Stats 115: Probabilistic Decision Making

Lectures

When : Winter 2020, Tuesdays & Thursdays, 2pm – 3:15pm
Where : Royce 162

Sections

When : Wednesdays, 5-5:50pm
Where : WGYOUNG 1044

Instructor Info

Professor : Tao Gao (Dept. of Statistics)
Office: : Math & Science Building 8917
Email : tao.gao@stat.ucla.edu (***must include " [Stats 115] " in the title of your emails***)
Office Hours : Tuesday 3:30 – 4:30pm, or by appointment

Teaching Assistant

TA : Kaiwen Jiang
Office: : Math & Science Building 8349
Email : kaiwenj@g.ucla.edu
Office Hours : Monday 5-6pm

Course Description

Decision matters. At every moment, numerous decisions are made by humans and machines. What is a "rational" decision ? How to make a rational decision when information is incomplete, when the future is uncertain ? Probabilistic decision making offers a unified framework for solving these challenges. Learn how to build models that can infer, predict and act.

Justification

Making decisions is at the core of both cognitive sciences and artificial intelligence. Yet most current courses on probability focus on analyzing data, testing hypotheses, instead of making decisions and taking actions. This course will introduce the mathematical foundation of rational decision making under uncertainty. It will reveal the deep connections between probabilistic inference and decision theory.

Course Objectives

(1) Formulating a decision making problem as probabilistic inference.
(2) Deriving algorithms that can make decisions under uncertainty.
(3) Implementing code that execute inference and decision.

Prerequisites

Taking this course without the recommended prerequisites is possible, but you will need my approval first.

- Introduction to Probability (Stats 100A)
- Experience with Python (or equivalent programming languages)

Grading

- Two Exams (50%)
- Eight weekly assignments (50%)

Late Policy

All assignments must be turned in via CCLE on time. We will allow a total of five late days cumulatively – no explanation required. We will not make any additional allowances for late assignments: the late days are intended to provide for exceptional circumstances, and students should avoid using them unless absolutely necessary. Any assignments that are submitted late (with insufficient late days remaining) will not be graded.

Topics & Tentative Schedule

Week	Date	Title
1	Jan 7	Introduction: Decision Under Uncertainty
	Jan9	Utility theory: Quantifying Rationality
2	Jan 14	Probability: Quantifying Uncertainty
	Jan 16	Making Simple Rational Decisions
3	Jan 21	Bayesian Probabilistic Inference
	Jan 23	The Value of Information
4	Jan 28	Bayesian Networks: Causality and Independence
	Jan 30	Exact Inference in Bayesian Networks
5	Feb 4	*Review and QA*
	Feb 6	**Mid-term Exam**
6	Feb 11	Decision Networks: Making Decisions with Bayesian Networks
	Feb 13	Markov Process: Probabilistic Reasoning Over Time
7	Feb 18	*Presidents' Day Break*
	Feb 20	Markov Decision Process
8	Feb 25	Value Iteration
	Feb 27	Q-learning: Model-free Reinforcement Learning
9	March 3	Monte-Carlo Tree Search: Model-based Planning
	March 5	Alpha-Zero: Integrating Model-based and Mode-free Methods
10	March 10	*Review and QA*
	March 12	**Final Exam**

Readings

- Artificial Intelligence: A Modern Approach (third edition), Chapters 13-17.
- Reinforcement Learning: An Introduction (second edition) http://incompleteideas.net/book/bookdraft2017nov5.pdf