CSE 573
Artificial Intelligence

Dan Weld
Xu Miao

www.cs.washington.edu/education/courses/cse573/04au

Logistics:

- Dan Weld weld@cs.washington.edu
- Xu Miao xm@cs.washington.edu
- Required Reading
  - Russell & Norvig "AIMA2"
  - Papers from WWW
- Grading:
  - Class Discussion
  - Mini Projects
  - Reviews on Reading
  - Midterm & Problem Sets

For You To Do

- Get on class mailing list
- Monitor class website for reading etc.
- Read
  - Ch 1 [History] is interesting, but optional
  - Ch 2 [Agents] is easy, but important
  - Ch 3 [Search] is crucial, but should be review

Goals of this Course

- To introduce you to a set of key:
  - Paradigms &
  - Techniques
- Teach you to identify when & how to use
  - Heuristic search
  - Constraint satisfaction
  - Machine learning
  - Logical inference
  - Bayesian inference
  - Policy construction
- Teach you how to evaluate (AI) papers
- Highlight directions for research

Outline

- Logistics
- Objectives
- What is AI?
- State of the Art
- Challenges
- Agents

Historical Perspective

- (4th C BC+) Aristotle, George Boole, Gottlob Frege, Alfred Tarski
  - formalizing the laws of human thought
- (16th C+) Gerolamo Cardano, Pierre Femat, James Bernoulli, Thomas Bayes
  - formalizing probabilistic reasoning
- (1950+) Alan Turing, John von Neumann, Claude Shannon
  - thinking as computation
- (1956) John McCarthy, Marvin Minsky, Herbert Simon, Allen Newell
  - start of the field of AI

Hardware

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

10^{7} transistors
10^{10} bits of RAM
cycle time: 10^{-9} sec

Computer vs. Brain

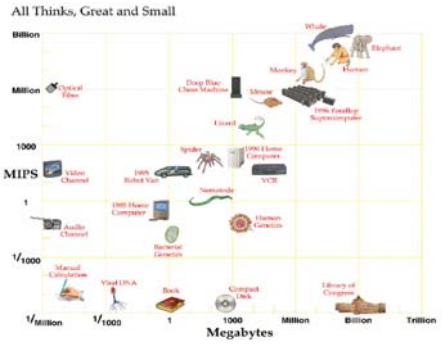

Evolution of Computers

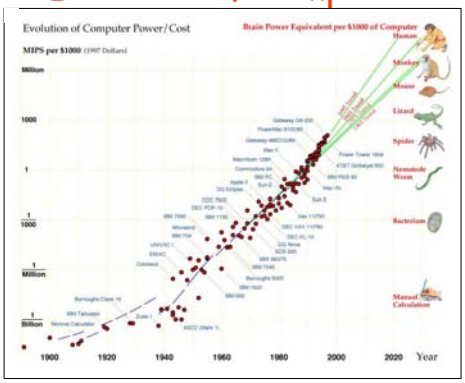

Projection

- 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.

What is Intelligence?

Dimensions of the AI Definition

human-like *vs.* rational

thought *vs.* behavior		
	Systems that think like humans	**Systems that think rationally**
	Systems that act like humans	**Systems that act rationally**

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
SPAM detection
Mobile robots, softbots & immobots
Data mining
Modeling biological systems
Medical expert systems…

State of the Art

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

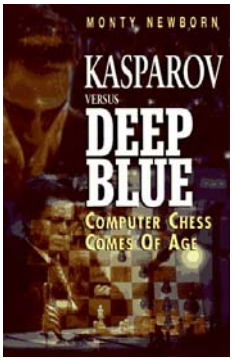

Saying Deep Blue doesn't really think about chess is like saying an airplane doesn't really fly because it doesn't flap its wings.

– Drew McDermott

Mathematical Calculation

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

Shuttle Repair Scheduling

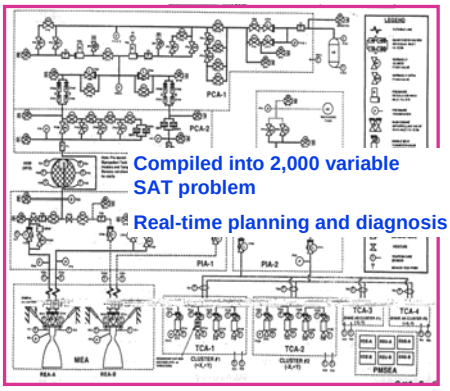

Compiled into 2,000 variable SAT problem

Real-time planning and diagnosis

2004 & 2009

Europa Mission ~ 2018

Credit Card Fraud Detection

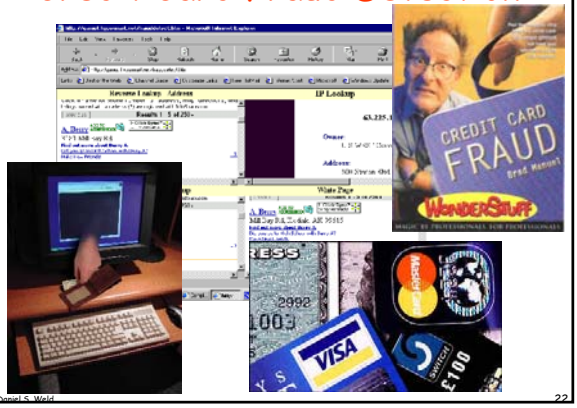

Speech Recognition

Autonomous Navigation: NAVLAB 1

NAVLAB 2

NAVLAB 11

NAVLAB 5

NAVLAB 7

NAVLAB 23?

