

Figures for Rebuttal (Paper ID: 5852, Recovering from Out-of-sample States via Inverse Dynamics in Offline Reinforcement Learning)

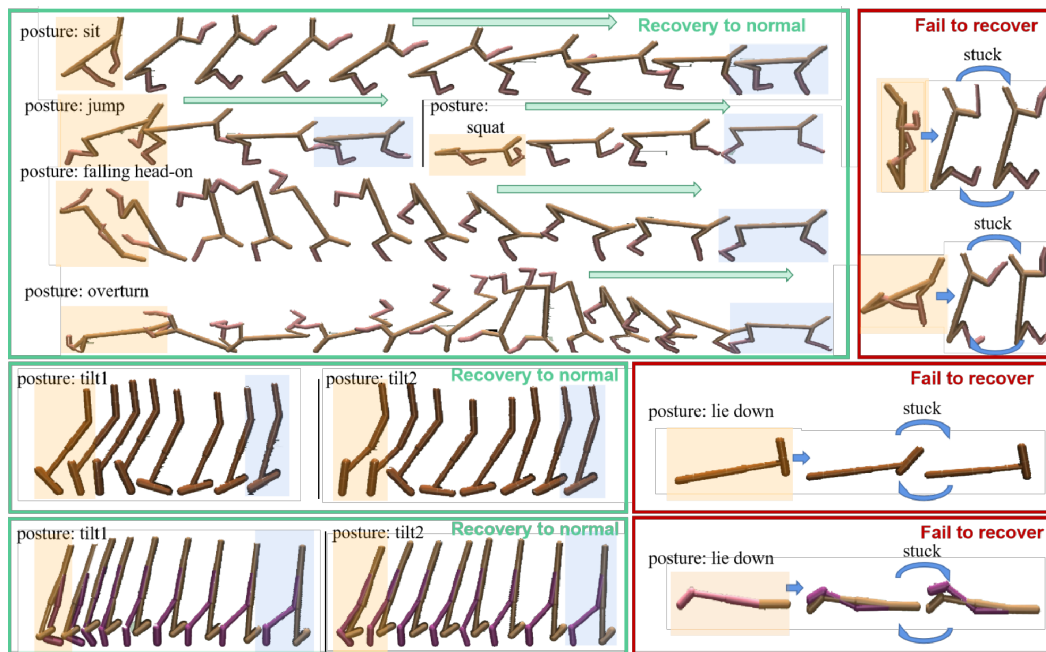


Figure 1: The visualization of OSR’s state recovery process on the real-world out-of-sample (OOS) Halfcheetah, Hopper and Walker2d benchmarks, where the initial states are generated by the modified generative adversarial network. The states in orange boxes represent the generated OOS states, and those in blue boxes are the recovered normal states. The trajectories in the left green box are samples successfully recovered while in the right red box are samples failed to recover.

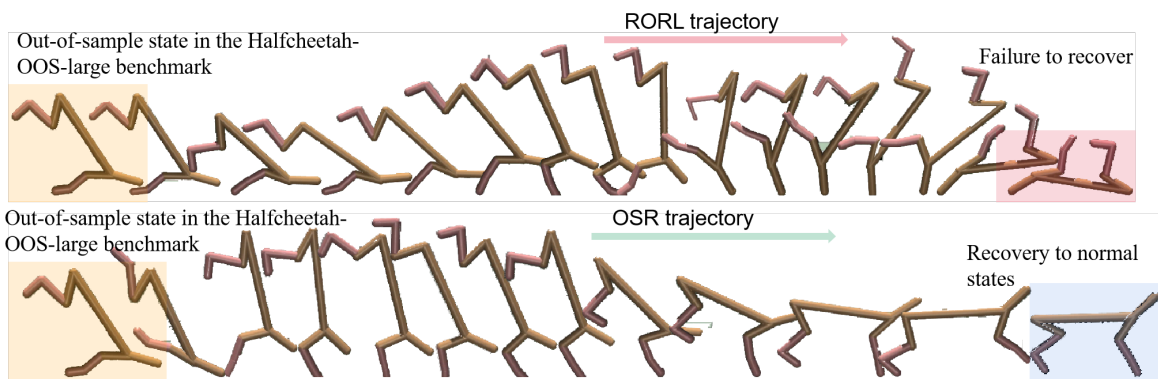


Figure 2: The visualization of OSR(ours) and RORL’s state recovery process on the Halfcheetah-OOS-large benchmark. The states in orange box are the out-of-sample state, the state in red box is the state the agent sticks in and the state in blue box is the normal state.