

Building Knowledge for AI Agents with Reinforcement Learning

Doina Precup
McGill University and DeepMind
Montreal, Canada
dprecup@cs.mcgill.ca

ABSTRACT

Reinforcement learning allows autonomous agents to learn how to act in a stochastic, unknown environment, with which they can interact. Deep reinforcement learning, in particular, has achieved great success in well-defined application domains, such as Go or chess, in which an agent has to learn how to act and there is a clear success criterion. In this talk, I will focus on the potential role of reinforcement learning as a tool for building knowledge representations in AI agents whose goal is to perform continual learning. I will examine a key concept in reinforcement learning, the value function, and discuss its generalization to support various forms of predictive knowledge. I will also discuss the role of temporally extended actions, and their associated predictive models, in learning procedural knowledge. Finally, I will discuss the challenge of how to evaluate reinforcement learning agents whose goal is not just to control their environment, but also to build knowledge about their world.

CCS Concepts/ACM Classifiers

- Computing methodologies - Reinforcement learning

Author Keywords

Reinforcement learning; knowledge representation

BIOGRAPHY

Doina Precup splits her time between McGill University/Mila, where she holds a Canada-CIFAR AI chair, and DeepMind, where she leads the Montreal research team formed in 2017. Her research focuses on reinforcement learning, deep learning, time series analysis, and diverse applications of machine learning with a special focus on health care. She completed her BSc/Eng (1994) degree in computer science at the Technical University Cluj-Napoca, Romania, and her MSc (1997) and PhD (2000) degrees at the University of Massachusetts, Amherst, where she was a Fulbright scholar. She became a senior member of AAAI in 2015,

a Canada Research Chair in 2016 and a Senior Fellow of CIFAR in 2017. Doina Precup is also involved in organizing activities aimed at increasing diversity in machine learning, such as the AI4Good summer lab and the Eastern European Machine Learning school.

