

Model-free control with deep learning

Marc Bellemare

Google DeepMind

In this talk I will present some recent developments in model-free reinforcement learning applied to large state spaces, with an emphasis on deep learning and its role in estimating action-value functions. The talk will cover a variety of model-free algorithms, including variations on Q-Learning, and some of the main techniques that make the approach practical. I will illustrate the usefulness of these methods with examples drawn from the Arcade Learning Environment, the popular set of Atari 2600 benchmark domains.