CS886 Winter 2005
Topics in AI:
Reasoning under Uncertainty

Instructor: Pascal Poupart

Email: ppoupart@cs.uwaterloo.ca

Office: DC2514

Lectures: Tu & Th 2:30-4 pm (DC3313)

Web: www.cs.uwaterloo.ca/~ppoupart/teaching/cs886.html

Rational Agent

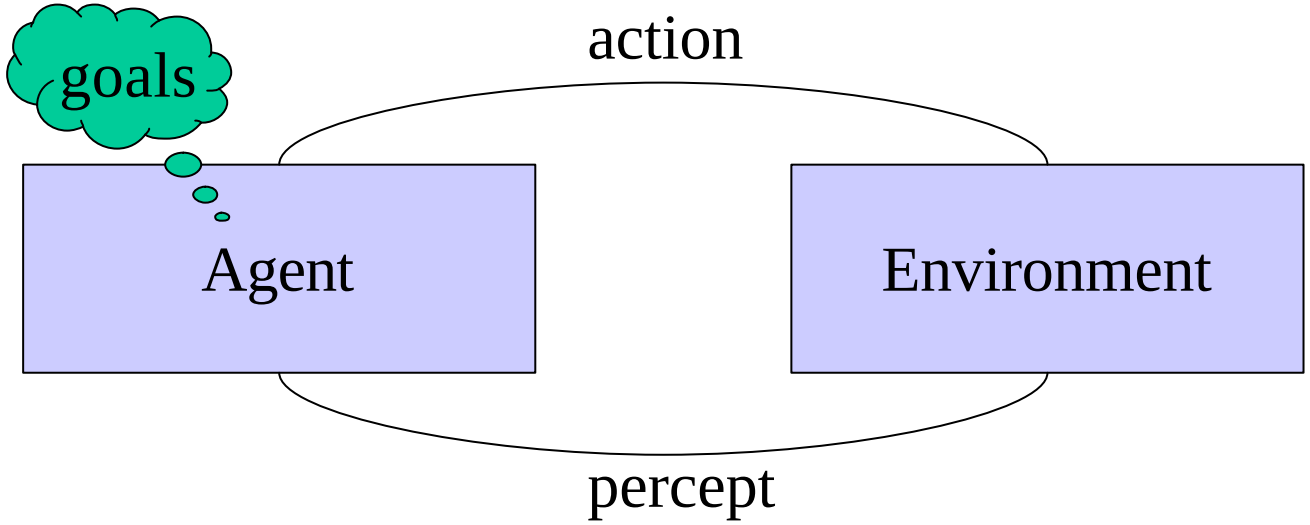

Robot Control

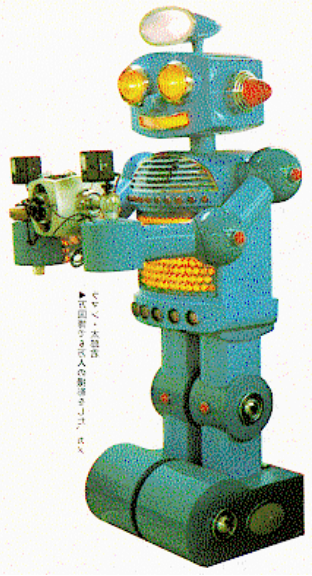

- Consider a multi-purpose robot
- How can we design a robust controller?
  - Interdependent sequence of decisions/controls
  - Imprecise grippers
  - Noisy cameras & microphones
  - Unknown user habits/preferences
  - Tradeoff: escort someone or deliver a parcel first?

Quality Assurance

- Chip manufacturing

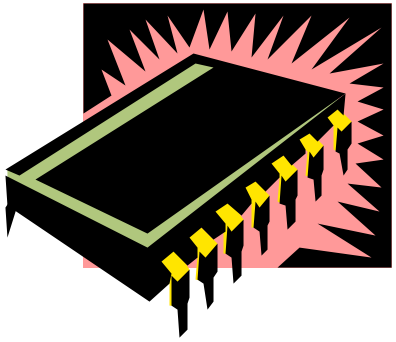

- What would be an effective testing schedule?
  - Tradeoff: testing costs vs. failure costs
  - Tests are noisy, but correlated
  - Sequential aspect: decide about future tests based on results of past tests

Assistive Technology

- People affected by Alzheimer's disease lose their autonomy

- What prompts should a system give to assist in simple daily tasks?
  - Interdependent sequence of prompts
  - Uncertain prompt effects
  - Imperfect / noisy task monitoring
  - Unknown user preferences / abilities

