CSE 473
Artificial Intelligence (AI)

Rajesh Rao (Instructor)
Yi-Shu Wei (TA)
Hunter Whalen (TA)

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

Based on slides by UW CSE AI faculty, Dan Klein, Stuart Russell, Andrew Moore

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
  - Probabilistic inference for reasoning under uncertainty
  - Machine learning (for pretty much everything)

CSE 473 Logistics

- E-mail:

Rajesh Rao	rao@cs
Yi-Shu Wei	yishuwei@uw.edu
Hunter Whalen	hwhalen@cs

- Required Textbook
  - Russell & Norvig's "AIMA3"

- Grading:

– Homeworks and projects	50%
– Midterm	20%
– Final	30%

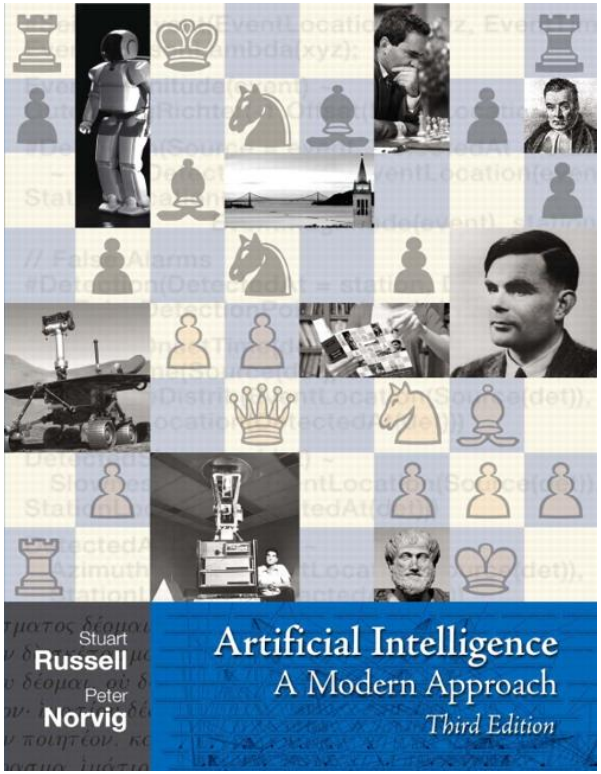

- Midterm on Monday, October 28, in class (closed book, except for one 8 ½" x 11"page of notes)

CSE 473 Topics

- Overview, agents, environments (Chaps 1 and 2)
- Search (Chaps 3 and 5)
- Knowledge representation and logic (Chaps 7-9)
- Uncertainty & Bayesian networks (Selected topics from Chaps 13-15 and 17)
- Machine Learning: Learning from examples (Chap 18)
- Machine Learning: Reinforcement learning (Chap 21)

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?

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
  - Brain-computer interfaces…

Hardware

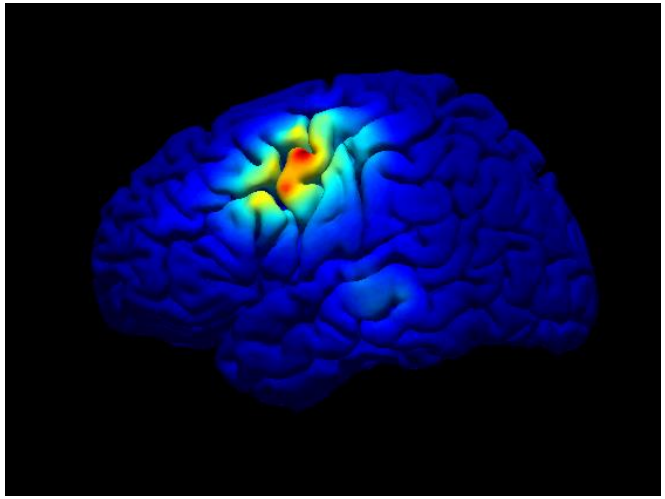

10^{11} neurons
10^{14} synapses
cycle time: 10^{-3} sec (1 kHz)

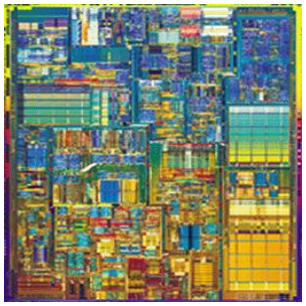

10^{10} transistors
10^{12} bits of RAM (125 GB)
cycle time: 10^{-10} sec (10 GHz)

