Artificial Intelligence II

Spring 2011

Paul Schrater

General Information

- Course Number: CSci 5512
- Class: M W 4:00-05:15 pm
- Web page: http://www-users.itlabs.umn.edu/classes/Spring- 2010/csci5512
  - Or go to www.schrater.org
  - Click on schrater’s homepage
  - Follow the AII 2 courselink

Course Content:

* Uncertainty (Chapter 13)
* Probabilistic Reasoning, Bayesian Networks (Chapter 14)
* Probabilistic Reasoning over Time (Chapter 15)
* Making Simple Decisions (Chapter 16)
* Making Complex Decisions, Markov Decision Processes (Chapter 17)
* Learning from Observations (Chapter 18)
* Statistical Learning Methods (Chapter 20)
* Reinforcement Learning (Chapter 21)
* If time, Language, Vision and Robotics (Chaps 22, 23, 24)

Coursework

- Homeworks
  - There will be 4
  - First one is posted and due in a week
  - Submit using the submit tool!
  - Writeup format: PDF
  - Programming: your choice- matlab, java, or C
  - Individual submission – and include names of people you discuss problems with
- One midterm
- One final project

Homework Schedule

Homework	***Post Date***	***Due Date***	***Due Time***	***Total Time***
HW1	Wed, Jan 19	Wed, Jan 26	4 pm	7 days
HW2	Mon, Feb 21	Mon, Mar 8	4 pm	14 days
HW3	Mon, Mar 28	Mon, Apr 11	4 pm	14 days
HW4	Mon, Apr 18	Mon, May 2	4 pm	14 days

Grading

- Homework: 50 % = 4 × 12.5 %
- Mid-Term: 20 %
- Final Project: 30%

Final Project

- Final Project Assignment: Your final project will involve one of the following
- 1) Simulation or experiments.
- 2) Literature survey (with critical evaluation) on a given topic.
- 3) Theoretical work (detailed derivations, extensions of existing work, etc)
- The project schedule is:
- Feb. 24: Topic selection. One or two pages explaining the project with a list of references.
- May 9: Final report (10 to 15 pages).

Final Project

- In all cases, the work should be written up as a 10-15 page paper. More difficult projects will get better grades if sucessfully completed. You will be evaluated in terms of the care with which you set up and thought through the goals and implementation, and in terms of the competence of the execution. Regardless of form the write up must include a survey of related literature results. This survey counts for 30% of your project grade and should show your ability to independently find, read, understand, and summarize papers in the primary literature related to your project topic.

Autonomous Agents

- **Artificial, Deterministic world**
  - Agents can be programmed to reason and interact in a known environment
- **Real, stochastic, partially observed world**
  - Environmental dynamics and consequences of actions are not fully determined or known (*uncertainty)*
  - Environment must be partially acquired by experience (*learning)*
  - Agents goals must be encoded at a level that permits learning and uncertainty handling (*reinforcement leanrning*

