



A Unified Framework for Sequential Decisions under Uncertainty



Sequential decisions are an almost universal problem class, spanning dynamic resource allocation problems, control problems, 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. A rich problem class involves systems that must actively learn about the environment, which arises in disease mitigation.

These problems have been addressed in the academic 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. The speaker is going to introduce a universal modeling framework that can be used for any sequential decision problem in the presence of different sources of uncertainty.

The speaker and his collaborators use a “model first” strategy that optimizes over policies for making decisions. They claim that there are four (meta)classes of policies that are the foundation of any solution approach that has ever been proposed for a sequential problem. Using a simple energy storage problem, the speaker shows that any of the four classes of policies might work best depending on the data.

Warren B. Powell, Ph.D.

Professor Emeritus, Princeton University
Chief Analytics Officer – Optimal Dynamics

Date: Monday, January 24, 2022

Time: 1:50 – 2:40 p.m. US Central Time

Zoom Meeting ID: 998 4499 3279

Passcode: 724615

Faculty host: Le Xie, ECEN

Biography

Dr. Warren B. Powell is Professor Emeritus at Princeton University, where he taught for 39 years, and is currently the Chief Analytics Officer at Optimal Dynamics. He is the founder and director of CASTLE Labs, which spans contributions to models and algorithms in stochastic optimization, with applications to energy systems, transportation, health, e-commerce, and the laboratory sciences (see www.castlelab.princeton.edu). He has pioneered the use of approximate dynamic programming for high-dimensional applications, and the knowledge gradient for active learning problems. His recent work has focused on developing a unified framework for sequential decision problems under uncertainty, spanning active learning to a wide range of dynamic resource allocation problems. He has authored books on Approximate Dynamic Programming and (with Ilya Ryzhov) Optimal Learning, and is nearing completion of a book Reinforcement Learning and Stochastic Optimization: A unified framework for sequential decisions.

You can also click this link to join the seminar <https://tamu.zoom.us/j/99844993279?pwd=TkJodWFVRURyMmkwakl4SWZGeVJTQT09>