Computer vs. Brain

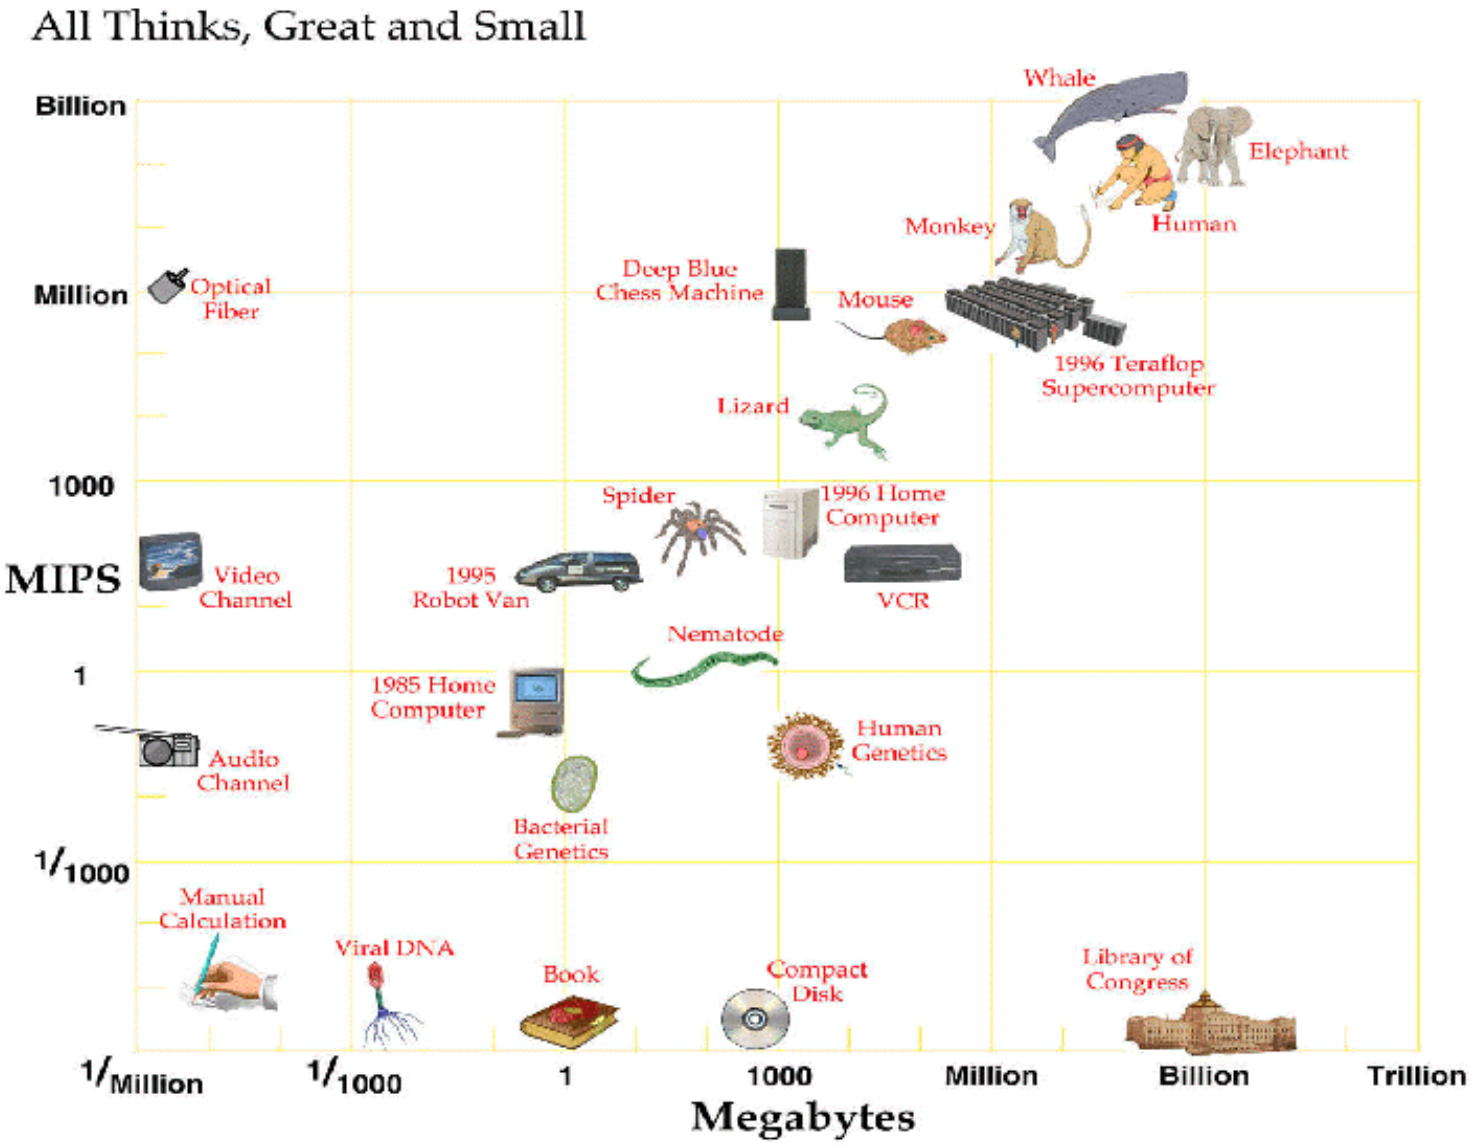

(from Moravec, 1998)

