



ZHAOHAN DANIEL GUO

Directed Exploration for Improved Sample Efficiency in Reinforcement Learning

Tuesday, December 18, 2018 – 9:00 a.m. – GHC 6501

A key challenge in reinforcement learning is how an agent can efficiently gather useful information about its environment to make the right decisions, i.e., how can the agent be sample efficient. This thesis proposes using a new technique called directed exploration to construct new sample efficient algorithms for both theory and practice. Directed exploration involves repeatedly committing to reach specific goals within a certain time frame. This is in contrast to dithering which relies on random exploration or optimism-based approaches that implicitly explore the state space. Using directed exploration can yield provably efficient sample complexity in a variety of settings of practical interest: when solving multiple tasks either concurrently or sequentially, algorithms can explore distinguishing state--action pairs to cluster similar tasks together and share samples to speed up learning; in large, factored MDPs, repeatedly trying to visit lesser known state--action pairs can reveal whether the current dynamics model is faulty and which features are unnecessary. Finally, directed exploration can also improve sample efficiency in practice for the deep reinforcement learning by being more strategic than dithering-based approaches and more robust than reward-bonus based approaches.

Thesis Committee:
Emma Brunskill, Chair
Drew Bagnell
Ruslan Salakhutdinov
Remi Munos, Google DeepMind