Autonomous Agents

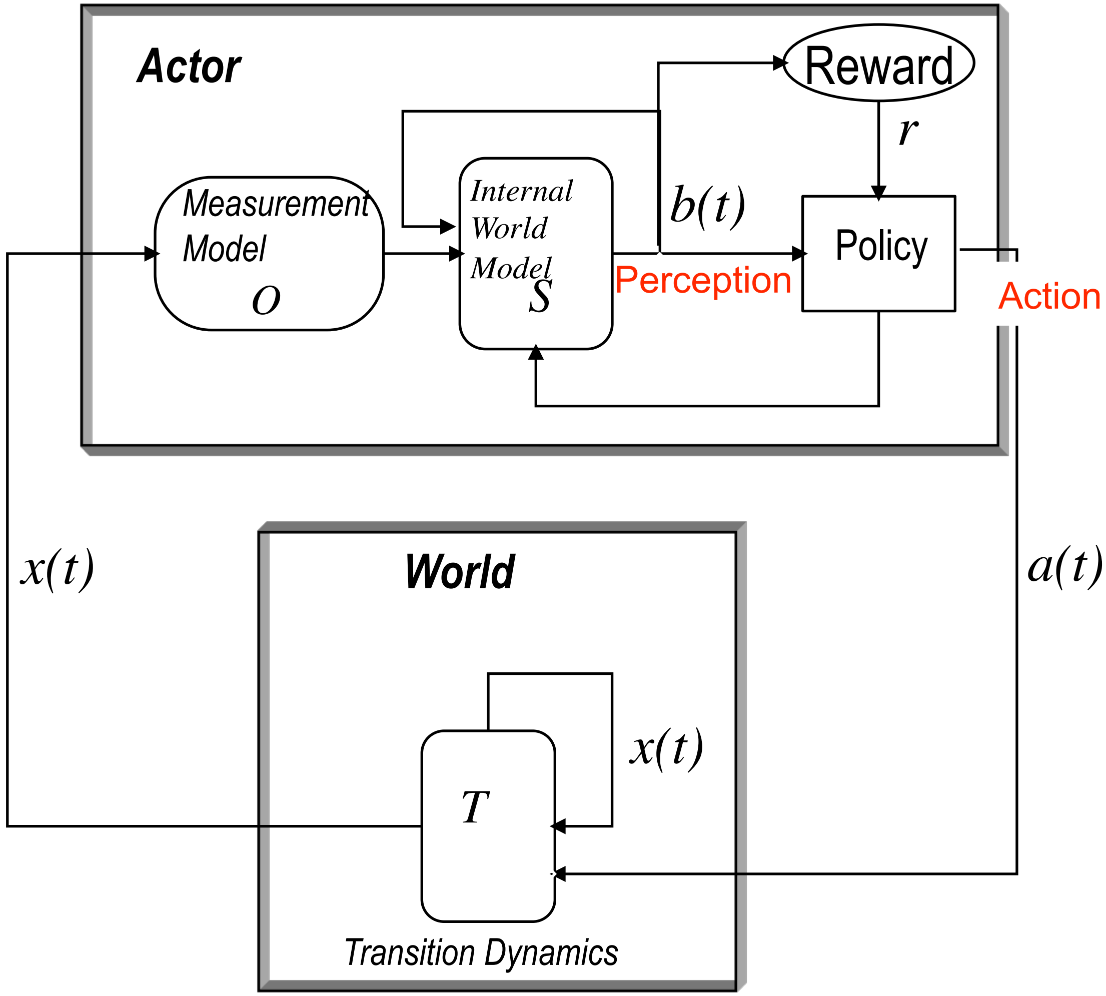

- X: set of states $[x_s, x_r]$
  - state component
  - reward component
- A: set of actions
- $T=P(x'|x,a)$: transition and reward probabilities
- O: Observation function
- b: Belief and info. state
- π: Policy

Topics

- Uncertainty (Probability)
- Probabilistic Reasoning (Bayesian Networks)
- Probabilistic Reasoning over Time (HMMs, DBNs)
- Making Simple/Complex Decisions (Utility, MDPs)
- Game Theory
- Learning from Observations
- Reinforcement Learning
- Latent Variable Models

Uncertainty

- Uncertainty inherent in decision problems

Uncertainty

- Uncertainty inherent in decision problems
    - Partial knowledge of environment

Uncertainty

- Uncertainty inherent in decision problems
  - Partial knowledge of environment
  - Environment may be complex or stochastic

Uncertainty

- Uncertainty inherent in decision problems
  - Partial knowledge of environment
  - Environment may be complex or stochastic
  - Existence of other agents

Uncertainty

- Uncertainty inherent in decision problems
  - Partial knowledge of environment
  - Environment may be complex or stochastic
  - Existence of other agents
- First-order logic is inappropriate for such domains

Uncertainty

- Uncertainty inherent in decision problems
  - Partial knowledge of environment
  - Environment may be complex or stochastic
  - Existence of other agents
- First-order logic is inappropriate for such domains
- Several different events are possible

