

Course wrap up

December 3, 2009

CS 486/686

University of Waterloo

Outline

- Course wrap up
- Final exam info
- Other AI courses
- AI jobs
- AI research


Topics Covered

- Knowledge representation
- Search algorithms
- Probabilistic Inference
- Decision Making under Uncertainty
- Machine Learning

Topics That We Didn't Cover

- Computer Vision
- Natural Language Processing
- Robotics
- Reinforcement Learning
- Multi-agent Systems

Agents and Environments



Agents include humans, robots, softbots, thermostats...

The **agent function** maps percepts to actions $f:P^* \rightarrow A$

The **agent program** runs on the physical architecture to produce f

Rational Agents

- Recall: A rational agent “does the right thing”
- Performance measure - success criteria
 - Evaluates a sequence of environment states
- A **rational agent** chooses whichever action maximizes the **expected** value of its performance measure **given the percept sequence to date**
 - Need to know performance measure, environment, possible actions, percept sequence
- Rationality \neq Omniscience, Perfection, Success
- Rationality \rightarrow exploration, learning, autonomy

Bounded Rationality

- What if the best strategy given past percepts cannot be implemented with today's computers?
- We have seen many theories for rational agents but what if those theories are intractable?
- **Bounded rationality**: find best **implementable** strategy given past percepts

Other AI courses

- CS489/689: Machine Learning (W10)
- CS798: Multiagent Systems (W10)
- CS886: Decision Making Under Uncertainty (W10)
- CS886: Perception as Bayesian Inference (W10)
- CS798: Logic in Computer Science or Theoretical Foundations of Clustering (S10)
- CS848: Machine Learning for Info Retrieval (F10)
- CS886: Bayesian Data Analysis (F10)
- STAT440/840: Computational Inference
- STAT441/841: Statistical Learning - Classification
- STAT442/890 Data visualization

CS489/689: Machine Learning

- Instructor: Pascal Poupart
- Term: Winter 2010
- Topics:
 - Theoretical Foundation of Machine Learning
 - Learning theory
 - Bayesian learning
 - Types of learning:
 - Supervised learning
 - Unsupervised & semi-supervised learning
 - Reinforcement learning
 - Specific models/algorithms
 - Perceptrons
 - Neural networks
 - Support vector machines
 - Gaussian processes
 - Ensemble learning

CS886: Decision Making Under Uncertainty

- Instructor: Pascal Poupart
- Term: Winter 2010
- Topics:
 - Reasoning under uncertainty
 - Sequential decision making
 - Markov decision processes
 - Reinforcement learning
 - Multi-agent systems

CS886: Bayesian Data Analysis

- Instructor: Pascal Poupart
- Term: Fall 2010
- Topics:
 - Bayesian learning
 - Gaussian processes
 - Dirichlet processes
 - Hierarchical priors

AI research group

- Web: ai.uwaterloo.ca
- Professors:
 - Shai Ben David (learning theory)
 - Chrysanne DiMarco (natural language processing)
 - Peter Van Beek (constraint programming)
 - Robin Cohen (multi-agent systems, user modeling)
 - Pascal Poupart (reasoning under uncertainty, machine learning, natural language processing)
 - Kate Larson (game theory, mechanism design)
 - Richard Mann (computational vision)

My research projects

- Decision Making under uncertainty
 - Partially observable Markov decision processes
 - Bayesian reinforcement learning
- Intelligent assistive technologies
 - Smart walkers
 - Symptom monitoring for Alzheimer's disease
- Natural language processing
 - Document clustering
 - Spoken dialogue systems
- Machine learning
 - Semi-supervised learning
 - Cluster Labeling

AI jobs

- Very few "AI companies"
- AI tends to be **embedded** in many applications
- Many companies have AI R&D groups
 - Intel, Microsoft, IBM, Google, NEC, Yahoo, HP
- AI is a growing industry
- Has the potential to revolutionize the computer industry!