Introduction to Computer Science

CSCI 109

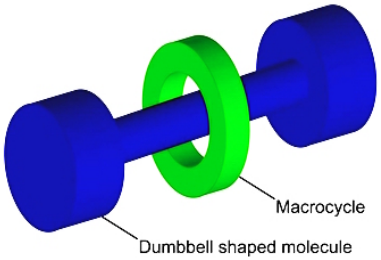

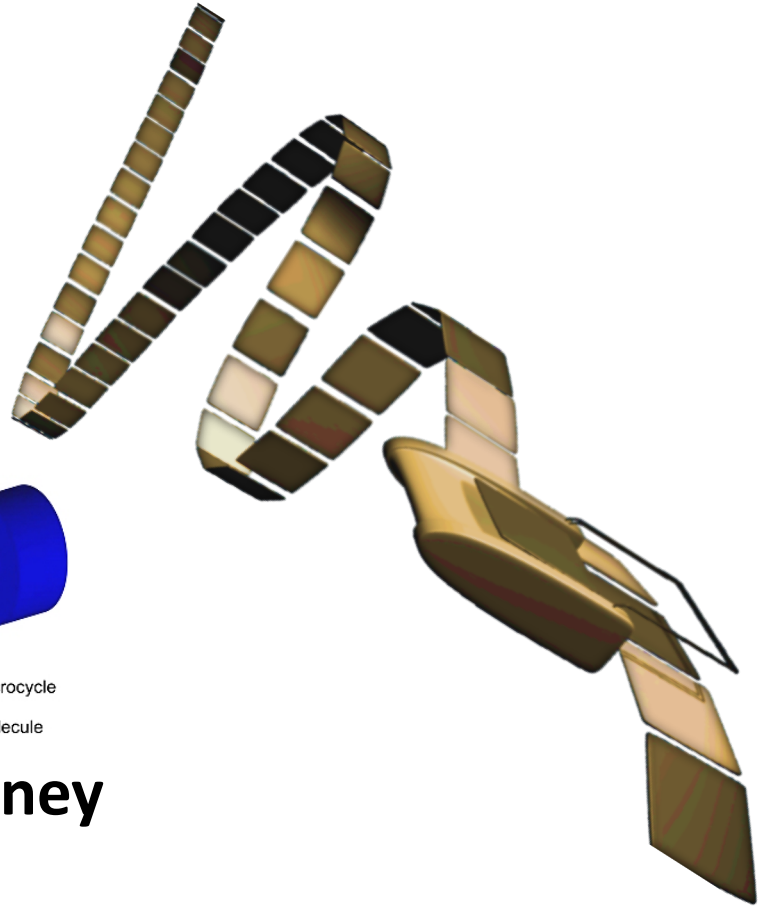

Andrew Goodney

Fall 2019

Lecture 10: Artificial Intelligence *Nov. 11th, 2019*

Schedule

Date	Topic		Assigned	Due	Quizzes/Midterm/Final	Slide Deck
26-Aug	Introduction	What is computing, how did computers come to be?				1
2-Sep	Labor day					
9-Sep	Computer architecture	How is a modern computer built? Basic architecture and assembly	HW1			2
16-Sep	Data structures	Why organize data? Basic structures for organizing data			**Quiz 1** on material taught in class 8/26 and 9/9	3
23-Sep	Data structures	Trees, Graphs and Traversals	HW2	HW1		4
30-Sep	More Algorithms/Data Structures	Recursion and run-time				5
7-Oct	Complexity and combinatorics	How "long" does it take to run an algorithm. P vs NP			**Quiz 2** on material taught in class 9/16 and 9/23	5
14-Oct	Algorithms and programming	Programming, languages and compilers		HW2	**Quiz 3** on material taught in class 9/30	7
21-Oct	Operating systems	What is an OS? Why do you need one?	HW3		**Quiz 4** on material taught in class 10/7	8
28-Oct	Midterm	Midterm			**Midterm** on all material taught so far.	
4-Nov	Computer networks	How are networks organized? How is the Internet organized?		HW3		9
11-Nov	Artificial intelligence	What is AI? Search, plannning and a quick introduction to machine learning			**Quiz 5** on material taught in class 9/4	10
18-Nov	The limits of computation	What can (and can't) be computed?	HW4		**Quiz 6** on material taught in class 11/11	11
25-Nov	Robotics	Robotics: background and modern systems (e.g., self-driving cars)			**Quiz 7** on material taught in class 11/18	12
2-Dec	Summary, recap, review	Summary, recap, review for final		HW4	**Quiz 8** on material taught in class 11/25	13
13-Dec	Final exam 11 am - 1 pm in SGM 123				**Final** on all material covered in the semester	

Reading:
St. Amant Ch. 9

What is Intelligence?