Uncertainty

- Uncertainty inherent in decision problems
  - Partial knowledge of environment
  - Environment may be complex or stochastic
  - Existence of other agents
- First-order logic is inappropriate for such domains
- Several different events are possible
- Each event

Uncertainty

- Uncertainty inherent in decision problems
  - Partial knowledge of environment
  - Environment may be complex or stochastic
  - Existence of other agents
- First-order logic is inappropriate for such domains
- Several different events are possible
- Each event
  - Has a different “probability” of happening

Uncertainty

- Uncertainty inherent in decision problems
  - Partial knowledge of environment
  - Environment may be complex or stochastic
  - Existence of other agents
- First-order logic is inappropriate for such domains
- Several different events are possible
- Each event
  - Has a different “probability” of happening
  - Has different “utility” or “payoffs”

Uncertainty

- Uncertainty inherent in decision problems
  - Partial knowledge of environment
  - Environment may be complex or stochastic
  - Existence of other agents
- First-order logic is inappropriate for such domains
- Several different events are possible
- Each event
  - Has a different "probability" of happening
  - Has different "utility" or "payoffs"

Uncertainty

- Uncertainty inherent in decision problems
  - Partial knowledge of environment
  - Environment may be complex or stochastic
  - Existence of other agents
- First-order logic is inappropriate for such domains
- Several different events are possible
- Each event
  - Has a different "probability" of happening
  - Has different "utility" or "payoffs"

Rational decisions maximize expected utility

Decision Theory ≡ Utility Theory + Probability Theory

Examples

- Game of Monopoly

Examples

- Game of Monopoly

- Pursuit with Constraints

Examples

- Game of Monopoly

- Pursuit with Constraints
    - Chasing in Manhattan

Examples

- Game of Monopoly

- Pursuit with Constraints
  - Chasing in Manhattan
  - Robotic teams for search/rescue

Examples

- Game of Monopoly

- Pursuit with Constraints
  - Chasing in Manhattan
  - Robotic teams for search/rescue
- The Stock Market

Probability

- Sample space Ω of events
- Each "event" $\omega \in \Omega$ has a associated "measure"
  - Probability of the event $P(\omega)$
- Axioms of Probability:
  - $\forall \omega, P(\omega) \in [0,1]$
  - $P(\Omega) = 1$
  - $P(\omega_1 \cup \omega_2) = P(\omega_1) + P(\omega_2) - P(\omega_1 \cap \omega_2)$

Random Variables

- Random variables are mappings of events (to real numbers)
  - Mapping $X : \Omega \rightarrow R$
  - Any event ω maps to $X(\omega)$
- Example:
  - Tossing a coin has two possible outcomes
  - Denoted by {H,T} or {0,1}
  - Fair coin has uniform probabilities

$$P(X = 0) = \frac{1}{2} \qquad P(X = 1) = \frac{1}{2}$$

- Random variables (r.v.s) can be
  - Discrete, e.g., Bernoulli
  - Continuous, e.g., Gaussian

Distribution, Density

- For a continuous r.v.
  - Distribution function $F(x) = P(X \le x)$
  - Corresponding density function $f(x)dx = dF(x)$
  - Note that

$$F(x) = \int_{t=-\infty}^{x} f(t)dt$$

- For a discrete r.v.
  - Probability mass function $f(x) = P(X = x) = P(x)$
  - We will call this the probability of a discrete event
  - Distribution function $F(x) = P(X \le x)$

Joint Distributions, Marginals

- For two continuous r.v.s X_1, X_2
  - Joint distribution $F(x_1,x_2) = P(X_1 \le x_1, X_2 \le x_2)$
  - Joint density function $f(x_1,x_2)$ can be defined as before
  - The marginal probability density

$$f(x_1) = \int_{x_2=-\infty}^{\infty} f(x_1,x_2)dx_2$$

- For two discrete r.v.s X_1, X_2
  - Joint probability $f(x_1,x_2) = P(X_1 = x_1, X_2 = x_2) = P(x_1,x_2)$
  - The marginal probability

$$P(X_1 = x_1) = \sum_{x_2} P(X_1 = x_1, X_2 = x_2)$$

- Can be extended to joint distribution over several r.v.s
- Many hard problems involve computing marginals