Common Characteristics

- Sequential decision problem
  - Control sequence, test sequence, prompt sequence
- Noisy sensors & imperfect actuators
  - Cameras, microphones, tests, grippers, prompts
- Partially unknown environment
  - User habits, preferences and abilities
- Tradeoff between complex concurrent goals
  - Testing costs vs. quality, delivery vs. escort
- General approach: Reasoning under uncertainty

Topics

- Reasoning under uncertainty
  - Probabilistic inference, statistical inference
    - Bayesian networks and Markov networks
  - Utility theory, decision theory
    - Influence diagrams or decision networks
- Temporal reasoning
  - Sequential inference
    - Dynamic Bayesian networks and hidden Markov models
  - Sequential decision making
    - Markov decision processes

Topics

- Adaptation
  - Machine learning, statistical learning
    - Reinforcement learning, Bayesian learning
- Multi-agent reasoning
  - Game theory
    - Stochastic games, equilibrium concepts

Some applications

- Jan 20: PROVERB: a crossword-puzzle solver
- Jan 25: office assistant (a.k.a. paper clip) and printer wizard in Microsoft Windows
- Feb 1: MINERVA: robotic tour-guide for museums
- Feb 15: TD-Gammon: a world-champion backgammon program
- March 3: Helicopter control

PROVERB: crossword puzzle solver

- Keim, Shazeer, Littman, Agarwal, CCheves, Fitgerald, Grosland, Jiang, Pollard, Weinmeister, *PROVERB: The Probabilistic Cruciverbalist*
- Shazeer, Littman, Keim, *Solving Crossword puzzles as probabilistic constraint satisfaction*

- Looks up online databases of clues
- Probabilistic constraint satisfaction
- Can solve New York Times crossword puzzles with 95% accuracy in 15 minutes

Microsoft Windows

- Horvitz, Breese, Heckerman, Hovel, Rommelse, *The lumiere project: Bayesian user modeling for inferring the goals and needs of software users*
- Breese, Heckerman, *Topics in decision-theoretic troubleshooting: repair and experiment*

- Printer wizard (Windows 95)
- Lumiere project: office assistant (Office 97)
- User modeling with Bayesian networks and decision networks

MINERVA: mobile tour-guide robot

- Fox, Burgard, Dellaert, Thrun, *Monte Carlo localization: efficient position estimation for mobile robots*
- Thrun, Bennewitz, Burgard, Dellaert, Fox, Haehnel, Rosenberg, Schulte, Schulz, *MINERVA: A second generation mobile tour-guide robot*

- Robot localization:
  - Problem: imprecise odometry
  - Solution: probabilistic inference

TD-Gammon: backgammon program

- Tesauro, *TD-Gammon, a self-teaching backgammon program achieves master-level play*
- Tesauro, *Temporal difference learning and TD-Gammon*

- Reinforcement learning with neural networks
- TD-Gammon is now the world champion
- TD-Gammon often serves as an oracle

Helicopter control

- Ng, Jordan, *PEGASUS: a policy search method for large MDPs and POMDPs*
- Ng, Kim, Jordan, Sastry, *Autonomous helicopter flight via reinforcement learning*
- Ng, Coates, Diel, Ganapathi, Schulte, Tse, Berger, Liang, *Inverted autonomous helicopter flight via reinforcement learning*

- POMDP and reinforcement learning techniques

Inverted flight

Many other applications

- Speech recognition
- Spoken dialogue systems
- Autonomic computing
- Database query optimization
- Stochastic resource allocation
- Maintenance scheduling
- Medical diagnosis
- Gene sequencing
- Assistive technologies
- Preference elicitation

Course Structure

- Every lecture:
  - Discuss one (or a few) papers
  - Present material for next paper(s) to read

- What you must do:
  - Read assigned papers
  - Course project
  - Participate in class

Course Evaluation

- Course project (50%)
  - Report (30%)
  - Presentation (20%)
- Paper critiques (40%)
  - 2 paper critiques (10% each)
  - 2 presentations of the critiqued papers (10% each)
- Participation (10%)
  - Discussions in class (5%)
  - Brief summary of every paper (5%)

Course Project

- Choose an application
- Report
  - Literature survey
  - Propose model/technique for application
  - Discuss advantages/disadvantages of model/technique
  - Optional: implement model/technique
- Presentation
  - 15 minutes
  - Describe application
  - Overview of model/technique proposed

Paper critiques

- Choose 2 papers from those assigned for reading
- Each critique
  - 2 pages maximum
  - Review for conference/journal
- Each presentation
  - 15 minutes
  - Overview
  - Lead discussion

Participation

- Paper summaries
  - For each paper assigned for reading
  - At most half a page
  - Brief summary that outlines the main contribution
- Contribute to discussions