Evolution of Computers

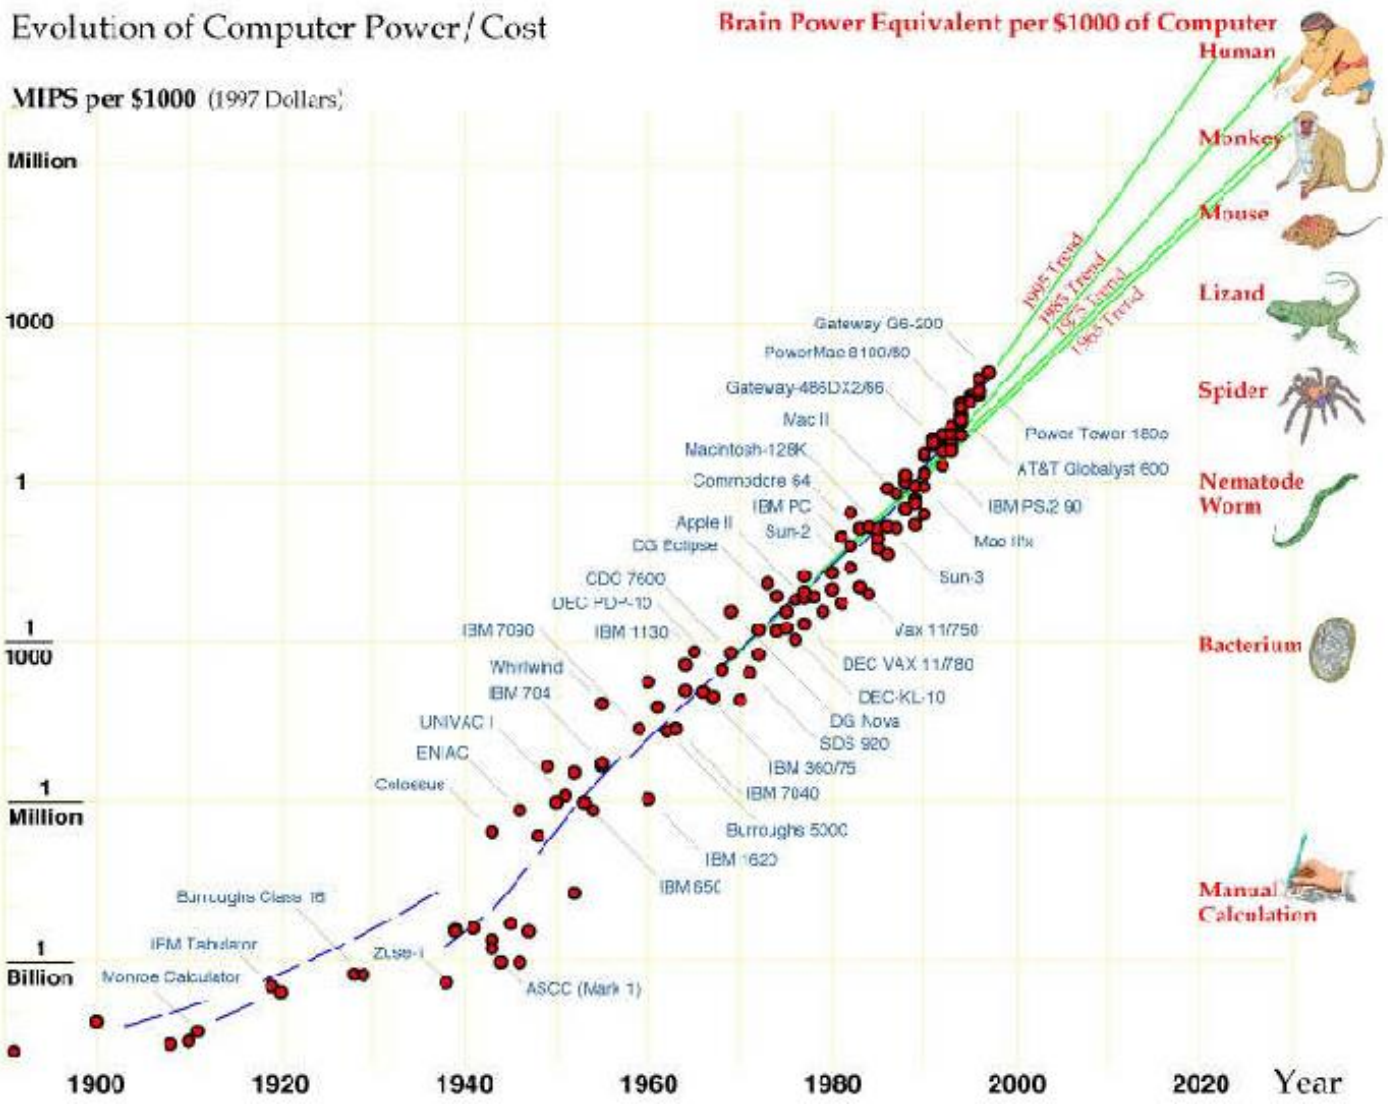

(from Moravec, 1998)

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

What is AI?

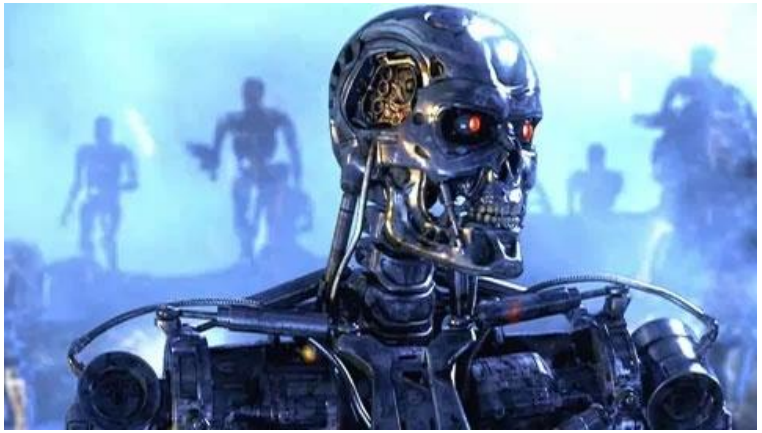

Defining AI

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

Rational: maximally achieving pre-defined goals

AI Prehistory

- Logical Reasoning: (4th C BC+) Aristotle, George Boole, Gottlob Frege, Alfred Tarski
- Probabilistic Reasoning: (16th C+) Gerolamo Cardano, Pierre Fermat, James Bernoulli, Thomas Bayes

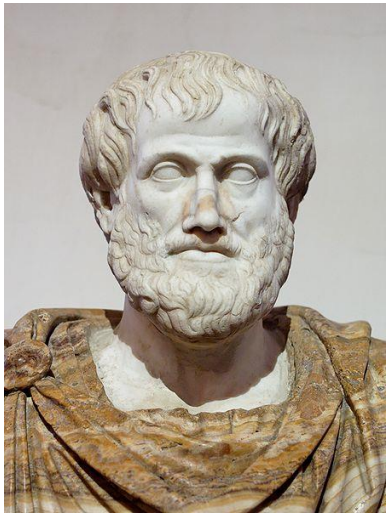

1940-1950: The Early Days

- 1943: McCulloch & Pitts: Boolean circuit model of brain

- 1950: Turing's “Computing Machinery and Intelligence”

I propose to consider the question, "Can machines think?" This should begin with definitions of the meaning of the terms "machine" and "think." The definitions might be framed...

-Alan Turing

The Turing Test

- Turing (1950) "Computing machinery and intelligence"
  - "Can machines think?" → "Can machines interact intelligently?"
  - The *Human Interaction Game:*

