Uncertainty in AI
COV884

Mausam

(Based on Slides by Stuart Russell, Henry Kautz, Subbarao Kambhampati, and UW-AI faculty)

Personnel

- Instructor: Mausam, SIT 402, mausam@cse.iitd.ac.in
- TAs:
  - Shivanshu Gupta

Logistics

- Timings: Wed 2-3:20 (9 lectures)
- Office hours
  - By appointment
- Course Website: www.cse.iitd.ac.in/~mausam/courses/cov884/spring2018
- Join class discussion group on Piazza (access code csl333) https://piazza.com/iit_delhi/spring2018/cov884/home
- Textbook:
  Artificial Intelligence: A Modern Approach (3rd edition), Russell and Norvig

Programming Assignments

- 1-2 programming assignments
  - some assignments may be done in teams of two (as per instructions)
  - no team can be repeated for a second assignment
  - late policy (penalty of 10% every day)
  - I/O error (penalty of 20%)
  - Logical error (penalty of 50% only under special permission)

Grading and Academic Integrity

- Grading:
  - 50% assignments
  - 50% exam
  - Extra credit: constructive class participation, and discussion group participation

- Academic Integrity
  - Cheating → negative penalty (and possibly more)
    - Exception: if one person/team is identified as cheater
      - Non-cheater gets a zero
  - http://www.willa.me/2013/12/the-top-five-unsanctioned-software.html

- Collaboration is good!!! Cheating is bad!!! Who is a cheater?
  - No sharing of part-code
  - No written/soft copy notes
  - Right to information rule
  - Kyunki saas bhi kabhi bahu thi Rule

Languages

- English ☺
- C++/Java/Python
  - Coding efficiency : python
  - Program efficiency : C++
- Your choice of language may give unfair disadvantage to you!

Who can take the course

Goals of this course

- Modern, Probabilistic AI
- General computer scientist
  - general tools to aid in attacking a new problem
- Serious AI enthusiast
  - A primer from which to launch advanced study

Theory vs. Modeling vs. Applications

- Lecture balance tilted towards modeling
- Assignment balance tilted towards applications
- Relatively few theorems and even fewer proofs

What is *artificial* intelligence?

thought *vs.* behavior	human-like *vs.* rational	
	"[automation of] activities that we associate with human thinking, activities such as decision making, problem solving, learning…" (Bellman 1978)	"The study of mental faculties through the use of computational models" (Charniak & McDertmott 1985)
	"The study of how to make computers do things at which, at the moment, people are better" (Rich & Knight 1991)	"The branch of computer science that is concerned with the automation of intelligent behavior" (Luger & Stubblefield 1993)

What is *artificial* intelligence?

thought *vs.* behavior \ human-like *vs.* rational	human-like *vs.* rational
thought	Systems that think like humans	Systems that think rationally
behavior	Systems that act like humans	Systems that act rationally

Acting rationally

- Rational behavior: doing the right thing
- Need not always be deliberative
  - Reflexive
- Aristotle (Nicomachean ethics)
  - Every art and every inquiry, and similarly every action and every pursuit is thought to aim at some good.

Rational Agents

- An agent should strive to do the right thing, based on what it can perceive and the actions it can perform. The right action is the one that will cause the agent to be most successful

- Performance measure: An objective criterion for success of an agent's behavior

- E.g., performance measure of a vacuum-cleaner agent could be amount of dirt cleaned up, amount of time taken, amount of electricity consumed, amount of noise generated, etc.

Ideal Rational Agent

"For each possible percept sequence, does whatever action is expected to maximize its performance measure on the basis of evidence perceived so far and built-in knowledge."

- Rationality vs omniscience?
- Acting in order to obtain valuable information

Modern AI

- Logic *vs.* Probability
  - In 1950s, logic dominates (McCarthy, …
    - attempts to extend logic
  - 1988 – Bayesian networks (Pearl)
    - efficient computational framework
  - Today, no longer rivals
    - Hot topic: combining probability & FOL

Topics of this Course

- Decision Making under Uncertainty
  - Decision theory
  - Markov Decision Processes
  - Reinforcement Learning
- Probabilistic KR
  - Bayesian Networks