

Sanjiban Choudhury

Title: To RL or not to RL

Abstract: Model-based Reinforcement Learning (MBRL) and Inverse Reinforcement Learning (IRL) are powerful techniques that leverage expert demonstrations to learn either models or rewards. However, traditional approaches suffer from a computational weakness: they require repeatedly solving a hard reinforcement learning (RL) problem as a subroutine. This requirement presents a formidable barrier to scalability. But is the RL subroutine necessary? After all, if the expert already provides a distribution of “good” states, does the learner need to explore? In this work, we demonstrate a more informed MBRL and IRL reduction that utilizes the state distribution of the expert to provide an exponential speedup in theory. In practice, we significantly speed up the prior art on continuous control tasks.