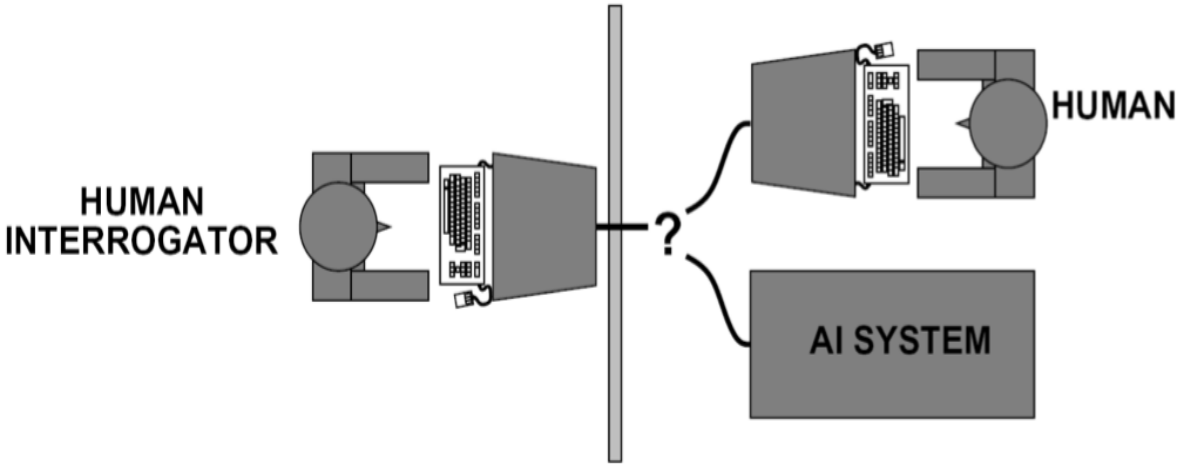

  - Suggested major components of AI: knowledge, reasoning, language understanding, learning
  - Missing: Physical interactions with the real-world

1950-1965: Excitement

- **1950s: Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine**

- **1956: Dartmouth meeting: "Artificial Intelligence" adopted**

- **1965: Robinson's complete algorithm for logical reasoning**

"Over Christmas, Allen Newell and I created a thinking machine."

-Herbert Simon

Battle for the Soul of AI

- **Minsky & Papert (1969) – *Perceptrons* book**
  - **Single-layer neural networks cannot learn XOR**
  - **Argued against neural nets in general**

- **Backpropagation learning algorithm**
  - **Invented in 1969 and again in 1974**
  - **Hardware too slow, until rediscovered in 1985**

- **Research funding for neural nets disappears**
- **Rise of knowledge based systems**

1970-1980: Knowledge Based Systems

- **1969-79: Early development of knowledge-based systems**
- **1980-88: Expert systems industry booms**
- **1988-93: Expert systems industry busts**

  “AI Winter”

1988-present: Statistical Approaches

- **1985-1990: Probability and Decision Theory become dominant**

  Pearl, Bayes Nets

- **1990-2000: Machine learning takes over subfields: Vision, Natural Language, etc.**

- **Agents, uncertainty, and learning systems…**

  "AI Spring"?

  "Every time I fire a linguist, the performance of the speech recognizer goes up"
  -*Fred Jelinek*, IBM Speech Team

Pop Quiz

Which of the following can be done by AI systems today?

- **Play a decent game of Soccer?**
- **Defeat a human in a game of Chess? Go? Jeopardy?**
- **Drive a car safely along a curving mountain road? On University Way?**
- **Buy a week's worth of groceries on the Web? At QFC?**
- **Make a car? Make a cake in your kitchen?**
- **Discover and prove a new mathematical theorem?**
- **Perform a heart bypass surgery?**
- **Unload a dishwasher and put everything away?**
- **Translate Mandarin Chinese into English in real time?**

Examples: Chess (Deep Blue, 1997)

