

16-831: Statistical Approaches to Robotics

Machine Learning and Probabilistic Approaches to Uncertainty and Adaptation

Instructor: Drew Bagnell, NSH 3111, (412) 681-8669, dbagnell@ri.cmu.edu
Office Hours: Thursdays, 8:00-9:00am and by appointment

Teaching Assistant: Boris Sofman, EDSH 232, (412) 268-3402, (bsofman@ri.cmu.edu)
Office Hours: Tuesdays, 1:30-2:30pm and by appointment

Time and Day: Fall 2008, Tuesdays / Thursdays, 12-1:20pm, NSH 1305

Why?

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'll explore modern learning techniques that are effective at learning **online**: i.e. throughout the robots operation. We'll explore how the twin ideas of uncertainty and adaptation are closely tied in both theory and implementation.

What? (Things we will cover)

Fundamentals of Uncertainty (A new look at probability and machine learning)


- Cox Axioms
- Maximum Entropy
- Online algorithms, regret minimizing

Filtering and Localization

- Importance Sampling
- Particle Filtering, Condensation, Go-with-the-winners
- Monte-Carlo Localization and Mapping
- Rao-Blackwellization, FastSLAM
- Kalman Filtering, Unscented and Sigma Point Filtering
- Assumed Density Filtering and Expectation Propagation

Online Learning

- Practical, powerful and linear-time learners
- Bayes rule as an online algorithm
- Self-supervision (benefits and pitfalls)
- Robustness, Dealing with decision making (costs, change of measure)
- Online convex programming, maximum margin planning
- Feature selection




Random Fields and Factor Graphs

- Domain knowledge and graphical models
- Efficient Inference
- Supporting surfaces detection and super-resolution

Planning/Decision making under Uncertainty

Value-functions and stochastic planning
Partially Observed Markov Decision Processes
Policy search algorithms for practical reinforcement learning

Value of information and active learning
Online Learning Algorithms with Multi-agent Guarantees
Correlated Equilibrium



Who?

This course is directed to students interested in developing robust and adaptive software that interacts with the world. Although much of the material will be driven by applications within mobile robotics, anyone interested in applications of learning and probabilistic techniques, mathematical techniques for operating under uncertainty, or an interest in building complex adaptive systems is welcome.

Prerequisites

As an advanced course, familiarity with basic ideas from probability (joint and conditional distributions, Bayes rule, likelihood, etc.) and familiarity with basic ideas in machine learning will all be helpful. As the course will be project driven, prototyping skills including C, C++, and/or Matlab will also be important. Creative thought and enthusiasm are required.

How?

The course will include a mix of homework assignments that exercise the techniques we study, quizzes to demonstrate proficiency with the theoretical tools, and a strong emphasis on a significant research project.

Textbooks

Textbook: [*Probabilistic Robotics*](#), Sebastian Thrun, Wolfram Burgard, Dieter Fox
Optional Textbook: [*Pattern Recognition and Machine Learning*](#), Chris Bishop
Optional Textbook: [*Gaussian Processes for Machine Learning*](#), Carl Rasmussen, Christopher Williams

Additional readings will be posted on the course website.



Grading

Final grades will be based on the homeworks (30%), midterm (15%), final project (30%), and final exam (25%)

Late homework policy:

You will be allowed 2 total late days without penalty for the entire semester. Once those days are used, you will be penalized according to the following policy:

- Homework is worth full credit at the beginning of class on the due date
- It is worth half credit for the next 48 hours.
- It is worth zero credit after that.

You must turn in all homework, even if for zero credit.

Collaboration on homeworks:

Unless otherwise specified, homeworks will be done individually and each student must hand in their own assignment. It is acceptable, however, for students to collaborate in figuring out answers and helping each other understand the underlying concepts. You must write on each homework the names of the students you collaborated with.

Project

Projects may be done in groups of up to 3 students. The project is an opportunity to make a significant exploration into the application of ideas from the course to a robotics problem. More information to follow.

Exams

There will be a midterm exam and a final exam. Both will be open book and open notes (no computers allowed).

Auditing

If you do not wish to take the class for credit, you must register to audit the class. To satisfy the auditing requirement, you must either:

- Do two homework assignments, at least one of which must be one of the homeworks requiring the implementation of algorithms discussed in class
- Take the final exam and get at least 50% of the points
- Do a class project