

Moving Closer to Real-world Reinforcement Learning

Yunjin Tang

Google Brain

Abstract

From super-human performance in games to ultra-efficient automatic chip layout design, we have witnessed the rise of deep reinforcement learning (RL) and its transition from proof-of-concept tasks to beneficial real world applications. In this presentation, I will show some research topics that are critical for us to further advance toward more and more real-world RL applications. The topics to be covered include model-based RL, which aims to address the data inefficiency and environmental uncertainty problems commonly encountered in deep RL, and offline RL techniques wherein we could only rely on the available dataset to learn policies when data collection is expensive. To shed light on interpreting agent behaviors, I will introduce works that combine attention with RL. Finally, manually designing tasks and tuning rewards to training policies can be exhausting, I will cover some works that focus on open-endedness RL that may train the desired policies and at the same time relieve us from the laborious task design.