AI SEMINAR SERIES

Reductionism in Reinforcement Learning

Dale Schuurmans
Professor,
University of Alberta
Research Scientist,
Google Brain

**WEDNESDAY
MAY 5
1-2 P.M. PACIFIC**

Zoom
tinyurl.com/392tj3hk
Calendar link
tinyurl.com/y3kahh9s

ABSTRACT

Learning to make sequential decisions from interaction raises numerous challenges, including temporal prediction, state and value estimation, sequential planning, exploration, and strategic interaction. Each challenge is independently difficult, yet reinforcement learning research has often sought holistic solutions that rely on unified learning principles. I will argue that such a holistic approach can hamper progress. To make the point, I focus on reinforcement learning in Markov Decision Processes, where the challenges of value estimation, sequential planning, and exploration are jointly raised. By eliminating exploration from consideration, recent work on off-line reinforcement learning has led to improved methods for value estimation, policy optimization, and sequential planning. Taking the reduction a step further, I reconsider the relationship between value estimation and sequential planning, and show that a unified approach faces unresolved difficulties whenever generalization is considered. Instead, by separating these two challenges, a clearer understanding of successful value estimation methods can be achieved, while the shortcomings of existing strategies can be overcome in some cases. I will attempt to illustrate how these reductions allow other areas of machine learning, optimization, control and planning to be better leveraged in reinforcement learning.

SPEAKER BIOGRAPHY

Dale Schuurmans is a Research Scientist at Google Brain, Professor of Computing Science at the University of Alberta, a Canada CIFAR AI Chair, and a Fellow of AAAI. He has served as an Associate Editor in Chief for IEEE TPAMI, an Associate Editor for JMLR, AIJ, JAIR and MLJ, and as a Program Co-chair for AAAI-2016, NIPS-2008 and ICML-2004. He has worked in many areas of machine learning and artificial intelligence, including model selection, on-line learning, adversarial optimization, boolean satisfiability, sequential decision making, reinforcement learning, Bayesian optimization, semidefinite methods for unsupervised learning, dimensionality reduction, and robust estimation. He has published over 200 papers in these areas, and has received paper awards at NeurIPS, ICML, IJCAI, and AAAI.

Accommodations for disabilities may be made by contacting patravaj@oregonstate.edu.

Oregon State University