Independence

- Joint probability $P(X_1 = x_1, X_2 = x_2)$

Independence

- Joint probability $P(X_1 = x_1, X_2 = x_2)$
  - X_1, X_2 are different dice

Independence

- Joint probability $P(X_1 = x_1, X_2 = x_2)$
  - X_1, X_2 are different dice
  - X_1 denotes if grass is wet, X_2 denotes if sprinkler was on

Independence

- Joint probability $P(X_1 = x_1, X_2 = x_2)$
  - X_1, X_2 are different dice
  - X_1 denotes if grass is wet, X_2 denotes if sprinkler was on
- Two r.v.s are independent if

$$P(X_1 = x_1, X_2 = x_2) = P(X_1 = x_1)P(X_2 = x_2)$$

Independence

- Joint probability $P(X_1 = x_1, X_2 = x_2)$
  - X_1, X_2 are different dice
  - X_1 denotes if grass is wet, X_2 denotes if sprinkler was on
- Two r.v.s are independent if

$$P(X_1 = x_1, X_2 = x_2) = P(X_1 = x_1)P(X_2 = x_2)$$

  - Two different dice are independent

Independence

- Joint probability $P(X_1 = x_1, X_2 = x_2)$
  - X_1, X_2 are different dice
  - X_1 denotes if grass is wet, X_2 denotes if sprinkler was on
- Two r.v.s are independent if

$$P(X_1 = x_1, X_2 = x_2) = P(X_1 = x_1)P(X_2 = x_2)$$

  - Two different dice are independent
  - If sprinkler was on, then grass will be wet $\Rightarrow$ dependent

Conditional Probability, Bayes Rule

	Grass Wet	Grass Dry
Sprinkler On	0.4	0.1
Sprinkler Off	0.2	0.3

- Inference problems:

Conditional Probability, Bayes Rule

	Grass Wet	Grass Dry
Sprinkler On	0.4	0.1
Sprinkler Off	0.2	0.3

- Inference problems:
  - Given ‘grass wet’ what is P(‘sprinkler on’)

Conditional Probability, Bayes Rule

	Grass Wet	Grass Dry
Sprinkler On	0.4	0.1
Sprinkler Off	0.2	0.3

- Inference problems:
  - Given ‘grass wet’ what is P(‘sprinkler on’)
  - Given ‘symptom’ what is P(‘disease’)

Conditional Probability, Bayes Rule

	Grass Wet	Grass Dry
Sprinkler On	0.4	0.1
Sprinkler Off	0.2	0.3

- Inference problems:
  - Given ‘grass wet’ what is P(‘sprinkler on’)
  - Given ‘symptom’ what is P(‘disease’)
- For any r.v.s X,Y, the conditional probability

$$P(x|y) = \frac{P(x,y)}{P(y)}$$

Conditional Probability, Bayes Rule

	Grass Wet	Grass Dry
Sprinkler On	0.4	0.1
Sprinkler Off	0.2	0.3

- Inference problems:
  - Given 'grass wet' what is P('sprinkler on')
  - Given 'symptom' what is P('disease')
- For any r.v.s X,Y, the conditional probability

$$P(x|y) = \frac{P(x,y)}{P(y)}$$

- Since P(x,y) = P(y|x)P(x), we have

$$P(y|x) = \frac{P(x|y)P(y)}{P(x)}$$

Conditional Probability, Bayes Rule

	Grass Wet	Grass Dry
Sprinkler On	0.4	0.1
Sprinkler Off	0.2	0.3

- Inference problems:
  - Given 'grass wet' what is P('sprinkler on')
  - Given 'symptom' what is P('disease')
- For any r.v.s X,Y, the conditional probability

$$P(x|y) = \frac{P(x,y)}{P(y)}$$

- Since P(x,y) = P(y|x)P(x), we have

$$P(y|x) = \frac{P(x|y)P(y)}{P(x)}$$

- Expressing 'posterior' in terms of 'conditional': Bayes Rule