Warm up…

- https://www.youtube.com/watch?v=WnzlbyTZsQY
- https://www.youtube.com/watch?v=vphmJEpLXU0

What is Measured by a Test/Standard

- "Intelligence is what is measured by intelligence tests." (*E. Boring)*
- Thought processes, or behavior, indistinguishable from what a human would produce (at some level of abstraction)
  - *Turing test*

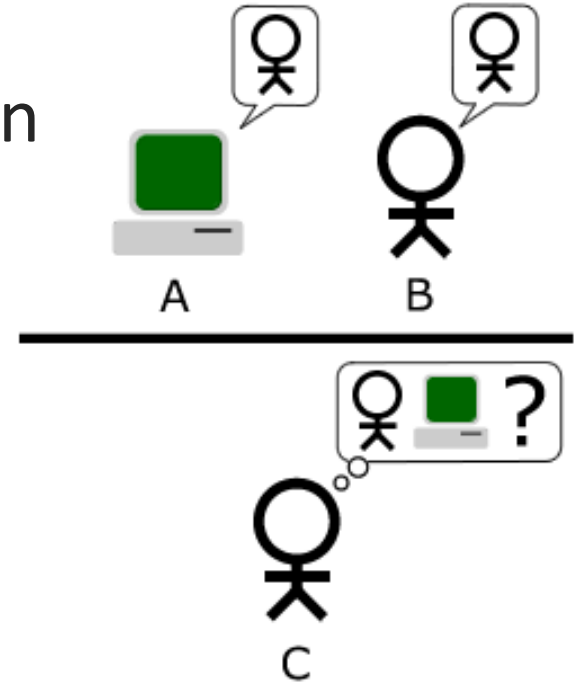

Conglomeration of Specific Capabilities

- "The general mental ability involved in calculating, reasoning, perceiving relationships and analogies, learning quickly, storing and retrieving information, using language fluently, classifying, generalizing, and adjusting to new situations" (*Columbia Encyclopedia)*
- "… a very general mental capability that, among other things, involves the ability to reason, plan, solve problems, think abstractly, comprehend complex ideas, learn quickly and learn from experience." (*Editorial in Intelligence with 52 signatories)*

A Single Focused Capability

- "The capacity to acquire and apply knowledge." (*The American Heritage Dictionary)*
- "The ability to plan and structure one's behavior with an end in view." (*J. P. Das)*
- "… the ability of an organism to solve new problems …" (*W. V. Bingham)*
- "The capacity to learn or to profit by experience." (*W. F. Dearborn)*
- "The ability to carry on abstract thinking." (*L. M. Terman)*
- "… ability to achieve goals in a wide range of environments." *(S. Legg & M. Hutter)*
- … ability to act rationally; that is, "does the 'right thing,' given what it knows." *(S. Russell & P. Norvig)*

Definition of Intelligence

- The common underlying capabilities that enable a system to be general, literate, rational, autonomous and collaborative
  - Can be combined into a ***Cognitive Architecture***
    - Defined in analogy to a computer architecture
    - Provides fixed ("programmable") structure of a *mind*

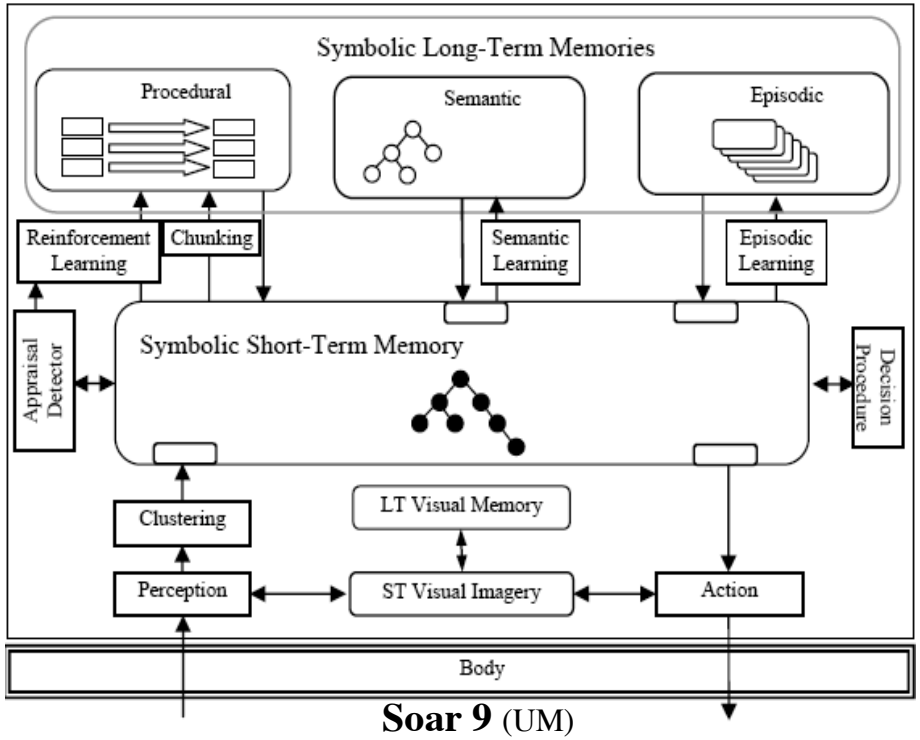

Soar 9 (UM)

The Study of Intelligence

- ***Cognitive Science*** is the interdisciplinary study of mind and intelligence in both natural and artificial systems
  - Although many limit it to just natural systems
- Disciplines involved include
  - **Philosophy**: Questions, concepts and formalisms
  - **Psychology**: Data and theories about natural systems
  - **Linguistics**: Study of language structure and use
  - **Neuroscience**: Data/theory that ground mind in brain
  - **Anthropology**: Intelligence in/across context/culture
  - **Sociology**: Data/theory on natural societies
  - **Computer science**: Study and construction of artificial systems, plus methods for modeling natural systems

What is Artificial Intelligence (AI)?

