

INFORMS Data Mining Society Online Seminar Series

11:30 – 12:30 US EST | February 24, 2023 Zoom Meeting Registration <u>https://us06web.zoom.us/webinar/register/WN_3trxS5NER76N697EsX1Bcw</u>

From Reinforcement Learning to Sequential Decision Analytics: Toward a Universal Framework for Sequential Decision Problems

Sequential decision problems are an almost universal problem class, spanning dynamic resource allocation problems, control problems, optimal stopping/buy-sell problems, active learning problems, as well as two-agent games and multiagent problems. Application settings span engineering, the sciences, transportation, health services, medical decision making, energy, e-commerce and finance, but in this talk I will emphasize applications in transportation and logistics.

These problems have been addressed in the research literature using a variety of modeling and algorithmic frameworks, including (but not limited to) dynamic programming, stochastic programming, stochastic control, simulation optimization, stochastic search, approximate dynamic programming, reinforcement learning, model predictive control, and even multiarmed bandit problems.

I will present a universal modeling framework that can be used for any sequential decision problem in the presence of different sources of uncertainty. I use a "model first" strategy that optimizes over policies for making decisions. I will present four (meta)classes of policies that are the foundation of any solution approach that has ever been proposed for a sequential problem, either in the research literature or used in practice (including policies that have not been invented yet).

I will close by making the case for teaching sequential decision analytics at both the undergraduate and graduate levels, including to students in fields centered on applications as well as methodology.



Speaker

Warren B Powell is Professor Emeritus at Princeton University, where he taught for 39 years, and is currently the Chief Innovation Officer at Optimal Dynamics. He was the founder and director of CASTLE Lab, which focused on stochastic optimization with applications to freight transportation, energy systems, health, e-commerce, finance and the laboratory sciences, supported by over \$50 million in funding from government and industry. He has pioneered a new universal framework that can be used to model any sequential decision problem, including the identification of four classes of policies that spans every possible method for making decisions. This is documented in his latest book with John Wiley: Reinforcement Learning and Stochastic Optimization: A unified framework for sequential decisions. He published over 250 papers, four books, and produced over 60 graduate students and post-docs. He is the 2021 recipient of the Robert Herman Lifetime Achievement Award from the Society for Transportation Science and Logistics, the 2022 Saul Gass Expository Writing Award. He is a fellow of Informs, and the recipient of numerous other awards.