Artificial Intelligence
CSE P573

Mausam

(Slides by Stuart Russell, Dan Weld, Henry Kautz, Oren Etzioni and other UW-AI faculty)

Logistics

- Instructor: Mausam, CSE 454, mausam@cs.washington.edu
- TA: David Broderick, dbroder@cs.washington.edu
- Course Website: www.cs.washington.edu/p573
- Join class mailing list (instructions on website)
- Text: Artificial Intelligence: A Modern Approach (3rd edition), Russell and Norvig
- Four Assignments
- Grading:
  - 50% assignments
  - 15% midterm
  - 25% final
  - 10% class participation, mailing list participation, extra credit, etc.

Goals of this course

- A brief intro to the philosophy of AI
- A brief intro to the breadth of ideas in AI

- General computer scientist
  - general tools to aid in attacking a new problem

- Serious AI enthusiast
  - A primer from which to launch advanced study

Science of AI

Physics: Where did the ***physical universe*** come from?
And what laws guide its dynamics?

Biology: How did ***biological life*** evolve?
And how do living organisms function?

AI: What is the nature of ***intelligent thought?***

AI as Engineering

- How can we make software systems more powerful and easier to use?
  - Speech & intelligent user interfaces
  - Autonomic computing
  - Mobile robots, softbots & immobots
  - Data mining
  - Medical expert systems
  - …

What is intelligence?

- Dictionary.com: *capacity for learning, reasoning, understanding, and similar forms of mental activity*

- Ability to perceive and act in the world
- Reasoning: proving theorems, medical diagnosis
- Planning: take decisions
- Learning and Adaptation: recommend movies, learn traffic patterns
- Understanding: text, speech, visual scene

Intelligence vs. humans

- Are humans intelligent?
- Are humans rational?
- Can non-human behavior be intelligent?

What is *artificial* intelligence?

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

Turing's Test

- If the human cannot tell whether the responses from the other side of a wall are coming from a human or computer, then the computer is intelligent.

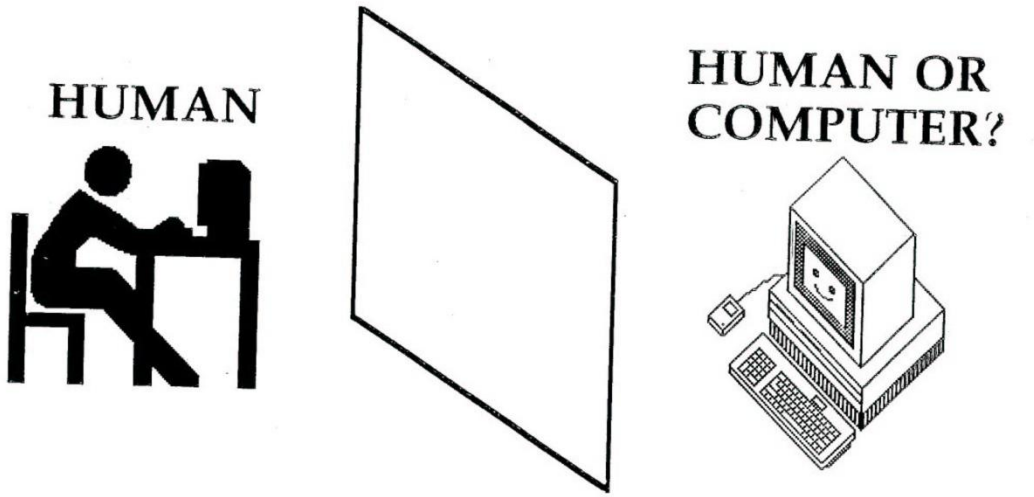

What is *artificial* intelligence (agent view)

- An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators

- Human agent:
  - eyes, ears, and other organs for sensors
  - hands, legs, mouth, and other body parts for actuators

- Robotic agent:
  - cameras and laser range finders for sensors
  - various motors for actuators

