



Zhaohan Daniel Guo

Improving Sample Efficiency for Reinforcement Learning through Smarter Exploration

Tuesday, January 16, 2018 – 9:00 a.m. – GHC 8102

This thesis proposes using more sophisticated exploration techniques to bring theory closer to practice for reinforcement learning algorithms. One technique, directed exploration, involves explicitly performing exploration for specific goals, which can be used to accumulate useful information that can narrow down the possibility space of unknown parameters. 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 for MDPs large and small, using data-dependent confidence intervals as a form of tempered optimism combined with explicit exploration towards gathering information about value gap between actions may result in more efficient, practical performance, along with tighter, problem-dependent bounds.

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