- Some bad (or perverse) definitions
  - "The study of how to make computers do things at which, at the moment, people are better." *(E. Rich & K. Knight)*
  - "The concept of making computers do tasks once considered to require thinking." *(Medford Police)*
  - "An algorithm by which the computer gives the illusion of thinking like a human." *(D. Gruber)*
  - "Making computers behave like humans." *(Webopedia)*

A Better Definition

- "The scientific understanding of the **mechanisms** underlying thought and intelligent behavior and their embodiment in machines." *(AAAI)*
- Overlaps strongly with Cognitive Science and its various subdisciplines, but also relates to:
  - **Mathematics**: Formalizations and analyses
  - **Economics**: Decision making
  - **Operations research**: Optimization and search
  - **Engineering**: Robotics
- The "what" is too hard, let's study the "how"

Systems of Interest

- Have goals to achieve
  - May concern internal or external situations
  - May be endogenous or exogenous
- Have capabilities to perceive and act in service of their goals
  - For external environments, might include eyes, ears, hands, legs, etc.
  - Or wheels, laser range finders, etc.
- Can embody "knowledge" concerning their goals, capabilities, and situations

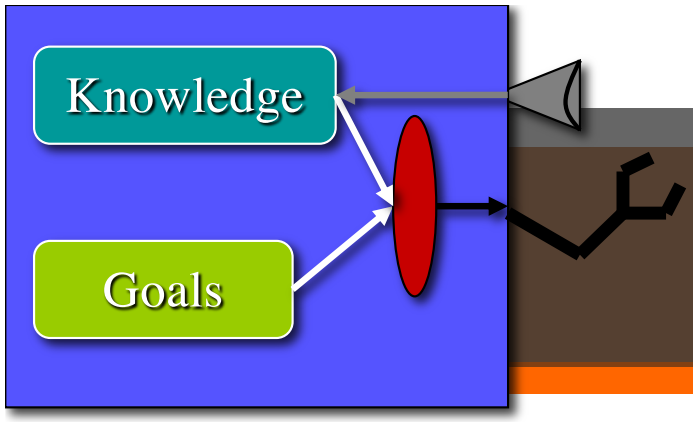

Agents

- Such systems are generally called ***Agents*** (or ***Intelligent Agents***) within AI
  - Differs from notion of agent in Hollywood and in the rest of CS, where the focus is on proxies (or representatives)
- May be embodied as *virtual humans* & *intelligent robots*
- Provides an integrative focus for AI
  - Although most of AI focuses on individual aspects
    - Search and problem solving, knowledge representation and reasoning, planning, machine learning, natural language and speech, vision and robotics, …


Willow Garage PR2


USC/ICT Ada & Grace

Some Relevant Agent Aspects

- **Generality**: Scope of goals and capabilities usable for them
  - Can the agent play both chess and tennis?
  - Can it solve math problems and drive a car?
  - Can it successfully perform full scope of adult human tasks?
- **Literacy**: Extent of knowledge available
  - Ignorance by itself is not lack of intelligence

- **Rationality**: Making best decisions about what to do given goals, knowledge and capabilities
  - Thermostats may be perfectly rational, but with limited generality
- **Autonomy**: Operating without assistance
- **Collaboration**: Working well with others

Some Examples

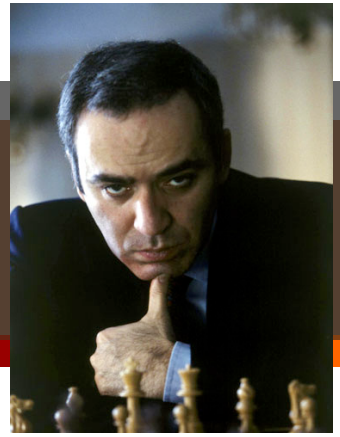

Deep Blue (IBM)

In 1997 Deep Blue became the first machine to win a match against a reigning world chess champion (by 3.5-2.5)

Some Chess Details

- 20 possible start moves, 20 “replies”
- 400 possible positions after 2 ply (1 B and 1 W)
- 197281 positions after 4 ply (2 B and 2 W)
- 7^13 positions after 10 moves
- Approximately 40 legal moves in any position
- Total of about 10^120 number of possible chess games

Search Trees

- Nodes are positions, edges are legal moves
- Leaf nodes are end positions that need to be evaluated
- Leaf nodes that end in check mate for the opponent are good
- Leaf nodes that don’t end in check mate need to be evaluated in some other way
- Each node gets a numeric evaluation score