- We will revisit this view in detail later in the course

Examples of Agents

- Robots
- Intelligent buildings
- Autonomous spacecraft
- Web agents

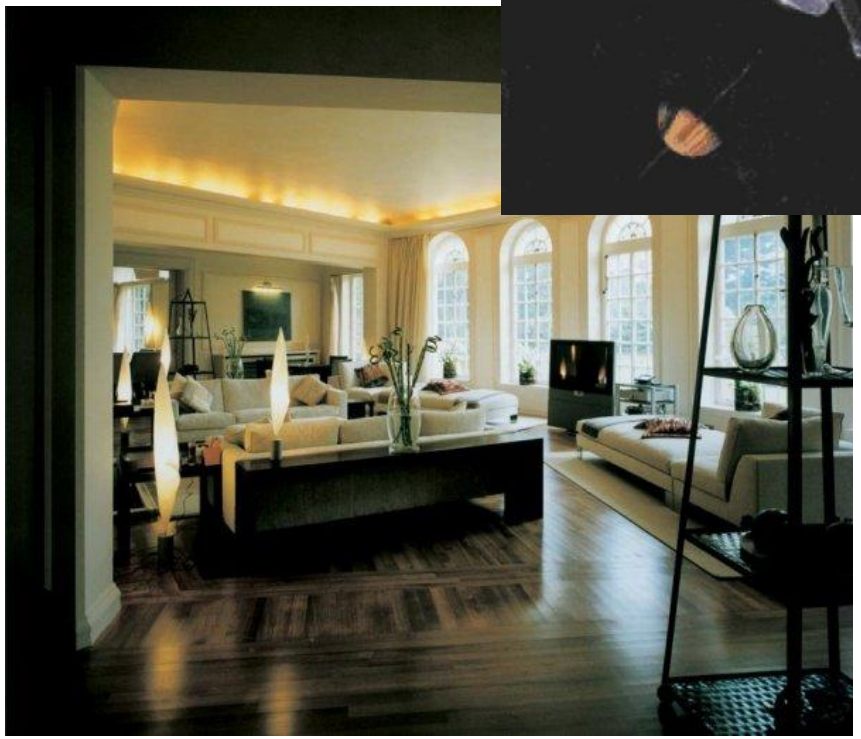

What is *artificial* intelligence (algorithmic view)

- A large number of problems are NP hard

- AI develops a set of tools, heuristics, …
  - to solve such problems in practice
  - for naturally occurring instances

- Search
- Game Playing
- Planning
- …

Examples: Mundane Tasks

- Perception
  - Vision
  - Speech
- Natural Language
  - Understanding
  - Generation
  - Translation
- Reasoning
- Robot Control

Examples: Formal Tasks

- Games
  - Chess
  - Checkers
  - Othello
- Mathematics
  - Logic
  - Geometry
  - Calculus
  - Proving properties of programs

Examples: Expert Tasks

- Engineering
  - Design
  - Fault Finding
  - Manufacturing planning
- Medical
  - Diagnosis
  - Medical Image Analysis
- Financial
  - Stock market predictions

Recurrent Themes

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

Recurrent Themes

- Weak *vs.* Strong Methods
  - Weak – general search methods (e.g., A* search)
    - primarily for problem solving
    - not motivated by achieving human-level performance
  - Strong -- knowledge intensive (e.g., expert systems)
    - more knowledge $\Rightarrow$ less computation
    - achieve better performance in specific tasks
  - How to combine weak & strong methods seamlessly?

Recurrent Themes

- Knowledge Representation
  - "In knowledge lies the power"
  - Feature engineering in Machine Learning
  - Reformulation
- Combinatorial Explosion
- Micro-world successes are hard to scale up.
- How to organize and accumulate large amounts of knowledge?

Mathematical Calculation

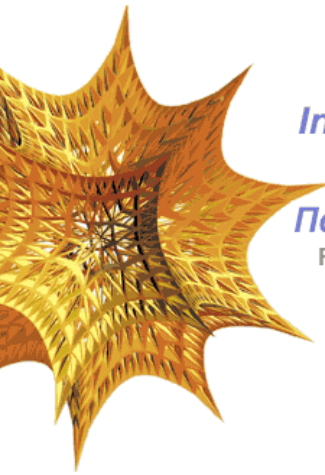

$$\partial_r^2 u = -\left[E' - \frac{l(l+1)}{r^2} - r^2\right] u(r)$$

$$e^{-2s}\left(\partial_s^2 - \partial_s\right) u(s) = -\left[E' - l(l+1)e^{-2s} - e^{2s}\right] u(s)$$

$$e^{-2s}\left[e^{\frac{1}{2}s}\left(e^{-\frac{1}{2}s}u(s)\right)'' - \frac{1}{4}u\right] = -\left[E' - l(l+1)e^{-2s} - e^{2s}\right] u(s)$$

$$e^{-2s}\left[e^{\frac{1}{2}s}\left(e^{-\frac{1}{2}s}u(s)\right)''\right] = -\left[E' - \left(l+\frac{1}{2}\right)^2 e^{-2s} - e^{2s}\right] u(s)$$

$$v'' = -e^{2s}\left[E' - \left(l+\frac{1}{2}\right)^2 e^{-2s} - e^{2s}\right] v$$

Success Story: Chess

"I could feel – I could smell – a new kind of intelligence across the table"

