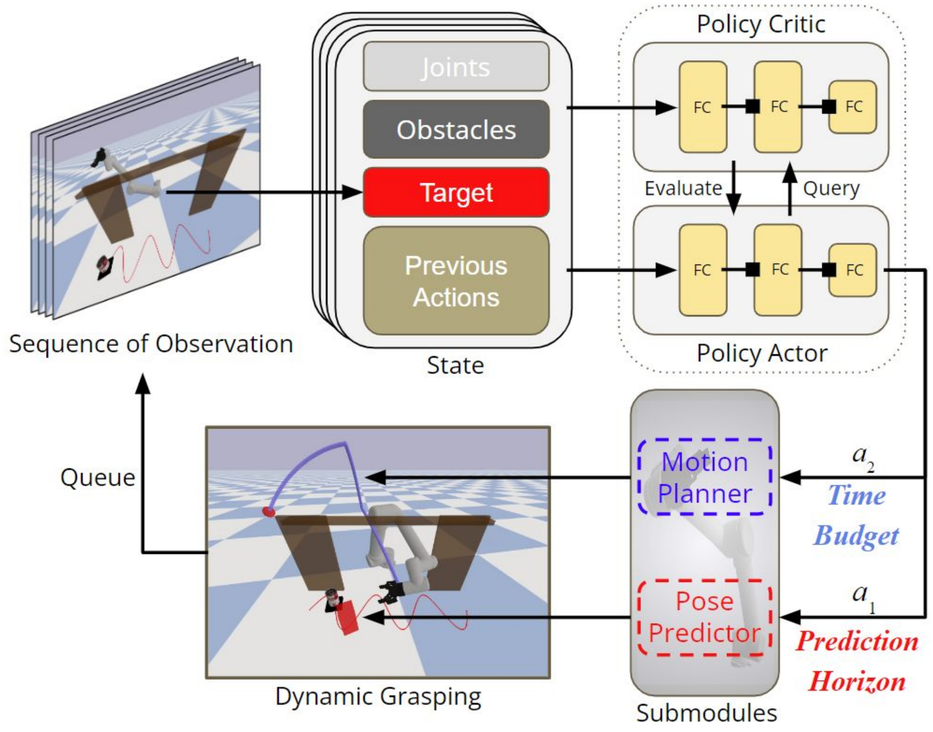

In this work, we learn a **meta-controller through reinforcement learning** to control the prediction horizon and time budget dynamically at each time step. Our meta-controller is trained with PPO and a sparse reward when the object is successfully picked up.

Experiments

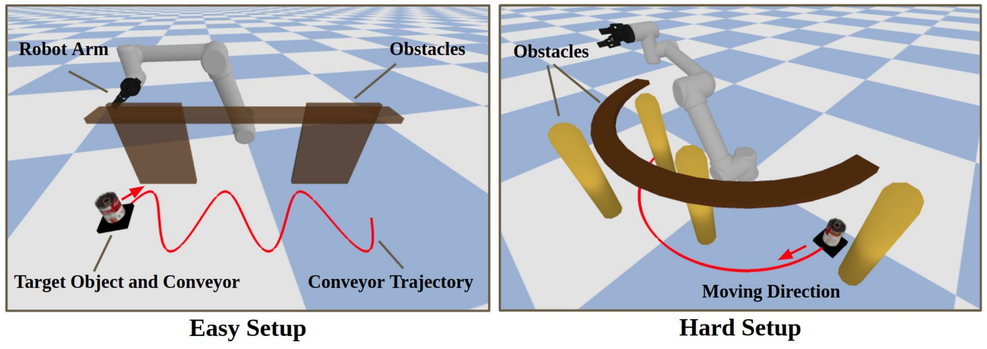

- ❖ Two setups (easy and hard) in the PyBullet simulator
- ❖ The trajectories of the object are covered and surrounded by obstacles whose poses are changed randomly in a certain range
- ❖ The object speed is sampled between 2 - 6 cm/s; The trajectory in the easy setup is selected within sinusoid, linear, and rectangular

Results

Quantitative Results

Our meta-controller improves the grasping **success rate** and reduces **grasping time**.

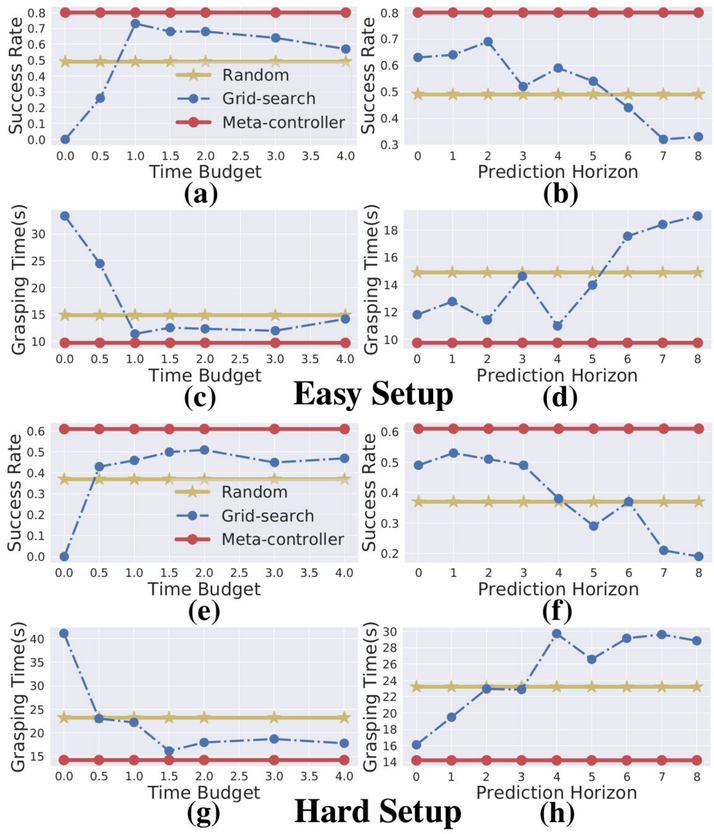

Qualitative Results

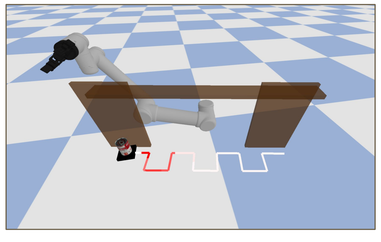

Meta-Controller outputs prediction horizon to constrain the predicted target pose within the **highly reachable space.**