Minimax: Basic search

- Computer assumes that both W and B play the ‘best’ move.
- Computer plays W and maximizes the score for W
- Choose child node with highest value if W to move
- Choose child node with lowest value if B to move
- About 40 branches at each position in a typical game
- If you want to look d ply ahead you need to search O(b^d)
- Heuristics

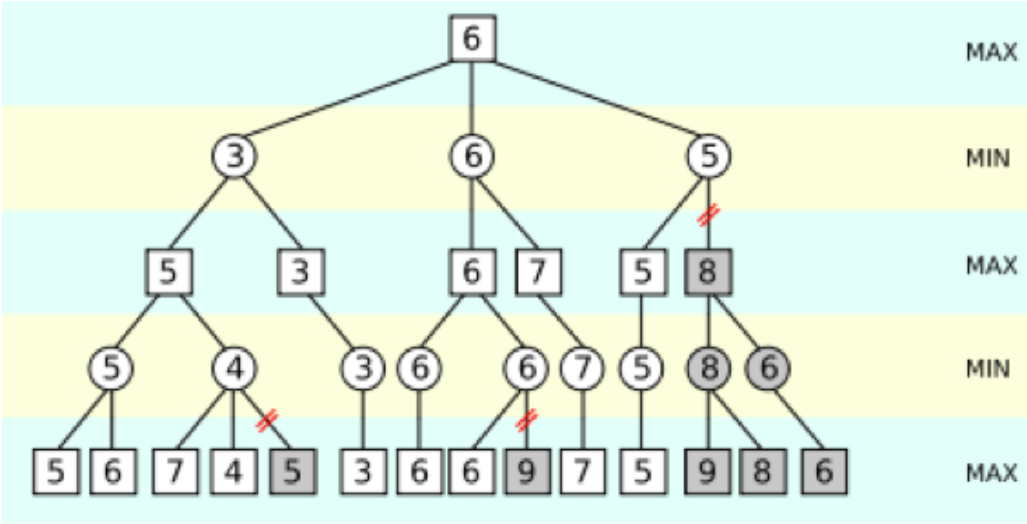

Tree Traversal

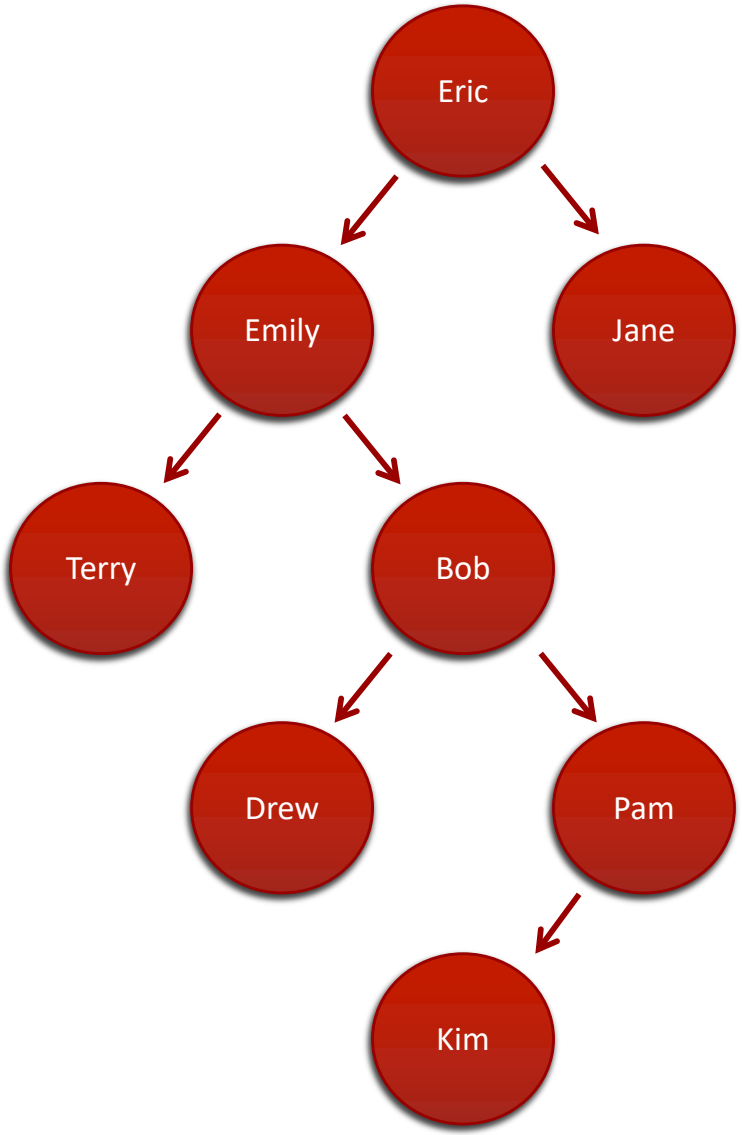

- Depth first traversal
  - Eric, Emily, Terry, Bob, Drew, Pam, Kim, Jane
- Breadth first traversal
  - Eric, Emily, Jane, Terry, Bob, Drew, Pam, Kim
- Best first traversal?
  - Follow edges to your best friend.