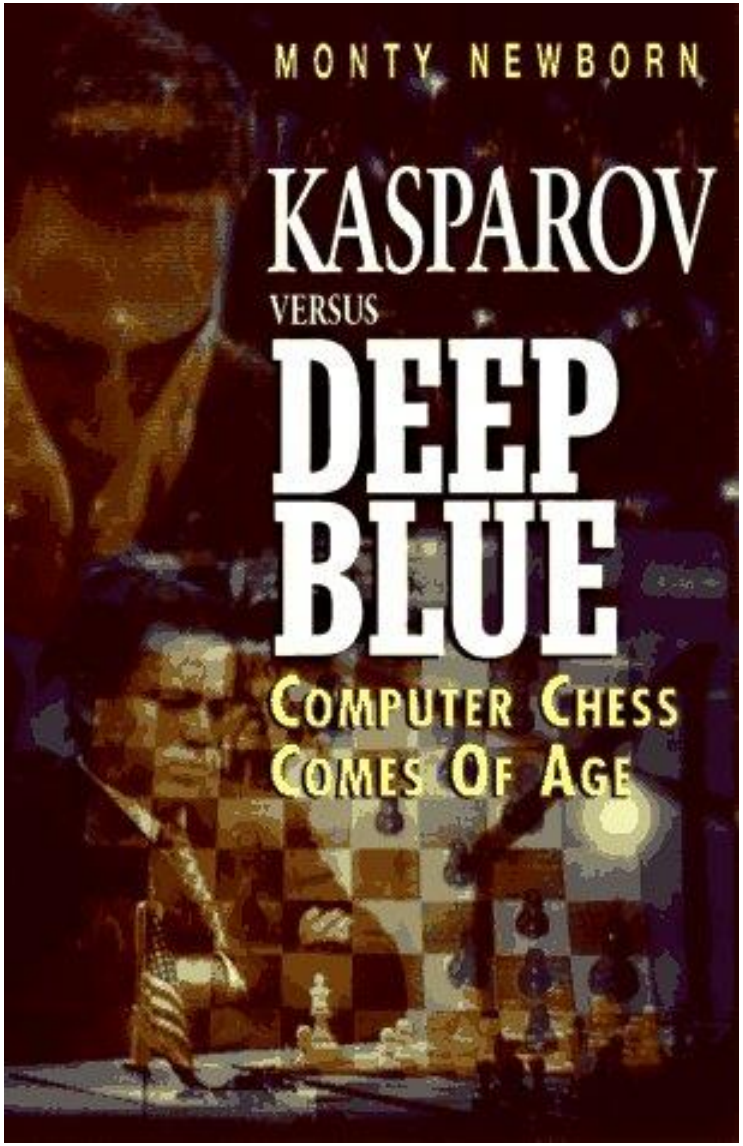

"I could feel –
I could smell –
a new kind of
intelligence
across the
table"

-Gary
Kasparov

Speech Recognition

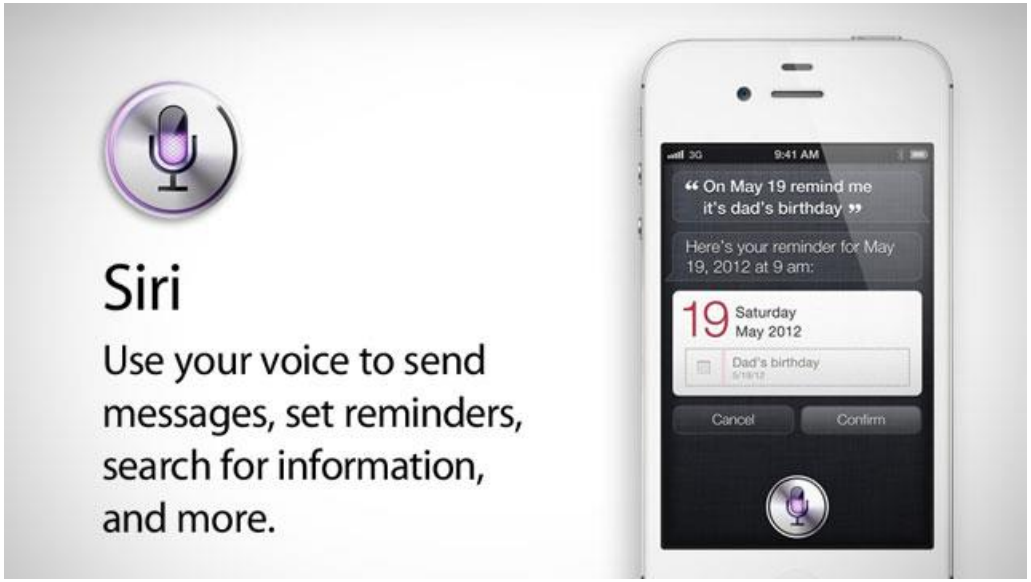

Navigation Systems

Automated call centers

Natural Language Understanding

- Speech Recognition
  - "word spotting" feasible today
  - continuous speech – limited success
- Machine Translation / Understanding
  - progress but not there yet

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

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

Mars Rovers (2003-now)

