



Deep Learning for Agents

Koray Kavukcuoglu
Google DeepMind, UK

Abstract

In recent years, deep learning has improved or replaced standard vision, speech and natural language processing pipelines on many supervised benchmark problems. At the same time, advances in deep reinforcement learning models made it possible to train agents that can solve complicated tasks in an end-to-end manner. However, one of the most challenging problems is for an agent to truly understand its environment without direct supervision or reward signal.

In my talk I will first introduce main approaches in deep reinforcement learning that can achieve state-of-the-art results on agent benchmarks. I will then focus on recent advances in generative models of images and videos. Lastly, I will discuss important future directions in deep learning and generative models for agents.