CSE 473
Artificial Intelligence (AI)

Rajesh Rao (Instructor)
Ravi Kiran (TA)

http://www.cs.washington.edu/473

Outline

- Goals of this course
- Logistics
- What is AI?
- Examples
- Challenges

CSE 473 Goals

- To introduce you to a set of key:
  - Concepts &
  - Techniques in AI
- Teach you to identify when & how to use
  - Heuristic search for problem solving and games
  - Logic for knowledge representation and reasoning
  - Bayesian inference for reasoning under uncertainty
  - Machine learning (for pretty much everything)

3

CSE 473 Topics

- Agents & Environments
- Search
- Logic and Knowledge Representation
- Planning
- Reasoning under Uncertainty
- Machine Learning

4

CSE 473 Logistics

- E-mail:

Rajesh Rao	rao@cs
Ravi Kiran	kiran@cs

- Required Textbook
  Russell & Norvig's "AIMA2"
- Grading:

Homeworks and projects	50%
Midterm	20%
Final	30%

- Midterm on Wednesday, Nov 1, in class (closed book, except for one 8 ½'' x 11''page of notes)

5

AI as Science

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 "***intelligence***" and what constitutes intelligent behavior?

6

AI as Engineering

- How can we make software and robotic devices more powerful, adaptive, and easier to use?
- Examples:
  - Speech recognition
  - Natural language understanding
  - Computer vision and image understanding
  - Intelligent user interfaces
  - Data mining
  - Mobile robots, softbots, humanoids
  - Medical expert systems…

7

Hardware

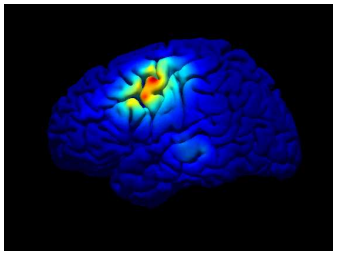

10^{11} neurons
10^{14} synapses
cycle time: 10^{-3} sec

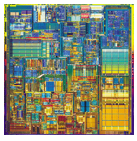

10^{7} transistors
10^{10} bits of RAM
cycle time: 10^{-9} sec

8

Computer vs. Brain

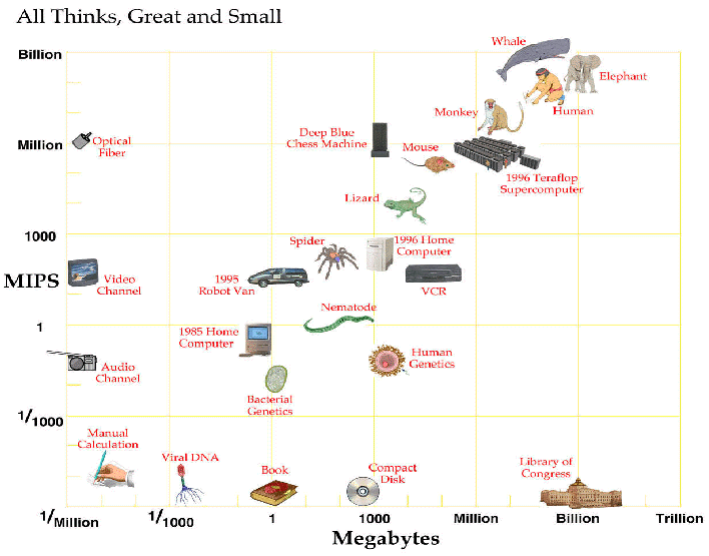

(from Moravec, 1998)

Evolution of Computers

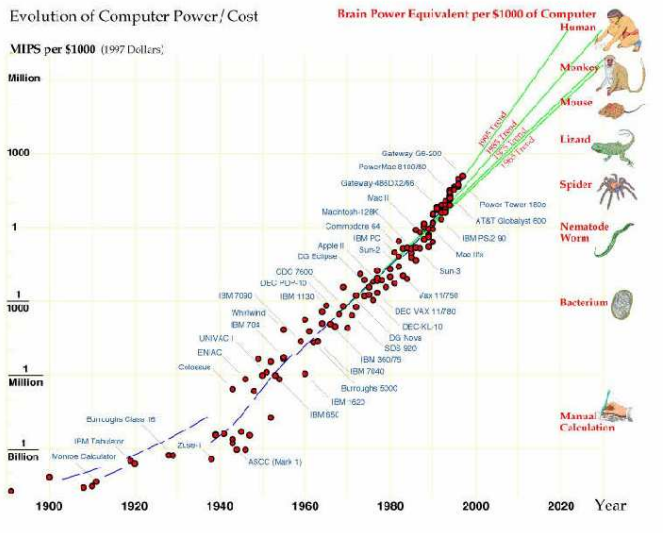

(from Moravec, 1998)

