16-899C:
Adaptive Control and Reinforcement Learning

Machine Learning Techniques for Decision Making, Planning and Control

Instructor: Drew Bagnell (dbagnell@ri.cmu.edu) and Chris Atkeson (cga@ri.cmu.edu)

Time and Day: Spring, 2008, Tuesday and Thursday

Why?

Machine learning has escaped from the cage of perception. A growing number of state-of-the-art systems from field robotics, acrobatic autonomous helicopters, to the leading computer Go player and walking robots rely upon learning techniques to make decisions. This change represents a truly fundamental departure from traditional classification and regression methods as such learning systems must cope with a) their own effects on the world, b) sequential decision making and long control horizons, and c) the exploration and exploitation trade-off.

In the last 5 years, techniques and understanding of these have developed dramatically. One key to the advance of learning methods has been a tight integration with optimization techniques, and as such our case studies will focus on this.

What? (Things we will cover)

Planning and Optimal Control Techniques

- Differential Dynamic Programming
- Elastic Bands and Functional Optimization over the Space of Trajectories
- Iterative Learning Control

Imitation Learning

- Imitation Learning as Structured Prediction
- Imitation Learning as Inverse Optimal Control
- LEARning to searCH and Maximum Margin Planning
- Maximum Entropy Inverse Optimal Control
- Personally customized routing navigation

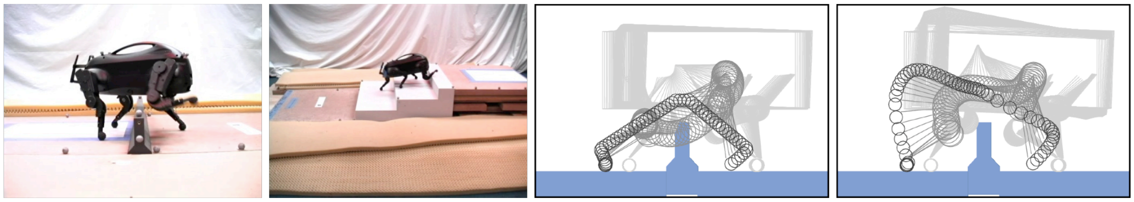

Reinforcement Learning and Adaptive Control

Exploration

- Bandit algorithms for limited feedback learning
- Contextual bandits and optimal decision making
    - "Sliding Autonomy" by contextual bandit methods
- Dual Control
    - "Bayesian" Reinforcement learning and optimal control for uncertain models
- "Unscented" linear quadratic regulation

Policy Search Methods

- Direct Policy Search Methods and Stochastic Optimization
    - Optimization of walking gaits and stabilizing controllers
- Conservative Policy Iteration
- Policy Search by Dynamic Programming
- REINFORCE and Policy Gradient Methods

Motion Planning

- Motion Planning that learns from experience
    - Trajectory libraries
    - Learning heuristics to speed planning

Design for Learnability

- Identifying feedback sources
- Modular learning design and structured problem
- Engineering insight as features and priors

Planning/Decision making under Uncertainty

Value-functions and stochastic planning
Partially Observed Markov Decision Processes and Information Space Planning
Belief Compression
Value of information and active learning

Who?

This course is directed to students—primarily graduate although talented undergraduates are welcome as well—interested in developing adaptive *software that makes decisions that affect the world.* Although much of the material will be driven by applications within mobile robotics, anyone interested in applications of learning to planning and control techniques or an interest in building complex adaptive systems is welcome.

Prerequisites

As an advanced course, familiarity with basic ideas from probability and machine learning are strongly recommended. 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 be 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.