

CS 287: Advanced Robotics Fall 2009

Lecture 1: Introduction

Pieter Abbeel
UC Berkeley EECS

WWW

- <http://www.cs.berkeley.edu/~pabbeel/cs287-fa09>

University of California at Berkeley
Dept of Electrical Engineering & Computer Sciences
CS 287: Advanced Robotics, Fall 2009

Instructor: Pieter Abbeel
Lectures: Tuesdays and Thursdays, 12:30pm-2:00pm, 405 Soda Hall
Office Hours: Thursdays 2:00-3:00pm (and by email arrangement) in 746 Suter/Dai Hall

Announcements

- [Announcements](#)
- [Assignments](#)
- [Course description](#)
- [Prerequisites](#)
- [Grading](#)
- [Assignment policy](#)
- [Syllabus and materials](#)
- [Related materials](#)

Announcements

- Communication:
 - Announcements: webpage
 - Email: pabbeel@cs.berkeley.edu
 - Office hours: Thursday 2-3pm + by email arrangement, 746 SDH
- Enrollment:
 - Undergrads stay after lecture and see me

Class Details

- Prerequisites:
 - Familiarity with mathematical proofs, probability, algorithms, linear algebra, calculus.
 - Ability to implement algorithmic ideas in code.
 - Strong interest in robotics
- Work and grading
 - Four large assignments (4 * 15%)
 - One smaller assignment (5%)
 - Open-ended final project (35%)
- Collaboration policy: Students may discuss assignments with each other. However, each student must code up their solutions independently and write down their answers independently.

Class Goals

- Learn the issues and techniques underneath state of the art robotic systems
- Build and experiment with some of the prevalent algorithms
- Be able to understand research papers in the field
 - Main conferences: ICRA, IROS, RSS, ISER, ISRR
 - Main journals: IJRR, T-RO, Autonomous Robots
- Try out some ideas / extensions of your own

Lecture outline

- Logistics --- questions? [textbook slide forthcoming]
- A few sample robotic success stories
- Outline of topics to be covered

Driverless cars

- Darpa Grand Challenge
 - First long-distance driverless car competition
 - 2004: CMU vehicle drove 7.36 out of 150 miles
 - 2005: 5 teams finished, Stanford team won
- Darpa Urban Challenge (2007)
 - Urban environment: other vehicles present
 - 6 teams finished (CMU won)
- Ernst Dickmanns / Mercedes Benz: autonomous car on European highways
 - Human in car for interventions
 - Paris highway and 1758km trip Munich -> Odense, lane changes at up to 140km/h; longest autonomous stretch: 158km

Kalman filtering, Lyapunov, LQR, mapping, (terrain & object recognition)

Autonomous Helicopter Flight

[Coates, Abbeel & Ng]



Kalman filtering, model-predictive control, LQR, system ID, trajectory learning

Four-legged locomotion


[Kolter, Abbeel & Ng]



inverse reinforcement learning, hierarchical RL, value iteration, receding horizon control, motion planning

Two-legged locomotion


[Tedrake +al.]



TD learning, policy search, Poincare map, stability

Mapping


[Video from W. Burgard and D. Haehnel]



"baseline" : Raw odometry data + laser range finder scans

Mapping

[Video from W. Burgard and D. Haehnel]



FastSLAM: particle filter + occupancy grid mapping

Mobile Manipulation

[Quigley, Gould, Saxena, Ng + al.]



SLAM, localization, motion planning for navigation and grasping, grasp point selection, (visual category recognition, speech recognition and synthesis)

Outline of Topics


- **Control:** underactuation, controllability, Lyapunov, dynamic programming, LQR, feedback linearization, MPC
- **Estimation:** Bayes filters, KF, EKF, UKF, particle filter, occupancy grid mapping, EKF slam, GraphSLAM, SEIF, FastSLAM
- **Manipulation and grasping:** force closure, grasp point selection, visual servo-ing, more sub-topics tbd
- **Reinforcement learning:** value iteration, policy iteration, linear programming, Q learning, TD, value function approximation, Sarsa, LSTD, LSPI, policy gradient, inverse reinforcement learning, reward shaping, hierarchical reinforcement learning, inference based methods, exploration vs. exploitation
- **Brief coverage of:** system identification, simulation, pomdps, k-armed bandits, separation principle
- **Case studies:** autonomous helicopter, Darpa Grand/Urban Challenge, walking, mobile manipulation.

1. Control

- Overarching theme: mathematically capture
 - What makes control problems hard
 - What techniques do we have available to tackle the hard problems
- E.g.: "Helicopters have underactuated, non-minimum phase, highly non-linear and stochastic (within our modeling capabilities) dynamics."
 - Hard or easy to control?

1. Control (ctd)

- Under-actuated vs. fully actuated
 - Example: acrobot swing-up and balance task



1. Control (ctd)

- Other mathematical formalizations of what makes some control problems easy/hard:
 - Linear vs. non-linear
 - Minimum-phase vs. non-minimum phase
 - Deterministic vs. stochastic
- Solution and proof techniques we will study:
 - Lyapunov, dynamic programming, LQR, feedback linearization, MPC

2. Estimation

- Bayes filters: KF, EKF, UKF, particle filter
- One of the key estimation problems in robotics: Simultaneous Localization And Mapping (SLAM)
- Essence: compute posterior over robot pose(s) and environment map given
 - (i) Sensor model
 - (ii) Robot motion model
- Challenge: Computationally impractical to compute exact posterior because this is a very high-dimensional distribution to represent
- [You will benefit from 281A for this part of the course.]

3. Grasping and Manipulation

- Extensive mathematical theory on grasping: force closure, types of contact, robustness of grasp
- Empirical studies showcasing the relatively small vocabulary of grasps being used by humans (compared to the number of degrees of freedom in the human hand)
- Perception: grasp point detection

4. Reinforcement learning

- Learning to act, often in discrete state spaces
- value iteration, policy iteration, linear programming, Q learning, TD, value function approximation, Sarsa, LSTD, LSPI, policy gradient, inverse reinforcement learning, reward shaping, hierarchical reinforcement learning, inference based methods, exploration vs. exploitation

5. Misc. Topics

- system identification: frequency domain vs. time domain
- Simulation / FEM
- Pomdps
- k-armed bandits
- separation principle
- ...

Reading materials

- Control
 - Tedrake lecture notes 6.832:
https://svn.csail.mit.edu/russt_public/6.832/underactuated.pdf
- Estimation
 - Probabilistic Robotics, Thrun, Burgard and Fox.
- Manipulation and grasping
 - -
- Reinforcement learning
 - Sutton and Barto, Reinforcement Learning (free online)
- Misc. topics
 - -

- Next lecture we will start with our study of control!