Tether

Autonomous Functional Play with Correspondence-Driven Trajectory Warping

William Liang, Sam Wang, Hung-Ju Wang, Osbert Bastani, Yecheng Jason Ma†, Dinesh Jayaraman†

tether-research.github.io

Method

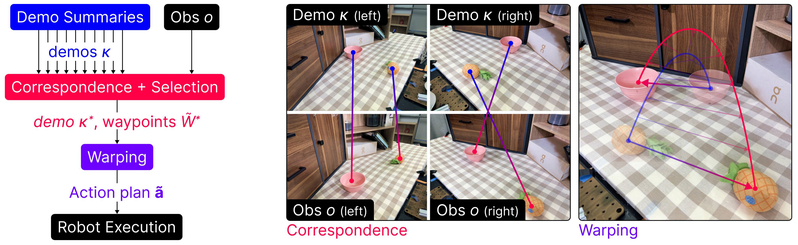

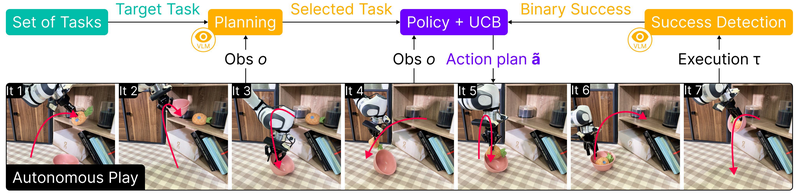

Running in a **continuous** VLM-guided multi-task cycle, Tether computes keypoint correspondences to warp demo trajectories for the current scene.

Experiments on Robust Imitation

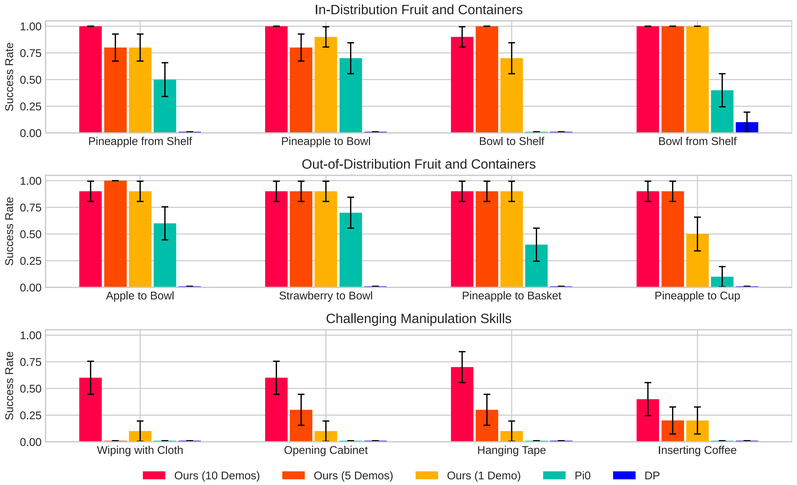

Tether's policy design is **sample efficient** and **robust** to spatial and semantic variation, surpassing imitation baselines in the low data regime.

