

Figure 1. Real-world success rates during the course of online fine-tuning for puck pushing.




Figure 2. Real-world success rates during the course of online fine-tuning for cloth pushing.




Figure 3. Real-world success rates during the course of online fine-tuning for ball pushing.




Figure 4. Real-world success rates during the course of online fine-tuning for puck pushing ablated over using domain randomization during training in simulation.




Figure 5. We plot normalized rewards for sim-2-sim fine-tuning task of walker walk, where rewards are normalized by the maximum reward achieved by any method.




Figure 6. We plot normalized rewards for sim-2-sim fine-tuning task of cheetah run, where rewards are normalized by the maximum reward achieved by any method.





Figure 7. We plot normalized rewards for sim-to-sim fine-tuning task of rope peg-in-hole, where rewards are normalized by the maximum reward achieved by any method.




Figure 8. We ablate over the choice of H and plot normalized rewards for sim-to-sim fine-tuning task of walker walk, where rewards are normalized by the maximum reward achieved by any method.




Figure 9. We ablate over the choice of H and plot normalized rewards for sim-2-sim fine-tuning task of rope peg-in-hole, where rewards are normalized by the maximum reward achieved by any method.




Figure 10. We plot simulation training curves for SAC and TDMPC2 for hammering task showing that the implementations are correct and able to learn eventually



Figure 11. Execution of real-world fine-tuning of a squishy toy ball



Figure 12. Execution of real-world fine-tuning of a towel




Figure 13. Execution of sim-to-sim Rope Peg-in-Hole task. The end effector is holding a rope attached to a peg, with the goal of getting the peg into the hole.




Figure 14. Real-world success rates during the course of online fine-tuning for insertion.




Figure 15. Real-world success rates during the course of online fine-tuning for hammering.