-Gary Kasparov

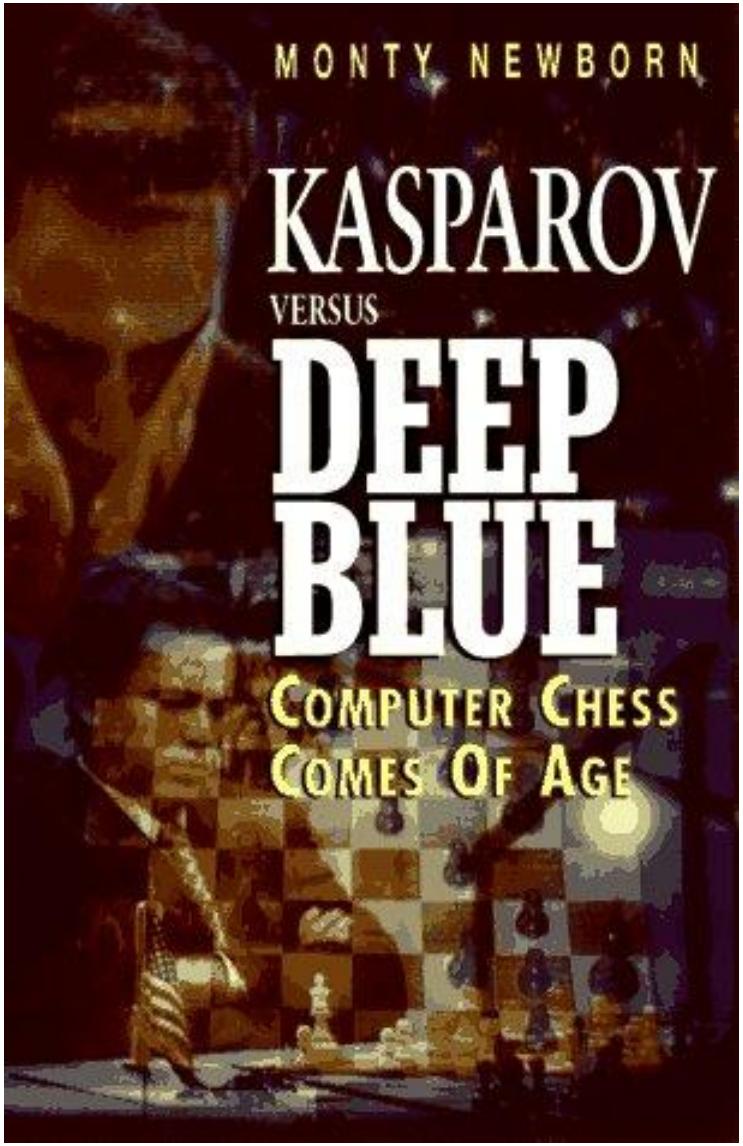

Saying Deep Blue doesn't really think about chess is like saying an airplane doesn't really fly because it doesn't flap its wings.

– Drew McDermott

Speech Recognition

Autonomous Systems

- In the 1990's there was a growing concern that work in classical AI ignored crucial scientific questions:
  - How do we integrate the components of intelligence (*e.g.* learning & planning)?
  - How does perception interact with reasoning?
  - How does the demand for real-time performance in a complex, changing environment affect the architecture of intelligence?

- Provide a standard problem where a wide range of technologies can be integrated and examined
- By 2050, develop a team of fully autonomous humanoid robots that can win against the human world champion team in soccer.

DARPA Urban Challenge: 11/2007

Success Story: Stanley

Software Robots (softbots)

- Softbots: 'intelligent' program that uses software tools on a person's behalf.

- Sensors = LS, Google, etc.
- Effectors = ftp, Amazon.com

- Software: not physical but not simulated.

courtesy JPL

2004 & 2009

Europa Mission ~ 2018

Limits of AI Today

- Today's successful AI systems
  - operate in well-defined domains
  - employ narrow, specialize knowledge

- *Commonsense Knowledge*
  - needed in complex, open-ended worlds
    - Your kitchen vs. GM factory floor
  - understand unconstrained natural language

Role of Knowledge in Natural Language Understanding

- WWW Information Extraction
- Speech Recognition
  - "word spotting" feasible today
  - continuous speech – rapid progress
- Translation / Understanding
  - limited progress

  The spirit is willing but the flesh is weak. ***(English)***

  The vodka is good but the meat is rotten. ***(Russian)***

How the heck do *we* understand?

- John gave Pete a book.
- John gave Pete a hard time.
- John gave Pete a black eye.
- John gave in.
- John gave up.
- John's legs gave out beneath him.
- It is 300 miles, give or take 10.

How to Get Commonsense?

- CYC Project (Doug Lenat, Cycorp)
  - Encoding 1,000,000 commonsense facts about the world by hand
  - Coverage still too spotty for use!

- Machine Learning
- Open Mind
- Mining from Wikipedia & the Web

Topics of this Course

- Search
- Planning
- Constraint Satisfaction
- Logic
- Uncertainty
- Machine Learning
- Guest Lectures: Applications