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
- Alternatives?

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 II

- 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 III

- 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?

Recurrent Themes IV

- Importance of Representation
  - Features in ML
  - Reformulation

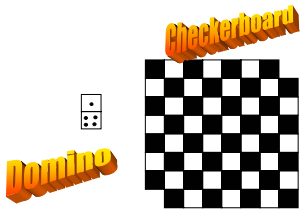

573 Topics

- Agents
- Search thru Problem Spaces & Constraint Sat
- Knowledge Representation
- Learning
- Planning
- Markov Decision Processes
- Reinforcement Learning

Intelligent Agents

- Have sensors, effectors
- Implement mapping from percept sequence to actions

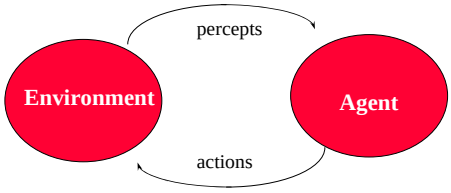

- Performance Measure

Defn: 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

Defn: Autonomy

An agent is autonomous to the extent that its behavior is determined by its own experience

Why is this important?

The parable of the dung beetle

Implementing ideal rational agent

- Table lookup agents
- Agent program
  - Simple reflex agents
  - Agents with memory
    - Reflex agent with internal state
    - Goal-based agents
    - Utility-based agents

Simple reflex agents

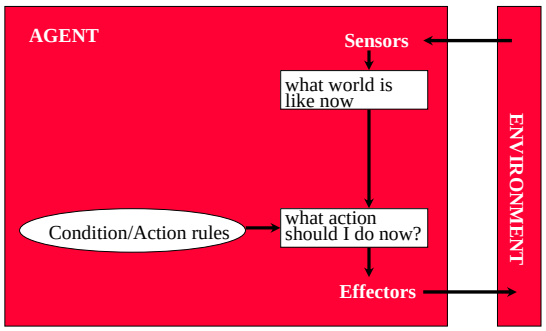

Reflex agent with internal state

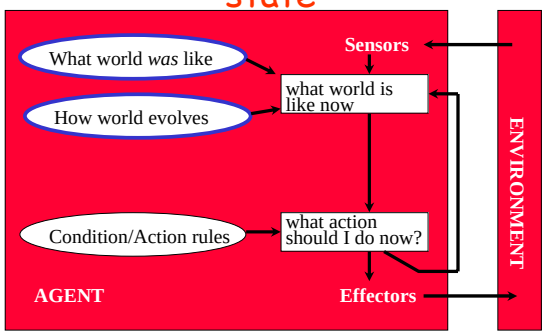

Goal-based agents

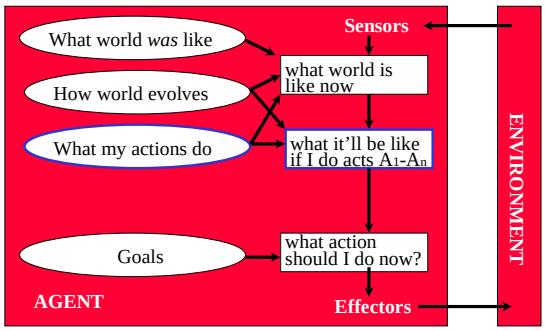

Utility-based agents

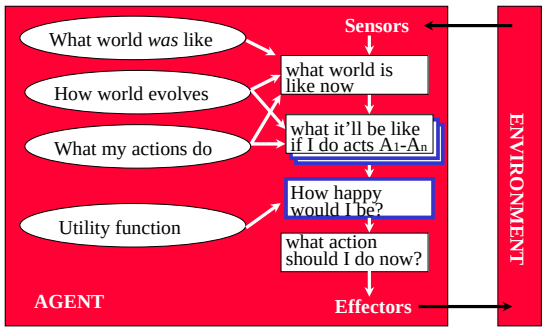

Properties of Environments

- Observability: full *vs.* partial *vs. non*
- Deterministic *vs.* stochastic
- Episodic *vs.* sequential
- Static *vs.* … *vs.* dynamic
- Discrete *vs.* continuous

- **Travel agent**
- **WWW shopping agent**
- **Coffee delivery mobile robot**