

Unsupervised Reinforcement Learning

Sergey Levine

Department of Electrical Engineering and Computer Science
University of California, Berkeley
svlevine@eecs.berkeley.edu

ABSTRACT

Conventionally, reinforcement learning algorithms are goal-directed: they aim to acquire policies that most effectively maximize a given reward signal. However, if we consider agents that must master very large repertoires of behaviors – such as general-purpose robots that must perform a diverse array of tasks in the real world – then it makes sense to instead frame the reinforcement learning process as an unsupervised learning procedure, which has the aim of extracting a large and diverse array of skills that can later be utilized for the many tasks that the agents may be asked to perform. Such a formulation not only makes it feasible to acquire diverse behaviors before any reward signal is actually observed, but can actually make learning much more tractable for tasks with delayed or sparse reward signals. In this talk, I will discuss recent advances in unsupervised reinforcement learning, many of which draw on an information-theoretic formulation for the unsupervised skill acquisition problem. I will discuss how this formulation can provide us with a principled view of unsupervised skill acquisition, and furthermore provides some tantalizing clues about how to quantify the usefulness of learned behaviors. I will also present experimental results showing that unsupervised reinforcement learning not only provides good results in a variety of simpler simulated environments, but in fact can be utilized with real-world robotic systems to learn sophisticated behaviors with minimal human input.

KEYWORDS

reinforcement learning, robotics

ACM Reference Format:

Sergey Levine. 2020. Unsupervised Reinforcement Learning. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*, Auckland, New Zealand, May 9–13, 2020, IFAAMAS, 1 pages.

BIOGRAPHY

Sergey Levine received a BS and MS in Computer Science from Stanford University in 2009, and a Ph.D. in Computer Science from Stanford University in 2014. He joined the faculty of the Department of Electrical Engineering and Computer Sciences at UC Berkeley in fall 2016. His work focuses on machine learning for decision making and control, with an emphasis on deep learning and reinforcement

learning algorithms. Applications of his work include autonomous robots and vehicles, as well as computer vision and graphics. His research includes developing algorithms for end-to-end training of deep neural network policies that combine perception and control, scalable algorithms for inverse reinforcement learning, deep reinforcement learning algorithms, and more.