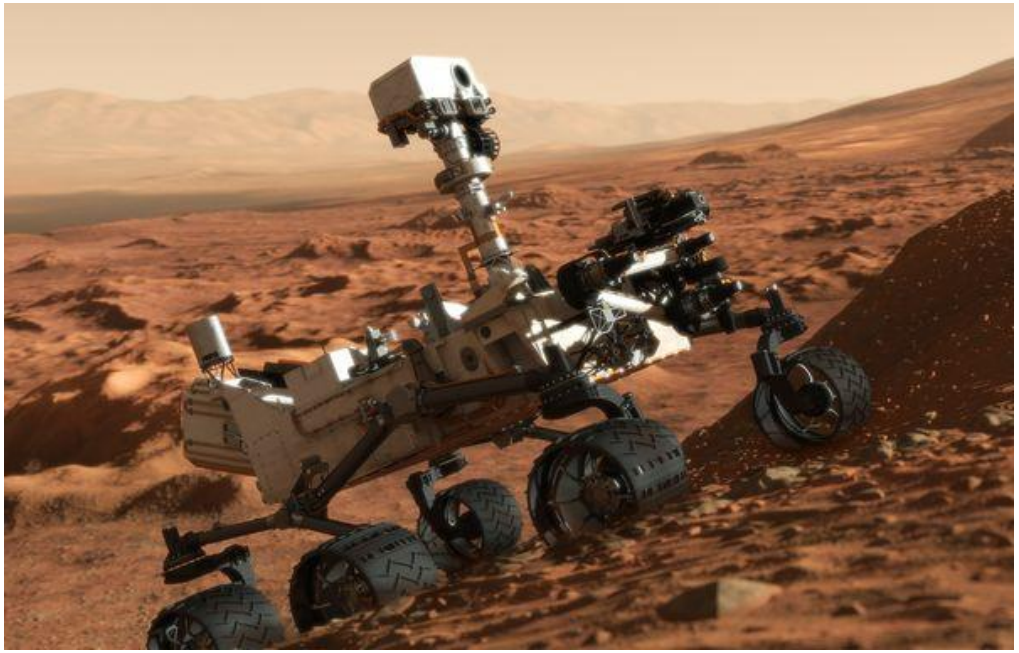

(See NASA website for latest updates)

