



Reinforcement Learning via Auxiliary Task Distillation

Consider the task of using a robot to rearrange an object in the house


- Fetch-Robot with 10-DOF and a suction gripper
- Requires diverse skills like Navigating, Opening a cabinet, Picking up, and Placing

Can long-horizon robot control be learnt end-to-end without using demonstrations or a curriculum?




Yes, by using Auxiliary Tasks!

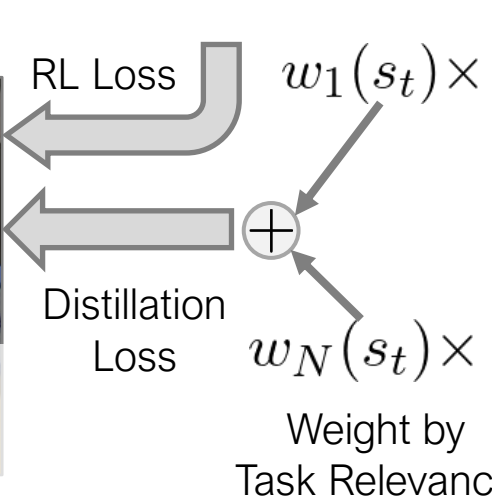

- Auxiliary tasks carry relevant behaviors which are easier to learn and transferred to the main task
- They are learnt simultaneously along with the main task



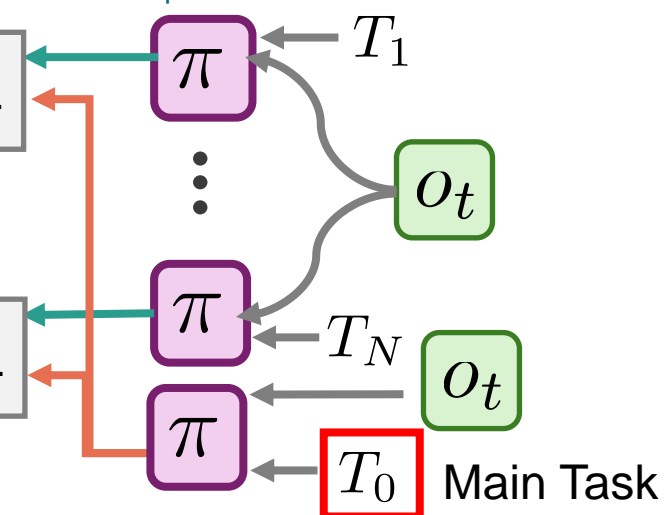
2. Collect Rollouts



3. Multi-Task RL loss across tasks



4. Compute Distillation Loss



Results

Outperforms a variety of end-to-end and hierarchical baselines by 2.3x

Easy: Episodes in which the object is placed in an open receptacle

Hard: Object is placed inside a closed receptacle

- M3 (+24%) → Hierarchical RL with STRIPS planner with Navigate, Pick and Place skills
- Mono (+73%) → end to end RL which directly maps observations to actions
- GALA (+24%): Scaling end to end RL with kinematic simulation (2B samples: x4 more than Aux-Distill)
- ST (+25%) → Transformer architecture for rearrangement using demonstrations

