
Faster Learning for Better Decisions*

Emma Brunskill†

Abstract: A fundamental goal of artificial intelligence is to create agents that learn to make good decisions as they interact with a stochastic environment. Some of the most exciting and valuable potential applications involve systems that interact directly with humans, such as intelligent tutoring systems or medical interfaces. In these cases, sample efficiency is highly important, as each decision, good or bad, is impacting a real person. I will describe our research on tackling this challenge, including transfer learning across sequential decision making tasks, as well as its relevance to improving educational tools.

Bio: Emma Brunskill is an assistant professor in the computer science department at Carnegie Mellon University. She is also affiliated with the machine learning department at CMU. She works on reinforcement learning, focusing on applications that involve artificial agents interacting with people, such as intelligent tutoring systems. She is a Rhodes Scholar, Microsoft Faculty Fellow and NSF CAREER award recipient, and her work has received best paper nominations in Education Data Mining (2012, 2013) and CHI (2014).

*Machine Learning External Seminar, Gatsby Unit, March 18 2015.

†<http://www.cs.cmu.edu/~ebrun/>