NRI: FND: Better robotic manipulation using state and action abstraction

PI: Robert Platt, Northeastern University
https://www.ccs.neu.edu/home/rplatt/

Problem: action abstraction in robotic manipulation environments with large action spaces.

- For example, robotic pick and place with each pick/place position is represented as a separate action.
- Better action abstraction leads to better transfer to new tasks and higher sample efficiency -- major challenges in reinforcement learning.

Our method distills the policies of experts for N training tasks and uses the gained information to learn a policy for a novel testing task quickly.

- An action prior suggest actions plausible under any expert.
- A task classifier decides if an expert is applicable in a given state (perhaps the state is outside the expert's domain).
- A model-free agent uses the action prior exploration policy to discover the goal state of the testing task.

We identify three potential **areas of impact:**

1. Reinforcement learning: better exploration policies, such as curiosity-driven exploration, are highly sought after.
2. Stochastic planning: efficient search algorithms prune action spaces in the search tree.
3. Imitation learning: modeling the expert can lead to higher sample efficiency.

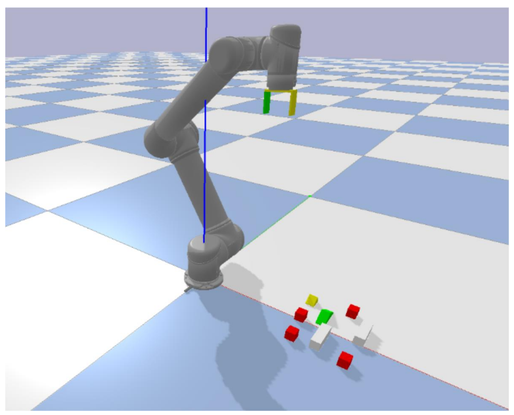

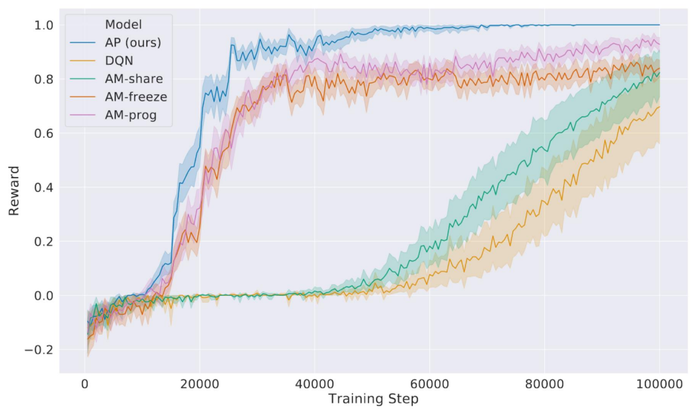

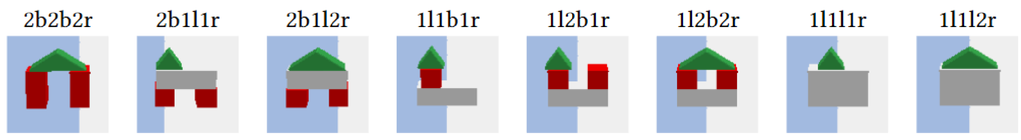

Left: our method can solve 14/16 block building tasks, where a heuristic baseline can solve only 4/16.

Right: our method (AP) outperforms Actor-Mimic (a transfer learning method) baselines in a proof-of-concept fruit picking domain.

Bottom: examples of structures in the block building domain.

2021 NRI & FRR Principal Investigators' Meeting
March 10-12, 2021

Award ID#: 80NSSC19K1474