ORF 411

Sequential Decision Analytics and Modeling

Fall, 2018

Course description: Students will learn how to model and solve general sequential decision problems under uncertainty. Sequential decision problems have five major components: states, decisions, exogenous information, transition function and objective function. Decisions are made using policies, and students will learn a structured approach to developing effective policies.

The course will be organized around a series of problems that are designed to bring out specific modeling features. Some examples include active learning of belief models, lagged decision processes, and dynamic forecasts. We will also address uncertainty modeling and quantification, and illustrating both policy search and lookahead policies. The figure on the next page lists the problems, while the columns highlight the modeling issues for each problem, and the types of policies that will be considered.

Key modeling elements will include:

- Modeling physical, informational and belief states; forecasts.
- Scalar and vector-valued decisions; static and lagged decisions.
- Uncertainty modeling and quantification.
- Model-based and model-free systems; belief transitions.
- Modeling performance metrics; multiobjectives; risk.
- The four classes of policies and policy evaluation.

Instructor: Warren B. Powell (powell@princeton.edu) – Sherrerd Hall Rm 230
Office hours: By appointment (please send request via email).

Teaching assistants: Grace Lee <gylee@princeton.edu>
Yongyi Guo <yongyig@princeton.edu>

Web site: All course materials will be made available on Blackboard.

Readings: The course will be centered on a new book, *Sequential Decision Analytics and Modeling* which will be available on Overleaf.

Prerequisites: Students are expected to have a background in basic probability and statistics (ORF 245 or ORF 309). While courses such as ORF 311, 360 and 418 are not prerequisites, this course will extend the type of material covered in these courses.

Requirements:

Problem sets (including attendance):	40 %
Midterm:	25 %
Take-home final:	35 %

ORF 411
Sequential Decision Analytics
Course outline
2018

Sept 12 – Sample applications
Canonical framework
Asset selling illustration

Sept 17 – Adaptive market planning – Policy function approximations
Derivative-based search
Stochastic gradients as a sequential decision problem

Sept 19 - Recursive statistical modeling
Lookup tables, parametric models, (briefly nonparametrics)

Sept 24 - Learning diabetes medications – Cost function approximations
Derivative free stochastic search
Stepsize policies

Sept 26 - Derivative-free stochastic search – direct lookahead policies
Thompson sampling
Knowledge gradient

Oct 1 - Multistep lookahead policies
Decision trees
Learning with a beta-bernoulli model
Stochastic shortest paths – Bellman's equation

Oct 3 - Stochastic shortest paths with observed costs
Bellman's equation with expanded state variable
The post-decision state variable
Approximate dynamic programming

Oct 8 - Dynamic shortest paths
Direct lookahead policies

Oct 10 - Designing policies
Policy search vs. lookahead policies
The four classes of policies

Oct 15 - Energy storage I
Basic model
A policy function approximation for energy arbitrage
Derivative-based policy search

Oct 17 - Energy storage II
A policy function approximation
A cost function approximation
Dynamic programming for energy storage:
Backward MDP
Backward ADP
Forward ADP
A direct lookahead
A parameterized direct lookahead

Oct 22 - Midterm review
Air Liquide case study

Oct 24 - Midterm

Fall break

Oct 29 - Energy systems modeling – Hugo Simao
Model of the PJM energy grid
Lagged decision processes
Modeling wind
Lookahead policy

Nov 31 - Modeling uncertainty
Empirical distributions
From general to normal and back
Time series models
Mean reversion
Jump diffusion

Hidden semi-Markov models

Nov 5 - Two-agent newsvendor problem
The two-agent newsvendor game
Basic model
Learning policies

Nov 7 - The beer game (in class game)

Nov 12 - Supply chain management for Pratt & Whitney
A parameterized lookahead policy
Policy search
Types of uncertainty

Nov 19 - Student decision problems
Modeling
Designing policies

Nov 21 – Thanksgiving break

Nov 26 - Ad-click optimization
Basic model
Bidding policies
Online learning
The ad-click game

Nov 28 - Student decision problems
Modeling
Designing policies

Dec 3 - Optimizing clinical trials
Basic model
A simple lookahead policy
A stochastic lookahead policy
An optimal policy

Dec 5 - Revenue management for hotels
Basic model
A lookahead policy

Dec 10 - Blood management problem
Basic model
Designing policies
Approximate dynamic

Dec 12 - Student decision problems