Projection

- In near future (~2020) computers will

  become cheap enough and have enough processing power and memory capacity to *match the general intellectual performance of the human brain*

- But…what "software" does the human brain run?

  Very much an open question

Defining AI Systems

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

12

History of AI: Foundations

- Logic: rules of rational thought
  - Aristotle (384-322 BC) – syllogisms
  - Boole (1815-1864) – propositional logic
  - Frege (1848-1925) – first-order logic
  - Hilbert (1962-1943) – "Hilbert's Program"
  - Gödel (1906-1978) – incompleteness
  - Turing (1912-1954) – computability, Turing test
  - Cook (1971) – NP completeness

13

History of AI: Foundations

- Probability & Game Theory
  - Cardoano (1501-1576) – probabilities
  - Bernoulli (1654-1705) – random variables
  - Bayes (1702-1761) – belief update
  - von Neumann (1944) – game theory
  - Richard Bellman (1957) – MDP

14

Early AI

- Neural networks
  - McCulloch & Pitts (1943)
  - Rosenblatt (1962) – perceptron learning
- Symbolic processing
  - Dartmouth conference (1956)
  - Newell & Simon – logic theorist
  - John McCarthy – symbolic knowledge representation
  - Samuel's Checkers Program

15

Battle for the Soul of AI

- Minsky & Papert (1969) – *Perceptrons*
  - Single-layer networks cannot learn XOR
  - Argued against neural nets in general
- Backpropagation
  - Invented in 1969 and again in 1974
  - Hardware too slow, until rediscovered in 1985
- Research funding for neural nets disappears
- Rise of rule-based expert systems

16

Knowledge is Power

- Expert systems (1969-1980)
  - Dendral – molecular chemistry
  - Mycin – infectious disease
  - R1 – computer configuration
- AI Boom (1975-1985)
  - LISP machines
  - Japan's 5th Generation Project

17

AI Winter

- Expert systems oversold
  - Fragile
  - Hard to build, maintain
- AI Winter (1985-1990)
- Science went on... looking for
  - Principles for robust reasoning
  - Principles for learning

18

AI Now

- Graphical probabilistic models
  - Pearl (1988) – Bayesian networks
- Machine learning
  - Quinlan (1993) – decision trees (C4.5)
  - Vapnik (1992) – Support vector machines
  - Schapire (1996) – Boosting
  - Neal (1996) – Gaussian processes

19

AI Now: Applications

- Countless AI systems in day to day use
  - Industrial robotics
  - Data mining on the web
  - Speech recognition
  - Face & Iris recognition
  - Market research
  - Computational biology
  - Hardware verification
  - Credit card fraud detection
  - Surveillance
  - Threat assessment
  - Etc.

20

Notable Examples: Chess (Deep Blue, 1997)

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

-Gary Kasparov

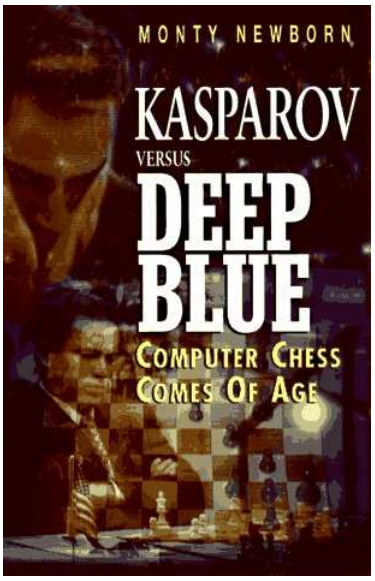

Saying Deep Blue doesn't really think about chess is like saying an airplane doesn't really fly because it doesn't flap its wings.

