# Lecture 1
# What is AI?

## CSE 473
## Artificial Intelligence

Oren Etzioni

# AI as Science

*What are the most fundamental scientific questions?*

# Goals of this Course

- To teach you the main ideas of AI.
  - Give you AI "color"
- To introduce you to a set of key techniques and algorithms from AI
- To introduce you to the applicability and limitations of these methods (problem sets)

# What is Intelligence?

# What is Artificial Intelligence?

# Hardware

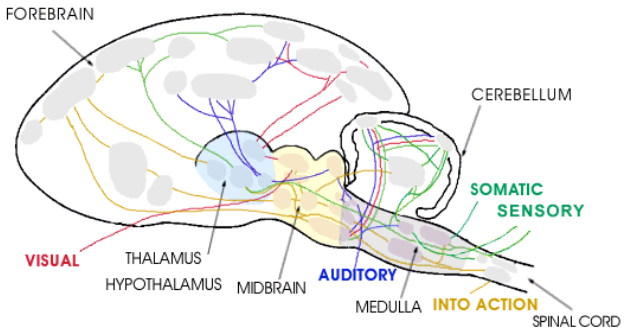


$10^{11}$ neurons
$10^{14}$ synapses
cycle time: $10^{-3}$ sec

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

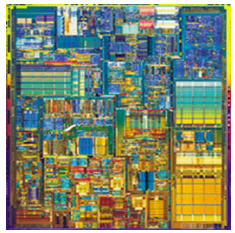

# Computer vs. Brain

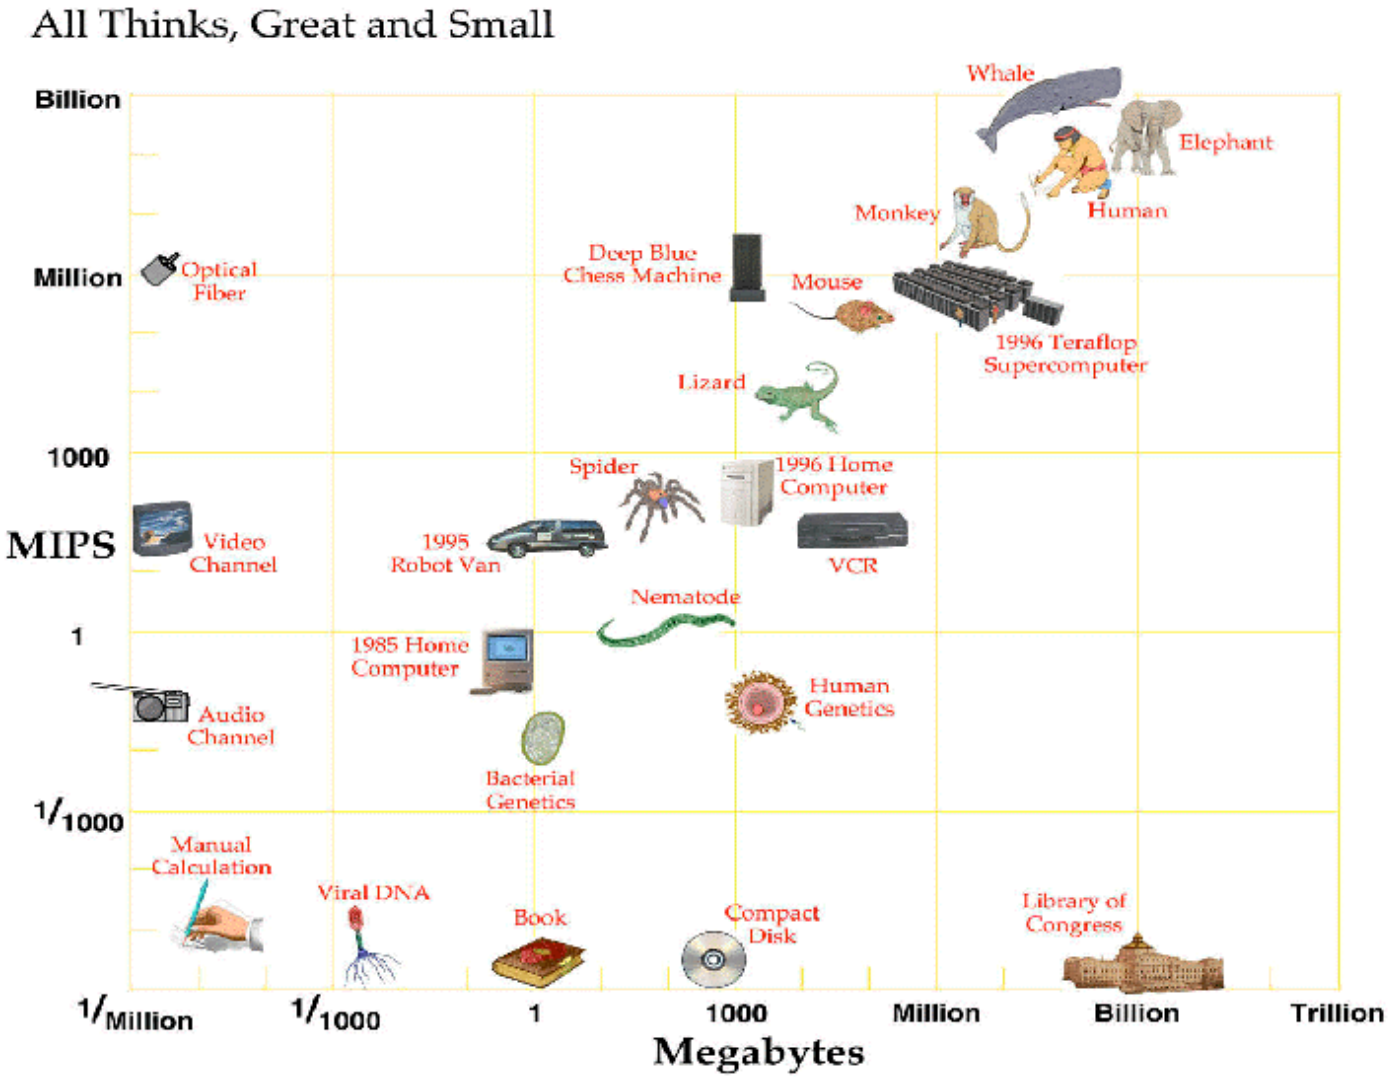

## Conclusion

- In near future we can have computers with as many processing elements as our brain, but:
  - far fewer interconnections (wires or synapses)
  - much faster updates.

**Fundamentally different hardware may require fundamentally different algorithms!**

- Very much an open question.
- Neural net research.

# What Level of Abstraction?

- Hardware (build brains)
- “network” (neural networks?)
- Algorithm + representation
- Intermediate Behavior (cognitive modeling)
- Task Performance (Deep Blue, Turing Test)
- Task Competence (Idealized view)

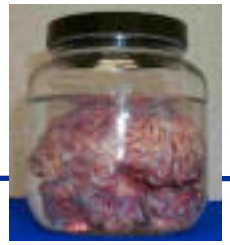

# Classical AI

The principles of intelligence are separate from any hardware / software / wetware implementation

Look for these principles by studying how to perform tasks that require intelligence

Can we rely on simple tasks? (e.g., 8-puzzle, tic tac toe)

# Success Story: Medical Expert Systems

Mycin (1980)

- Expert level performance in diagnosis of blood infections

Today: 1,000's of systems

- Everything from diagnosing cancer to designing dentures
- Often outperform doctors in clinical trials
- Major hurdle today – non-expert part – doctor/ machine interaction

