

Figure 5: Normalized performance of CSIL against baselines for offline imitation learning for tasks. Uncertainty intervals depict quartiles over 10 seeds. To assess convergence across seeds, performance is chosen using the highest filtered 25th percentile of the episode return during learning.

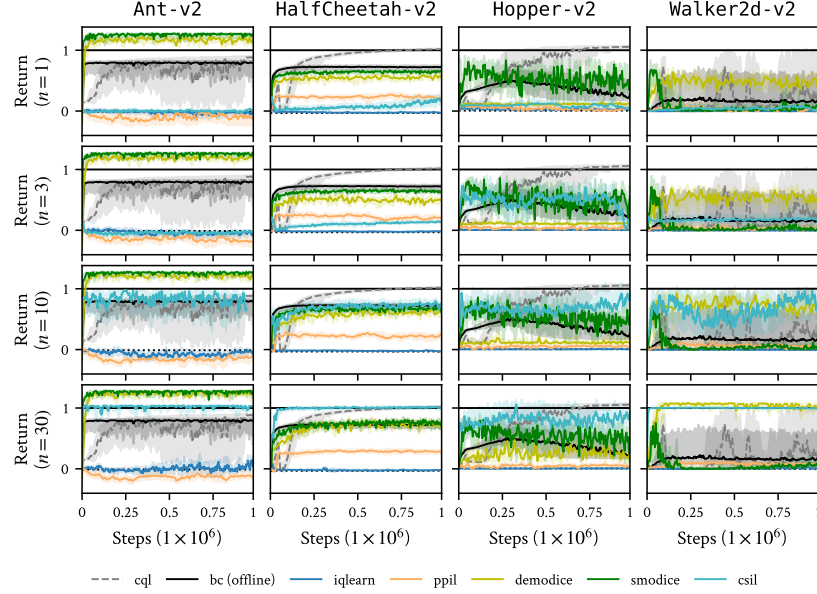


Figure 29: Normalized performance of CSIL against baselines for *offline* imitation learning for gym tasks . Uncertainty intervals depict quartiles over 10 seeds. n refers to demonstration trajectories.

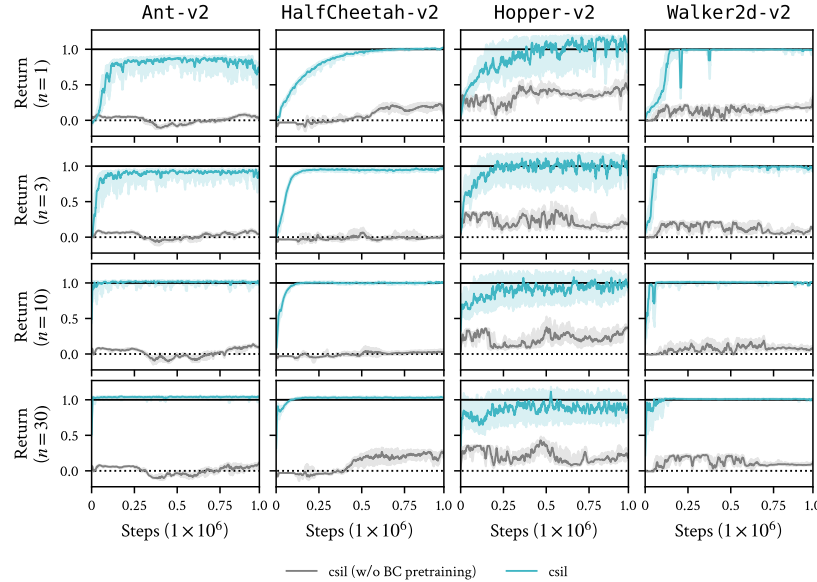


Figure 48: Online MuJoCo gym tasks where CSIL has no BC pretraining. Uncertainty intervals depict quartiles over 5 seeds. n refers to demonstration trajectories.