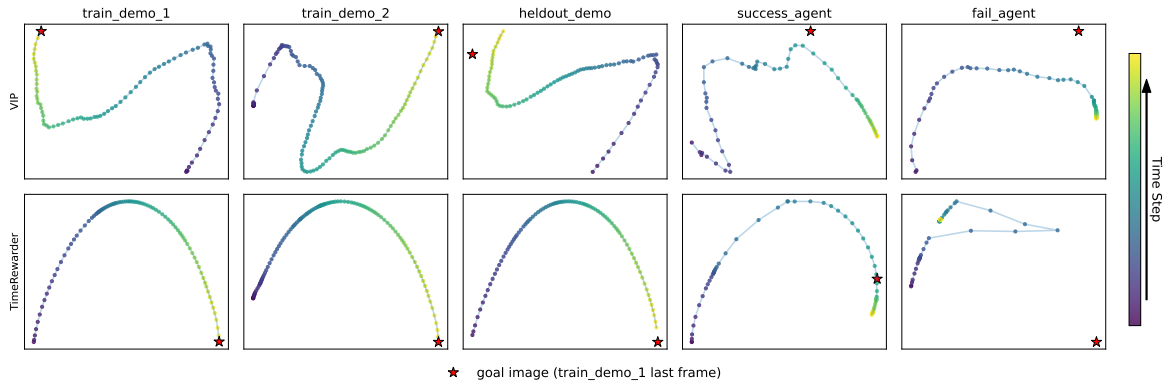

Figure 1. PCA visualization of VIP and TimeRewarder representations on trajectories of the `window-open` task. `train_demo_1` and `train_demo_2` are expert demonstrations from the training set; `heldout_demo` is a different expert demonstration used to evaluate generalization (relates to TimeRewarder paper Fig. 3). `success_agent` and `fail_agent` are trajectories from RL rollouts, illustrating the agent's successful and failed attempts (relates to TimeRewarder paper Fig. 4).

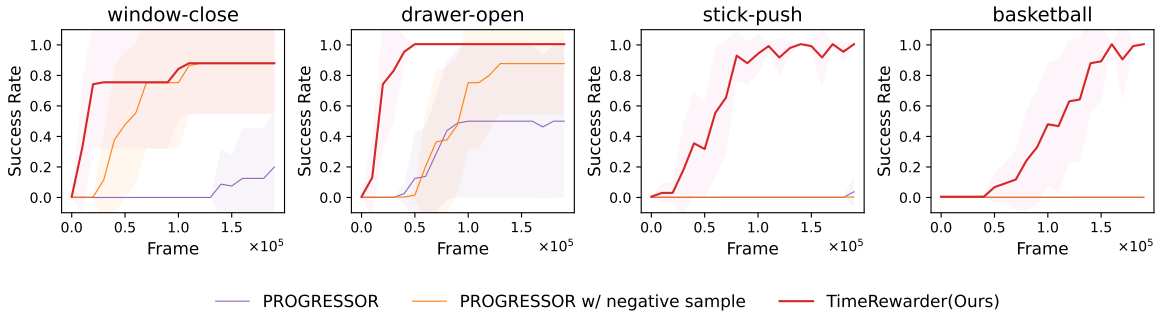

Figure 2. Ablation study of PROGRESSOR. Curves show mean ± s.d. over eight seeds. While adding negative sampling significantly improves performance on `window-close` and `drawer-open`, tasks like `basketball` and `stick-push` still underperform compared to TimeRewarder, since PROGRESSOR's structure does not naturally support antisymmetric (inverse) representation modeling.

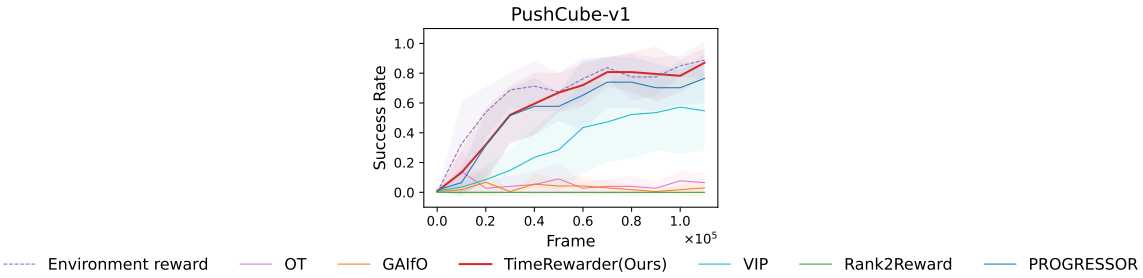

Figure 3. RL performance on `PushCube-v1` task in ManiSkill3 suite. Curves show mean ± s.d. over eight seeds.