Course wrap up

July 25, 2006

CS 486/686

University of Waterloo

Outline

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

Agents and Environments

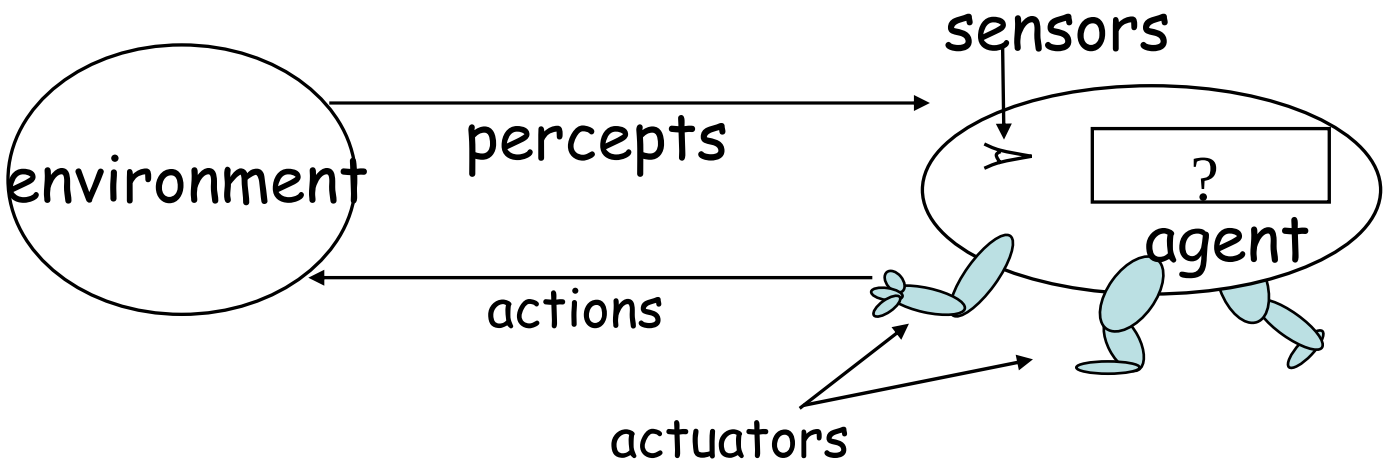

Agents include humans, robots, softbots, thermostats…

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 ≠ Omniscience, Perfection, Success
- Rationality ➔ 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

- CS497: Frontiers of Computer Science
- CS498: Foundations of Machine Learning
- CS498: Image and vision computing
- CS785: Intelligent Computer Interaction
- CS886: Reasoning under Uncertainty
- CS886: Knowledge representation
- CS886: Theoretical Foundations of clustering
- CS886: Game-Theoretic Methods in Computer Science

CS497: Frontiers of Computer Science

- Instructor: Shai Ben David
- Term: Winter 2007
- Objectives:
  - The purpose of this course is to expose undergraduate students to research. It is meant to be a stimulating course where students get a taste of various research topics as well as the opportunity to explore in more depth a topic of their choice.

CS498: Foundations of Machine Learning

- Instructor: Shai Ben David
- Term: Winter 2007
- Objectives:
  - The course is aimed to familiarize the students with the basic theoretical tools and issues underlying some of the most useful machine learning techniques. The theory of machine learning draws from several established mathematical areas including statistics, geometry, combinatorics, and computational complexity.

CS498: Image and Computer Vision

- Instructor: Richard Mann
- Term: ???
- Topics: TBA

CS785 Intelligent Computer Interaction

- Instructor: Robin Cohen
- Term: ???
- Topics:
  - multiagent systems,
  - intelligent tutoring systems and knowledge-based systems,
  - datamining,
  - user modeling,
  - natural language generation and dialogue,
  - plan recognition

CS886: Topics in AI: Reasoning under Uncertainty

- Instructor: Pascal Poupart
- Term: ???
- Objectives:
  - This course will focus on the principles of probabilistic reasoning and sequential decision making for a wide range of settings including adaptive and multi-agent systems. The modeling techniques that will be covered are quite versatile and can be used to tackle a wide range of problems in many fields including robotics (e.g., mobile robot navigation, control), computer systems (e.g., autonomic computing, query optimization), human-computer interaction (e.g., spoken dialog systems, user modeling), bioinformatics (e.g., gene sequencing, design of experiments), operations research (e.g., resource allocation, maintenance scheduling, planning), etc. Hence, the course should be of interest to a wide audience beyond artificial intelligence.

CS886: Topics in AI: Knowledge Representation

- Instructor: Chrysanne DiMarco
- Term: ???
- Topics: TBA

CS886: Topics in AI: Theoretical Foundations of clustering

- Instructor: Shai Ben David
- Term: Fall 2006
- Topics: TBA

CS886: Topics in AI: Game-Theoretic Methods in Computer Science

- Instructor: Kate Larson
- Term: Fall 2006
- Topics: TBA

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

- http://www.cs.uwaterloo.ca/~ppoupart/projects.html

- Partially observable Markov decision processes
- Intelligent assistive technologies
- Spoken dialogue systems
- Bayesian reinforcement learning
- Trust modeling in electronic markets
- Ontology learning
- Preference elicitation

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!