# Success Story: Chess

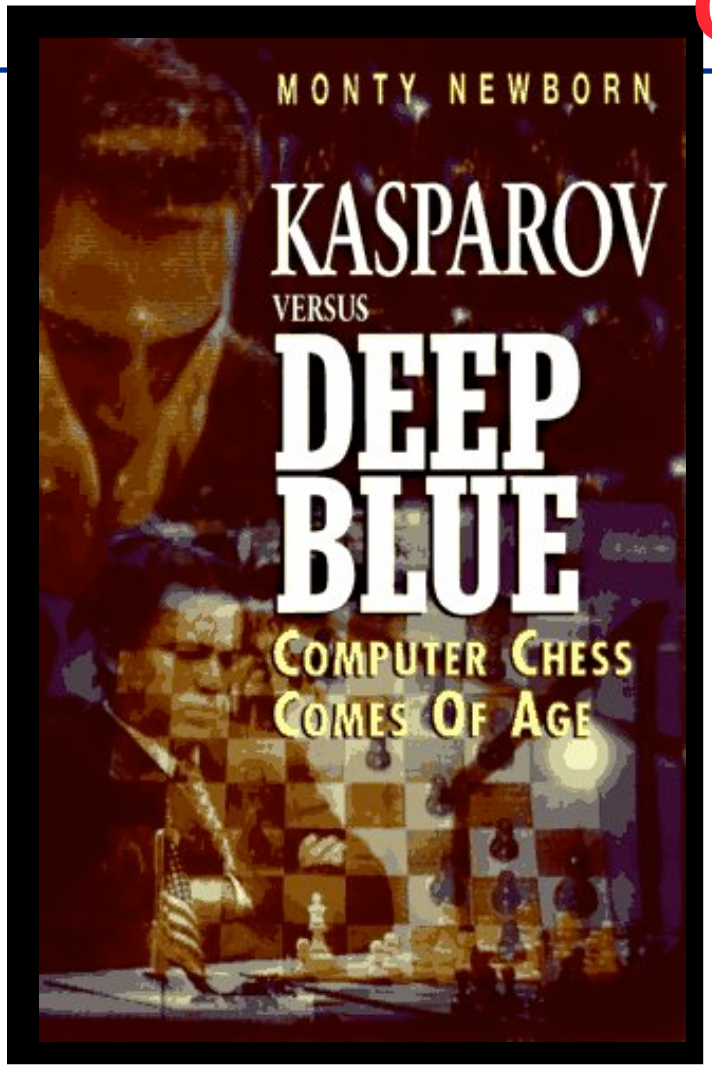


*I could feel – I could smell – a new kind of intelligence across the table*
**- Kasparov**

- Examines 5 billion positions / second
- Intelligent behavior emerges from brute-force search

# 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.

# 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!)

# Key Hard Problem for AI

**Today's successful AI systems**

- **operate in well-defined domains**
- **employ narrow, specialize knowledge**

***Commonsense Knowledge***

- **needed to operate in messy, complex, open-ended worlds**
  - **Your kitchen vs. GM factory floor**
- **understand unconstrained Natural Language**

# Role of Knowledge in Natural Language Understanding

**Speech Recognition**

- **"word spotting" feasible today**
- **continuous speech – rapid progress**
- **turns out that "low level" signal not as ambiguous as we once thought**

**Translation / Understanding**

- **very limited progress**

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

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

# Syntactic, Semantic, Analogical Knowledge

**Time flies like an arrow.**

**Fruit flies like a banana.**

**Fruit flies like a rock.**

# 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!**

**Alternatives?**

- **Open Mind**
- **KnowItAll**

# 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

# Recurrent Themes

## Neural nets vs AI

- McCulloch & Pitts 1943
- Died out in 1960's, revived in 1980's
  - Neural nets vastly simplified model of real neurons, but still useful & practical – massive parallelism
  - particular family of learning and representation techniques

## Logic vs Probability

- In 1950's logic seemed more computationally & expressively attractive (McCarthy, Newell)
  - attempts to extend logic "just a little" to deal with the fact that the world is uncertain!
- 1988 – Judea Pearl's work on Bayes nets
  - provided efficient computational framework
- Today – no longer rivals
  - hot topic: combining probability & first-order logic

# Recurrent Themes, cont.

## Weak vs Strong Methods

- Weak – general search methods
  - A* search, constraint propagation, ...
- Rise of "knowledge intensive" approach
  - expert systems
  - more knowledge, less computation
- Today: resurgence of weak methods
  - desktop supercomputers
  - in highly competitive domains (Chess) exceptions to the general rules are most important!
- How to combine weak and strong methods seamlessly?

# (Re-)Current Themes

- Combinatorial Explosion
- Micro-world successes don't scale up.
- How to Organize and accumulate large amounts of knowledge?
- How to translate from informal, ill-structured statements to formal reasoning (e.g., understand a story)?
- What are reasonable simplifying assumptions?