Robots that Learn

Before Learning

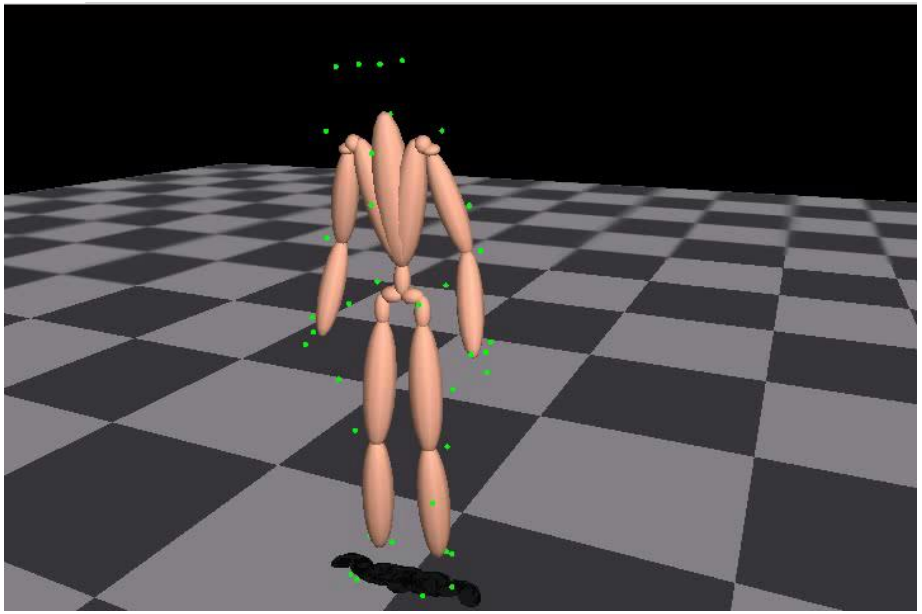

Human Motion Capture

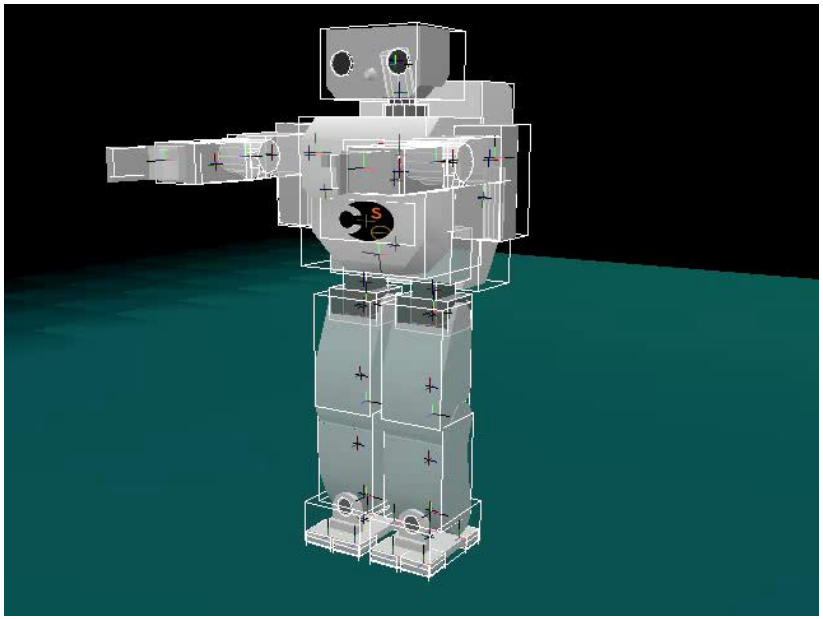

Attempted Imitation

Robots that Learn

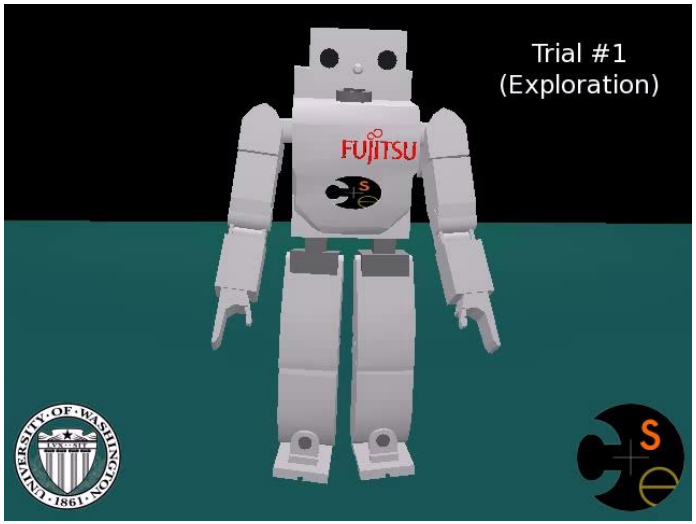

Learning

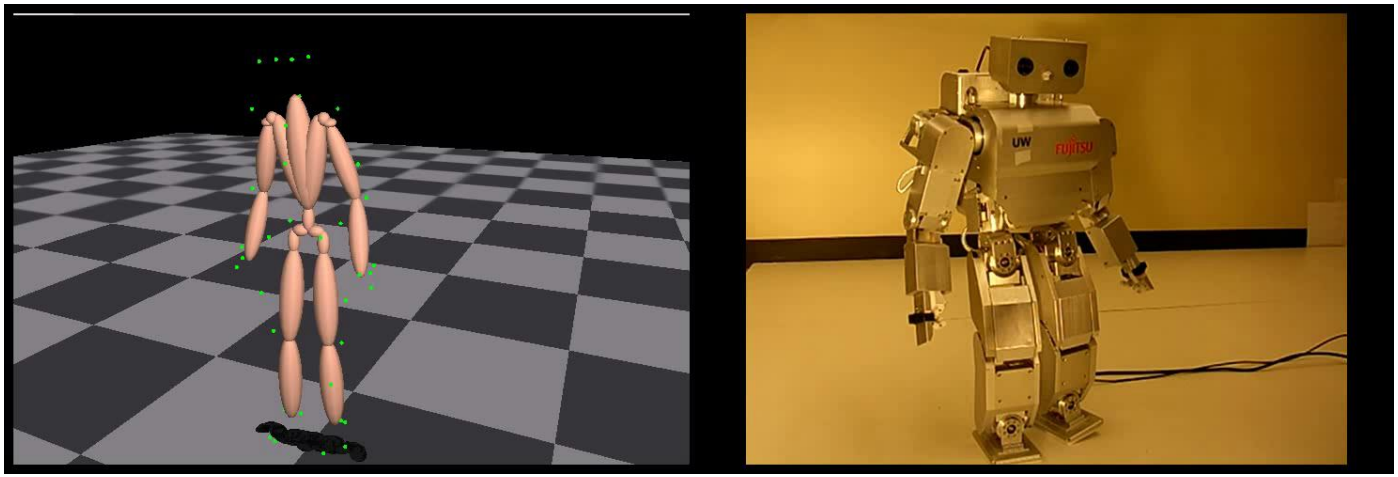

After Learning

(Work by UW CSE PhD David Grimes)

Muscle-Activated Robotics

(Work by UW CSE undergrad Beau Crawford)

Brain-Computer Interfaces

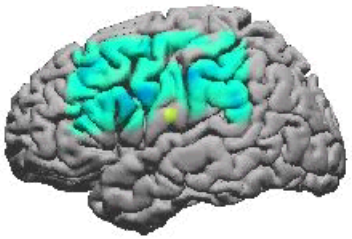

(Work by UW MD-PhD Kai Miller)

Limitations of AI Systems Today

- Today's successful AI systems
  - operate in well-defined domains
  - employ narrow, specialized hard-wired knowledge
- *Missing*: 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

For You To Do

- Browse CSE 473 course web page
- Do Project 0: Python tutorial
- Read Chapters 1 and 2 in text
- Project 1 to be assigned on Monday