

Statistical Techniques in Robotics (16-831) Fall 2017 REV06

Days TR

Room GHC 4215

Time 1:30 - 2:50 PM

Lecturer Kris Kitani

TAs Wen Sun, Sibi Venkatesan

Class Discussion and Slides <https://piazza.com/cmu/fall2017/16831/home>

Description

Probabilistic and learning techniques are now an essential part of building robots (or embedded systems) designed to operate in the real world. These systems must deal with uncertainty and adapt to changes in the environment by learning from experience. Uncertainty arises from many sources: the inherent limitations in our ability to model the world, noise and perceptual limitations in sensor measurements, and the approximate nature of algorithmic solutions. Building intelligent machines also requires that they adapt to their environment. Few things are more frustrating than machines that repeat the same mistake over and over again. We will explore (1) modern learning techniques that are effective at learning online: i.e. throughout the robots operation, (2) reinforcement learning based techniques built to learn from expert demonstrations and (3) probabilistic inference algorithms for maintaining an estimate of itself despite uncertainty.

Prerequisites

Linear Algebra, Multivariate Calculus, Probability theory

Grading

- (1) Assignments 60% (15% per assignment, 4 assignments total);
- (2) Final Project 40%.

Late Submissions

Submit before deadline 100% credit. Submit within 48 hours of deadline 50% credit. After 48 hours 0% credit but you must submit something to pass the class.

Educational Outcomes

- (1) Able to design, implement and prove the regret bounds of a Multiplicative Weights Algorithm.
- (2) Apply concepts of online learning and understand the theoretical implication to incremental supervised learning algorithms.
- (3) Implement and analyze the difference between linear programming, matrix games, quadratic programming and entropy maximization formulations of inverse reinforcement learning.
- (4) Implement and analyze the optimality of multi-armed bandit problems given various environments.

Academic Integrity

All encouraged to work together BUT you must do your own work (code and write up). If you work with someone, please include their name in your write up and inside any code that has been discussed. If we find highly identical write-ups or code without proper accreditation of collaborators, we will take action according to university policies.

Take care of yourself

Do your best to maintain a healthy lifestyle this semester by eating well, exercising, avoiding drugs and alcohol, getting enough sleep and taking some time to relax. This will help you achieve your goals and cope with stress. All of us benefit from support during times of struggle. You are not alone. There are many helpful resources available on campus and an important part of the college experience is learning how to ask for help. Asking for support sooner rather than later is often helpful. If you or anyone you know experiences any academic stress, difficult life events, or feelings like anxiety or depression, we strongly encourage you to seek support. Counseling and Psychological Services (CaPS) is here to help: call 412-268-2922 and visit their website at <http://www.cmu.edu/counseling/>. Consider reaching out to a friend, faculty or family member you trust for help getting connected to the support that can help.

Class Schedule

Online Experts

Date	Topic	Lecturer	Project (Released/Due/Graded)
08-29	Overview, sequential prediction	K. Kitani	
08-31	Prediction with Expert Advice	K. Kitani	
09-05	Randomized Greedy, Regret	K. Kitani	
09-07	Weight Majority Algorithm	K. Kitani	WMA (Released)

Online Classification

09-12	Perceptron	K. Kitani	
09-14	Winnow Algorithm	K. Kitani	

Online Convex Optimization

09-19	Convex Optimization	K. Kitani	WMA (Due)
09-21	FTL, FTRL and OGD	K. Kitani	

Incremental Supervised Learning

09-26	MLPs and SGD	K. Kitani	WMA (Graded)
09-28	Online SVM, AdaBoost	K. Kitani	Super (Released)
10-03	1-2-1 Final Project meetings		

Final Project Pitch

10-05	Elevator Pitch	All	
-------	----------------	-----	--

Reinforcement Learning I

10-10	Reinforcement Learning	D. Held	
10-12	IRL as LP, IRL as Matrix Game	K. Kitani	Super (Due)
10-17	IRL as QP, IRL as MaxEnt	K. Kitani, N. Rhinehart	Bandit (Released)
10-19	Multi-Armed Bandit	S. Venkatesan	Super (Graded)
10-24	Imitation Learning	W. Sun	
10-26	Buffer	TBD	

Mid-Term Presentations

10-31	Presentations	All	Bandit (Due)
11-02	Presentations	All	

Reinforcement Learning II

11-07	Policy Gradient Methods	D. Held	Bandit (Graded)
11-09	Robot Learning	D. Held	IRL (Released)

State Estimation and SLAM

11-14	Temporal State Models, HMM	K. Kitani	
11-16	Kalman Filtering, EKF, MonoSLAM	K. Kitani	
11-21	NO CLASS		
11-23	NO CLASS (Thanksgiving)		

Final Project Presentations

11-28	NO CLASS		
11-30	NO CLASS		
12-05	Presentations	All	
12-07	Presentations	All	