Best First Search

OPEN = [initial state] (game states are the nodes of the graph)

CLOSED = []

while OPEN is not empty do

1. Remove the best node from OPEN, call it n, add it to CLOSED.
2. If: n is the goal state, backtrace path to n (through recorded parents) and return path.
3. Else: Create n's successors.
4. For each successor do:
   a. If it is not in CLOSED and it is not in OPEN: evaluate it, add it to OPEN, and record its parent.
   b. Otherwise, if this new path is better than previous one, change its recorded parent.
      i. If it is not in OPEN add it to OPEN.
      ii. Otherwise, adjust its priority in OPEN using this new evaluation.

Greedy Best First Search

- What does it mean “best”?
- Evaluation function is a heuristic that attempts to predict how close the end of a path is to a solution
- Paths which are judged to be closer to a solution are extended first.
- This specific type of search is called greedy best-first search.

A* search: Best-first with $f = g + h$

For every node the evaluation is a knowledge-plus-heuristic cost function *f(x)* to determine the order in which the search visits nodes.

The cost function is a sum of two functions:

- past path-cost function, which is the known distance from the starting node to the current node *x* (usually denoted *g(x)*)
- future path-cost function, which is an admissible "heuristic estimate" of the distance from *x* to the goal (usually denoted *h(x)*).

Admissible means that *h* must not overestimate the distance to the goal.

Deep Blue Combined

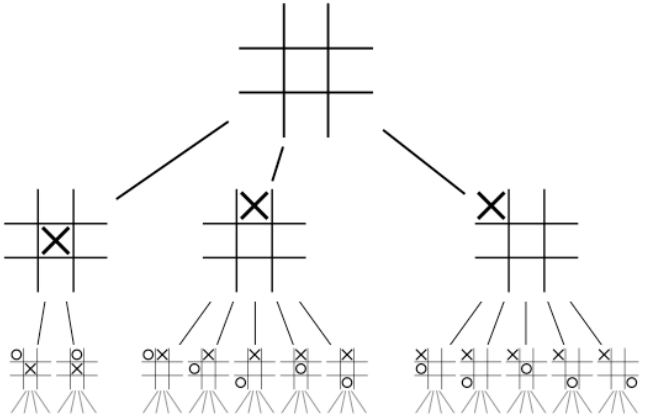

- Parallel and special purpose hardware
  - A 30-node IBM RS/6000, enhanced with
  - 480 special purpose VLSI chess chips
- A heuristic game-tree search algorithm
  - Capable of searching 200M positions/sec (out of 10^{43} total)
  - Searched 6-12 moves deep on average, sometimes to 40
- Chess knowledge
  - An opening book of 4K positions
  - An endgame database for when only 5-6 pieces left
  - A database of 700K GM games
  - An evaluation function with 8K parts and many parameters that were tuned by learning over thousands of Master games

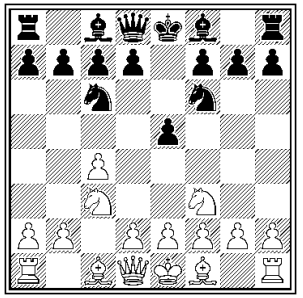

Watson (IBM)

- Compete (and win!) on Jeopardy
  - Question answering (or answer questioning)
- Parallel hardware
  - 2880 IBM POWER7 processor cores with 16 Terabytes of RAM
- Natural language understanding and generation
- A large knowledge base derived via machine learning from 200 million pages

Watson (IBM)