– Drew McDermott

21

Speech Recognition

Navigation Systems

Automated call centers

22

Natural Language Understanding

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

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

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

23

Museum Tour-Guide Robots

Rhino, 1997

Minerva, 1998

24

Mars Rovers (2003-now)

25

Europa Mission ~ 2018?

Humanoid Robots

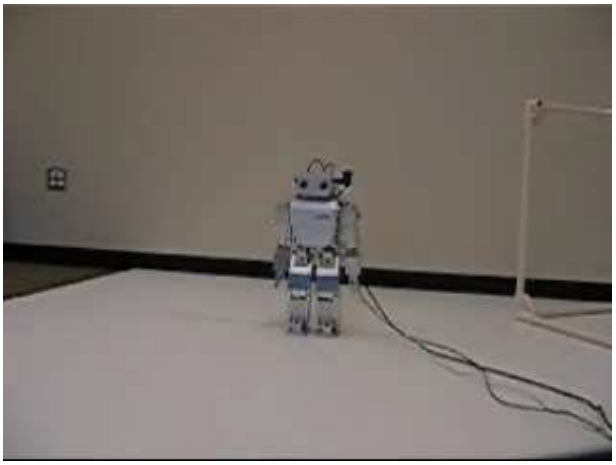

27

Robots that Learn

Before Learning

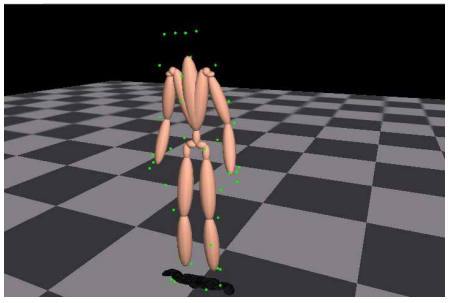

Human Motion Capture

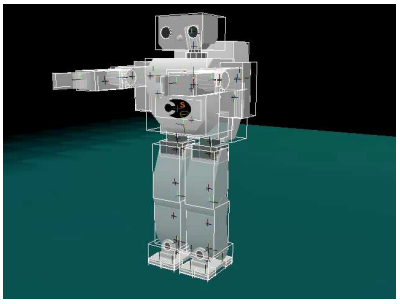

Attempted Imitation

28

Robots that Learn

After Learning

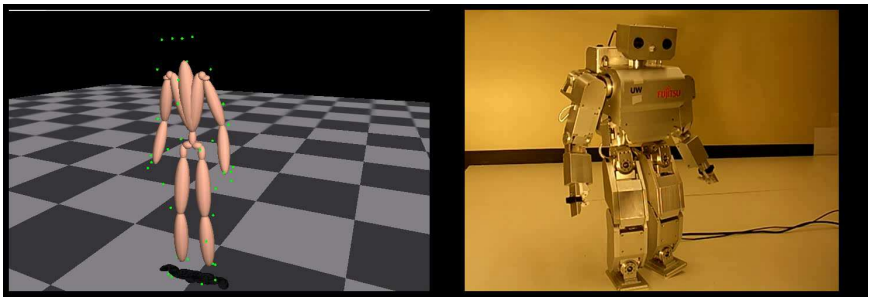

Movie

29

Chess Playing vs. Robots

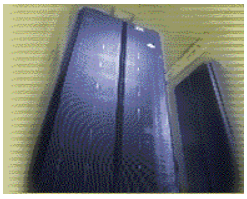

Deep Blue

Robot

- Static
- Deterministic
- Turn-based

Dynamic
Stochastic
Real-time

30

Robotic Prosthetics

31

Brain-Computer Interfaces

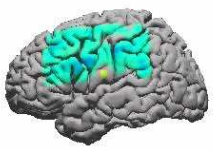

32

Limitations of AI Systems Today

- Today's successful AI systems
  - operate in well-defined domains
  - employ narrow, specialized hard-wired knowledge
- *Needed:* Ability to
  - Operate in complex, open-ended dynamic worlds
    - E.g., Your kitchen vs. GM factory floor
  - Adapt to unforeseen circumstances
  - Learn from new experiences
- In this class, we will explore some potentially useful techniques for tackling these problems

33

For You To Do

- Browse CSE 473 course web page
- Get on class mailing list
- Read Chapters 1 and 2 in text
- HW #1 to be assigned on Monday

34