

FiRST: Fine-tuning model-based RL for Similar-Task Transfer

Rosella Liu

May 2024

1 Abstract

Large models have inspired the need for knowledge transfer and continual learning, which are hampered by plasticity loss widely observed across a variety of models. This paper seeks to experiment and investigate into the possibility of knowledge transfer of a state-of-the-art model based reinforcement learning algorithm: TD- MPC, by explicitly reusing its world model. In contrast to our expectation, our experiments show that freezing world model tremendously undermines its fine-tuning performance. Despite this unforeseen result, it gives rise to several insights and questions into future directions of world model knowledge transfer: i) Should the states in the world model be exact observations or abstract latent states? ii) Should the world model be task oriented? iii) Given a pre-trained world model, how to effectively reuse and fine-tune its parameters? We wish this work will be helpful in designing world models that are portable to new tasks.