- Search via *generate and test*

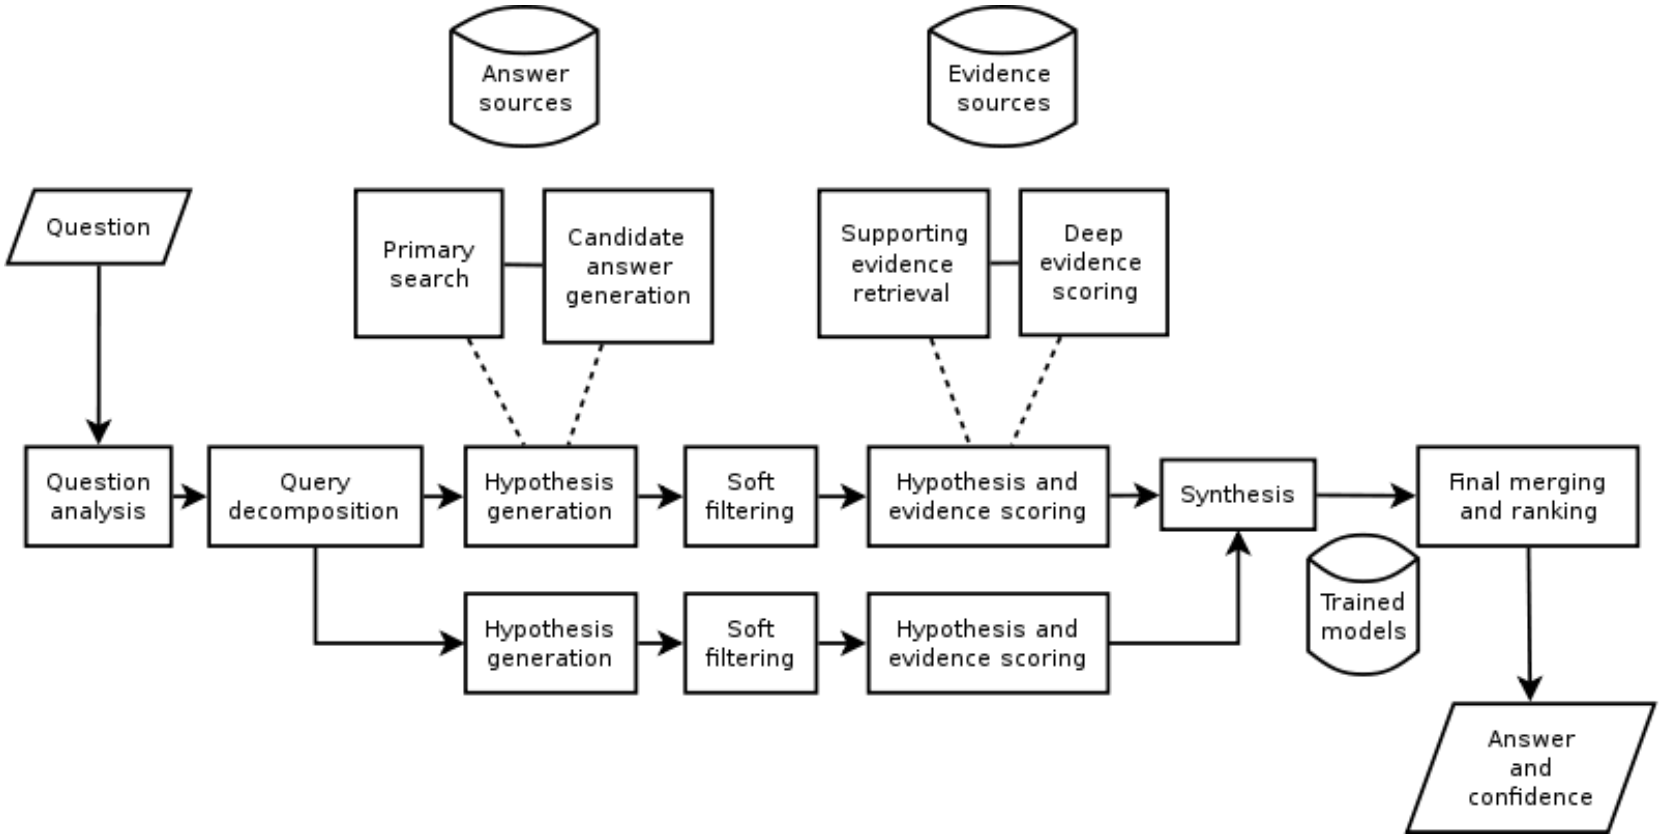

Go

- Players take turns to place black or white stones on a board
- Try to capture the opponent's stones or surround empty space to make points of territory
- Humans play primarily through intuition and feel
- 1,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000 **possible positions**

Google DeepMind AlphaGo

- AlphaGo combines advanced tree search with two deep neural networks
- Advanced tree search is a Monte-Carlo search
- Deep neural networks
  - take a description of the Go board as an input and process it through 12 different network layers containing millions of neuron-like connections
  - “policy network,” selects the next move to play
  - “value network,” predicts the winner of the game

Neural Network Training

- Neural network trained on 30 million moves from games played by human experts, until it could predict the human move 57 percent of the time
- AlphaGo “learned” to discover new strategies, by playing thousands of games between its neural networks, and adjusting the connections in the networks using a trial-and-error process known as reinforcement learning.
- LOTS of computing power -> extensive use of Google Cloud Platform.

Beating the world's top player

- In March 2016 AlphaGo took on Lee Sedol, the world's top Go player, in the Google DeepMind challenge
- Final score: AlphaGo 4 - Lee Sedol 1
- Human: great game play without extensive training
- Machine: better than human game play with orders of magnitude more training and essentially infinite recall

Virtual Humans (USC/ICT)


