CSE 473 Artificial Intelligence

Dieter Fox
Colin Zheng

www.cs.washington.edu/education/courses/cse473/08au

Goals of this Course

- To introduce you to a set of key:
  - Paradigms &
  - Techniques
- Teach you to identify when & how to use
  - Agents & Problem Spaces
  - Heuristic search
  - Constraint satisfaction
  - Knowledge representation
  - Planning
  - Uncertainty
  - Machine learning
  - Dynamic Bayesian networks & particle filters
  - Robotics

AI as Science

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

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

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?

Hardware

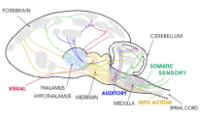

10^{11} neurons
10^{14} synapses
cycle time: 10^{-3} sec

10^{8} transistors
10^{12} bits of RAM
cycle time: 10^{-9} sec

Computer vs. Brain

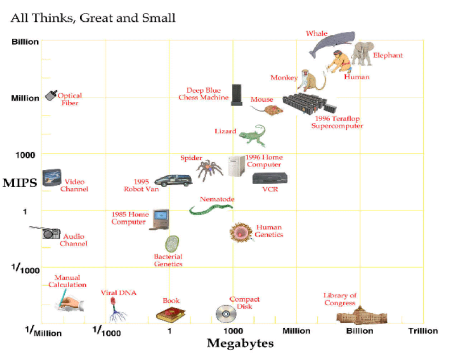

Evolution of Computers

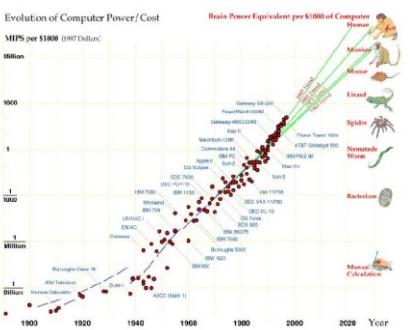

Projection

- In near future computers will have
  - As many processing elements as our brain,
  - But far fewer interconnections
  - Much faster updates.
- Fundamentally different hardware
  - Requires fundamentally different algorithms!
  - Very much an open question.

Dimensions of the AI Definition

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

Mathematical Calculation

$$
\begin{aligned}
 & = -\left[E' - \frac{l(l+1)}{r^2} - r^2\right] u(r) \\
e^{-2s}\left(\partial_s^2 - \partial_s\right) u(s) & = -\left[E' - l(l+1)e^{-2s} - e^{2s}\right] u(s) \\
e^{-2s}\left[e^{\frac{1}{2}s}\left(e^{-\frac{1}{2}s}u(s)\right)'' - \frac{1}{4}u\right] & = -\left[E' - l(l+1)e^{-2s} - e^{2s}\right] u(s) \\
e^{-2s}\left[e^{\frac{1}{2}s}\left(e^{-\frac{1}{2}s}u(s)\right)''\right] & = -\left[E' - \left(l+\frac{1}{2}\right)^2 e^{-2s} - e^{2s}\right] u(s) \\
v'' & = -e^{2s}\left[E' - \left(l+\frac{1}{2}\right)^2 e^{-2s} - e^{2s}\right] v
\end{aligned}
$$

State of the Art

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

-Gary Kasparov

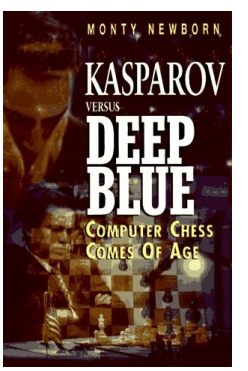

Saying Deep Blue doesn't really think about chess is like saying an airplane doesn't really fly because it doesn't flap its wings.

– Drew McDermott

Speech Recognition

Shuttle Repair Scheduling

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?

RoboCup

- 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.

Software Robots (softbots)

- Softbots: 'intelligent' program that uses software tools on a person's behalf.
- Sensors = LS, Google, etc.
- Effectors = RM, ftp, Amazon.com
- Software: not physical but not simulated.
- Active: not a help system (softbot safety!)

courtesy JPL

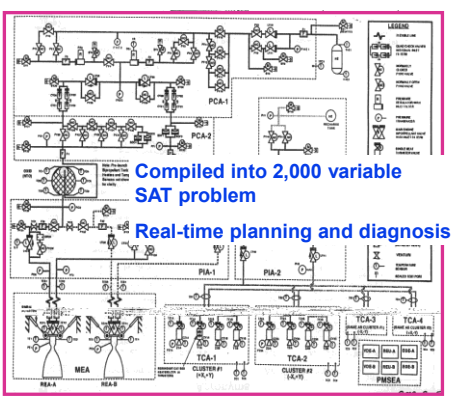

Compiled into 2,000 variable SAT problem

Real-time planning and diagnosis

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!
  - (But see Digital Aristotle project)
- Machine Learning
- Open Mind
- Mining from Wikipedia & the Web
- ???

Recurrent Themes

- Representation *vs.* Implicit
  - Neural Nets - McCulloch & Pitts 1943
    - Died out in 1960's, revived in 1980's
    - Simplified model of real neurons, but still useful; parallelism
  - Brooks "Intelligence without Reprsentation"

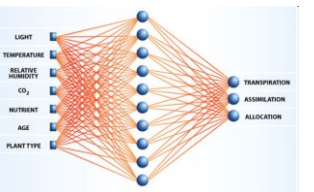

Recurrent Themes

- Logic *vs.* Probability
  - In 1950's, logic dominates (McCarthy, …
    - attempts to extend logic "just a little" (e.g. nomon)
  - 1988 – Bayesian networks (Pearl)
    - efficient computational framework
  - Today's hot topic: combining probability & FOL

Recurrent Themes

- Weak *vs.* Strong Methods
  - Weak – general search methods (e.g. A* search)
  - Knowledge intensive (e.g expert systems)
    - more knowledge ⇒ less computation
  - Today: resurgence of weak methods
    - desktop supercomputers
  - How to combine weak & strong?
- Importance of Representation
  - "In knowledge lies power"
  - Features in ML
  - Reformulation

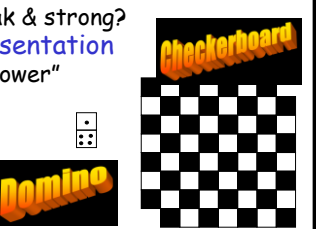

Recurrent Themes

- Combinatorial Explosion
- Micro-world successes are hard to scale up.
- How to organize and accumulate large amounts of knowledge?

Historical Perspective

- (4th C BC+) Aristotle, George Boole, Gottlob Frege, Alfred Tarski
  formalizing the laws of logical reasoning
- (16th C+) Gerc… Pierre Femat, James Bernou… …yes
  formalizing p… …oning
- (1950+) Alan … …on Neumann, Claude Shann…
  thinking as c…
- (1956) John M… …vin Minsky, Herbert Simon, Allen Newell
  start of the field of AI

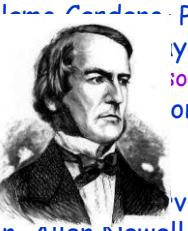

Logistics:

- See website
  www.cs.washington.edu/education/courses/cse473/08au
- Two small projects
  Othello
  TBD
- Grading:
  60% homeworks and mini-projects
  10% midterm
  20% final
  10% class participation, extra credit, etc

For You To Do

- Get on class mailing list
  www.cs.washington.edu/education/courses/cse473/08au
- Dan's Suggestion:
  Start reading Ch 2 in text
  Ch 1 is good, but optional