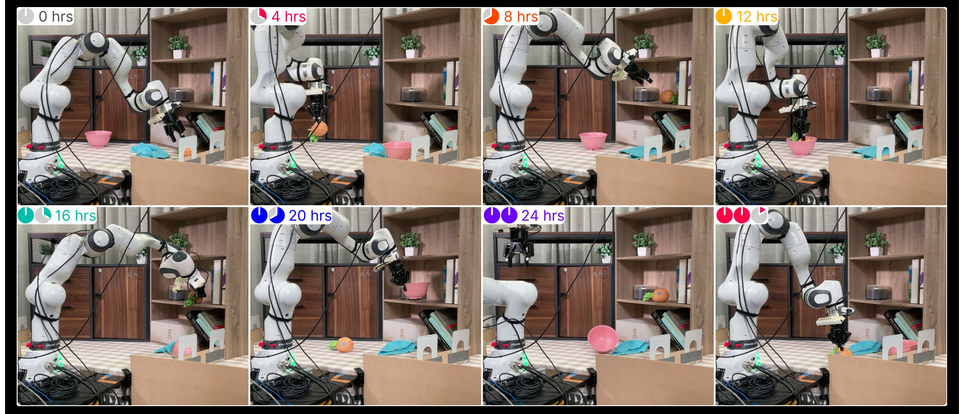

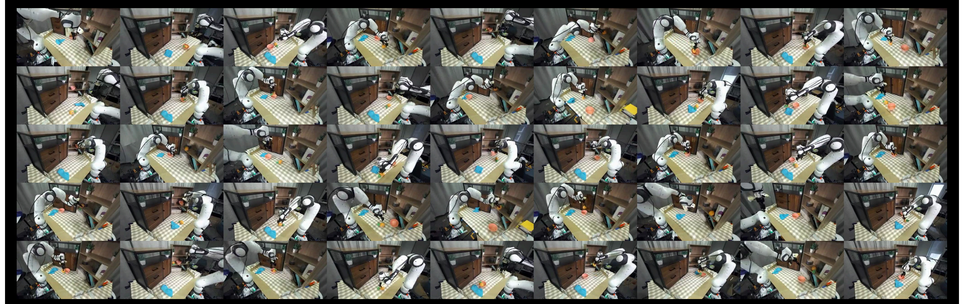

What happens if you let a robot play by itself... for over 24 hours?

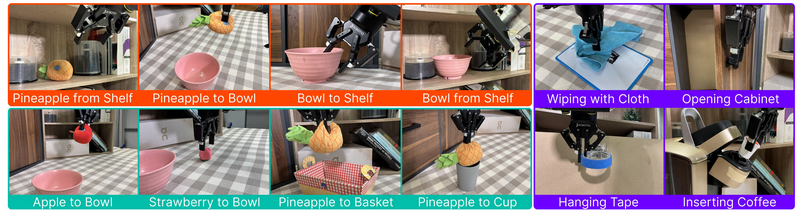

Experiments on Autonomous Play

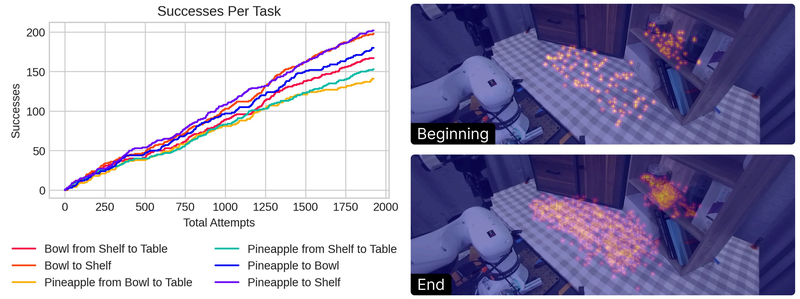

Tether plays for **26**+ hours and produces **1000**+ expert-level trajectories.

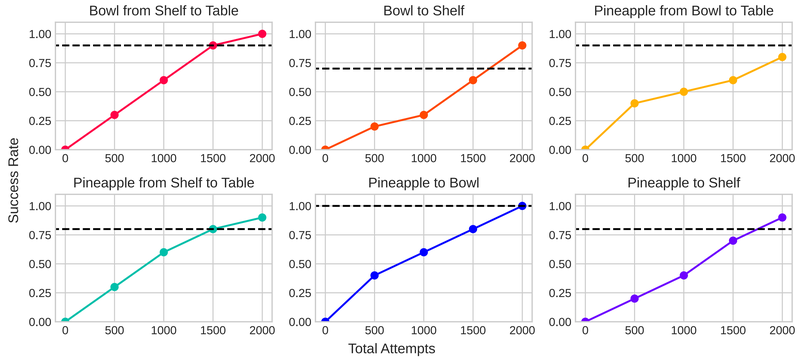

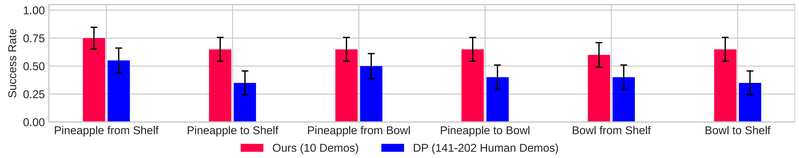

Tether's data stream **consistently improves** downstream policy training performance over time, with results comparable to human-collected data.

Penn
Berkeley
UNIVERSITY OF CALIFORNIA