Ada & Grace


Gunslinger

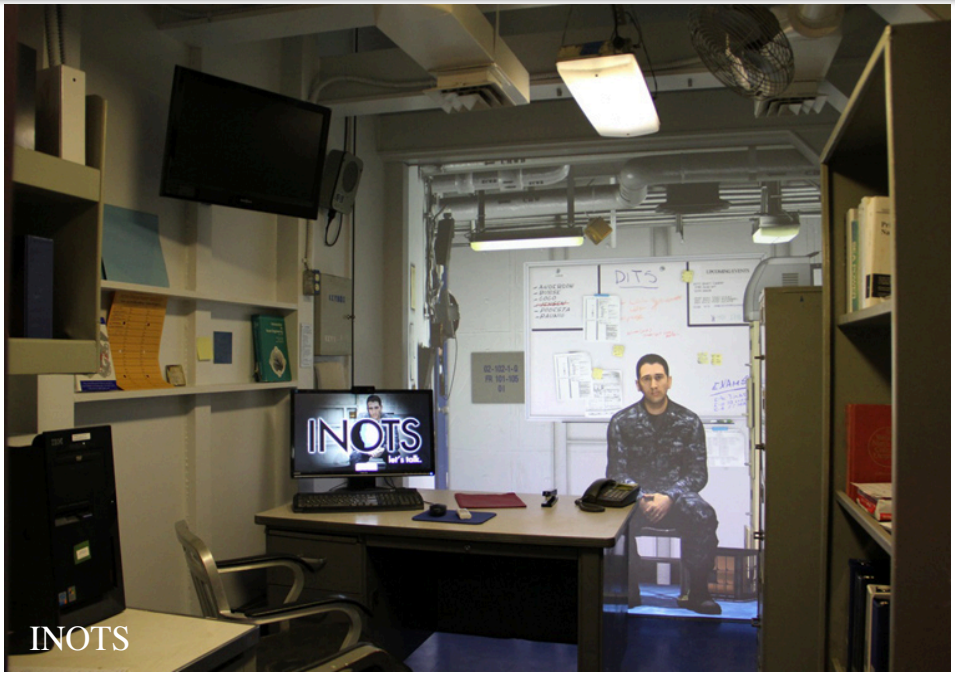

INOTS


SASO

Virtual Humans Combine

- Graphical human bodies with movement and gesture
- Speech, natural language and dialogue
  - May also have ability to visually sense state of human
- Models of actions that can be performed
  - Knowledge about how to choose among them
  - Plans comprising sequences of them
- Emotion models

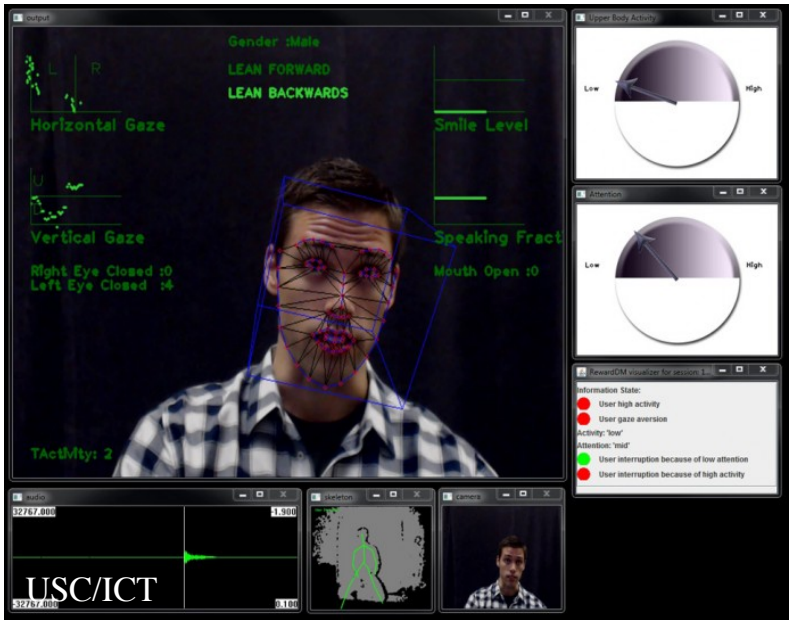

The Big Three Topics within AI

- **Deciding** what to do next
  - Search over possibilities to see which succeed (or are best)
    - A major focus in Deep Blue
    - *Book describes several basic search algorithms*
  - Create and execute plans
    - Used extensively in virtual humans
  - Integrate knowledge about available actions
    - Watson has a major focus on this

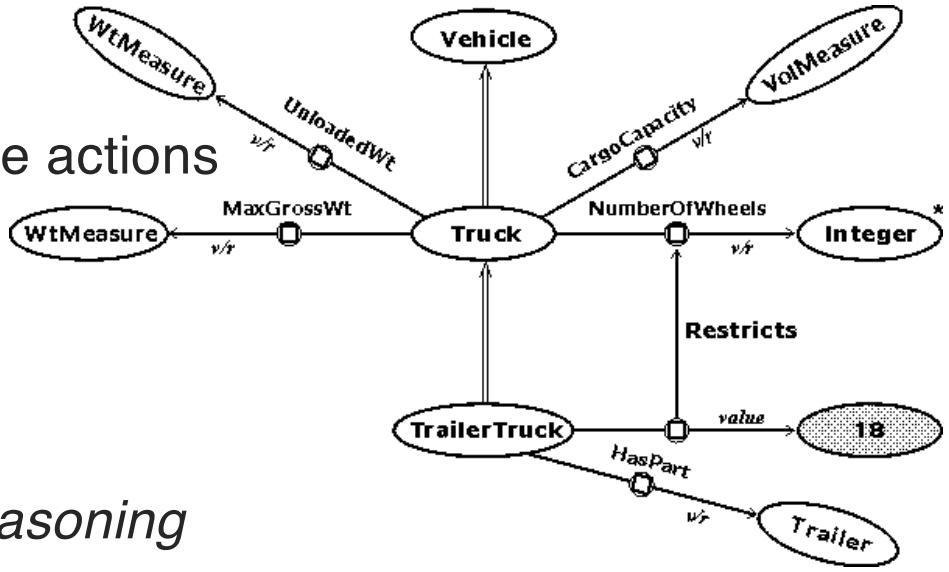

- **Reasoning** about situations
  - Knowledge representation
  - Logical and probabilistic reasoning
  - *Book describes basics of logical reasoning*
- **Learning** from experience and interactions with others
  - Watson and AlphaGo have a major focus on learning
  - *Book describes one basic algorithm*

Others

- Communication
  - Verbal: Speech and natural language
  - Nonverbal: Gesture, expression, …
- Perception
  - Audition, vision, …
- Action (Robotics)
  - Movement/mobility, manipulation (arms and hands)
- Social
  - Cooperative, competitive, …
  - Affect
- Integration (Architectures)
- Applications

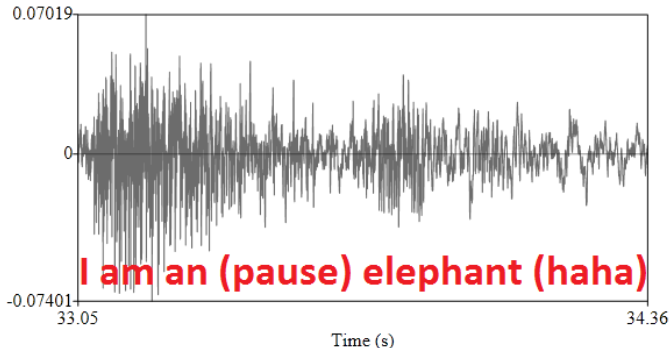

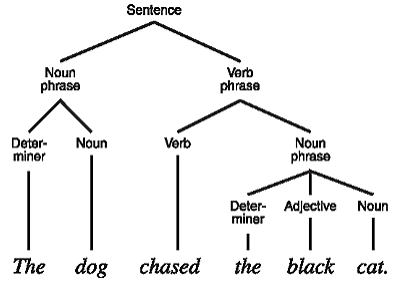

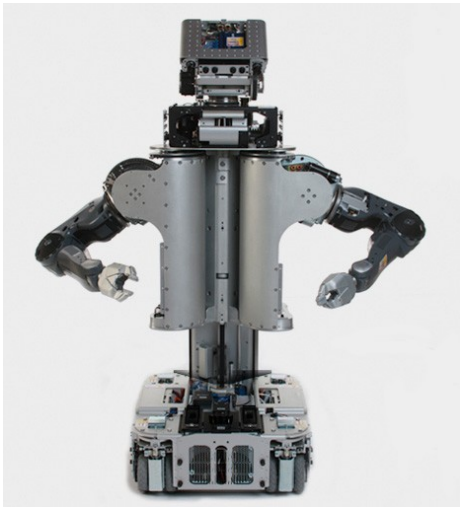

AI vs. Machine Learning

- BOTH extremely hot topics in CS
  - Want to “make a difference” and $200k/yr doing so?
- Often used interchangeably by press, non-Computer Scientists
- Tl;dr
  - AI = Actions
  - Machine Learning = Data
- AI is about actions: an intelligent system (agent) choosing what to do in a “smart” way
- Machine learning is about data: automatically analyzing large amounts of data to discover patterns so predictions can be made when presented with new data
- Many AI systems use algorithms trained with machine learning to inform their decisions

Philosophical Issues

- Is AI Possible?
  - Only act as if intelligent (Weak AI)
  - Can actually be intelligent [Think] (Strong AI)

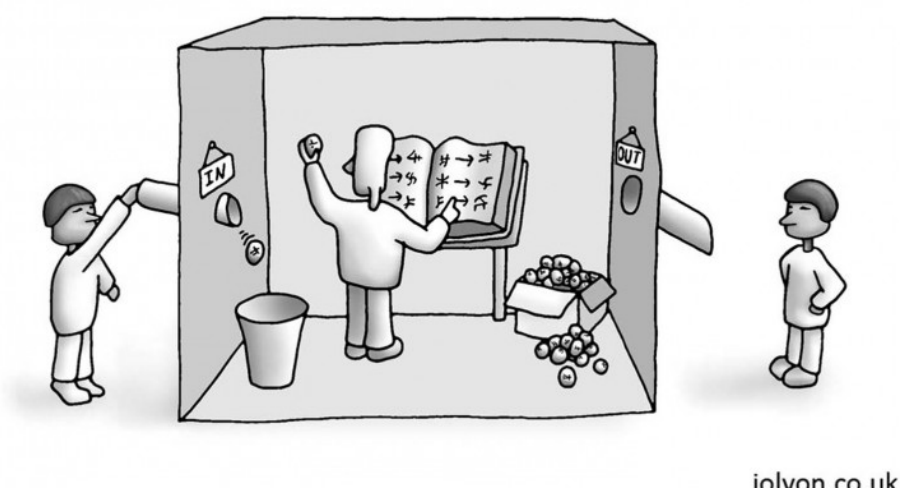

- What are the moral issues in AI?
  - With respect to humans
  - With respect to machines
  - Beyond humans and machines


Borg (Paramount)