Self-improvement for dummies (Machine Learning)

COS 116

4/27/2006

Instructor: Sanjeev Arora

Artificial Intelligence

- Definition of AI (Merriam-Webster):
  1. The capability of a machine to imitate intelligent human behavior (Next time)

Today:

  2. Branch of computer science dealing with the simulation of intelligent behavior in computers

- Definition of Learning:
  - To gain knowledge or understanding of or skill in by study, instruction, or experience

Lecture organization

- Brief look at learning in humans/animals
- Brief look at human brain
- Brief look at how today's machines learn

Caveat: imitating nature may not be best strategy

- Examples:

Birds

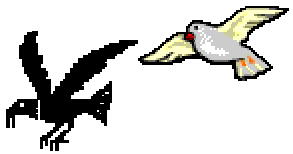

vs

Airplanes

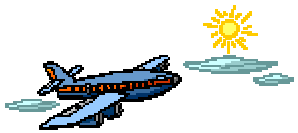

Cheetahs

vs

Race cars

Intelligence in animal world

Is an ant intelligent?

- Build huge, well-structured colonies organized using chemical-based messaging (“Super-organism”)

What about dogs?

Deep mystery: How do higher animals (including humans) learn?

How does  turn into 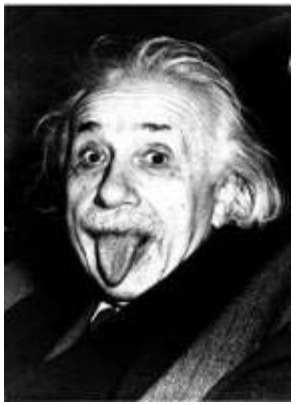

A crude first explanation: Behaviorism [Pavlov 1890's, Skinner 1930's]

- Animals and humans can be understood in a "black box" way as a sum total of all direct conditioning events

- Bell → "Food is coming" →Salivate

- "This person likes me more if I call her "Mama" and that one likes me more if I call him "Papa".

Aside: What does behaviorism imply for societal organization?

More thoughts on behaviorism

Original motivation: “Can’t look inside the working brain anyway, so theory that assumes anything about its working is not scientific or testable.”

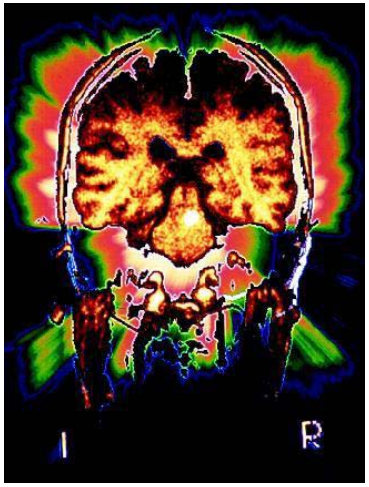

Today,…

Gives little insight into how to design machines with } intelligence. How did dogs, rats, humans sort through sensory experiences to understand reward/punishment?

Chomsky's influential critique of Behaviorism [1957]

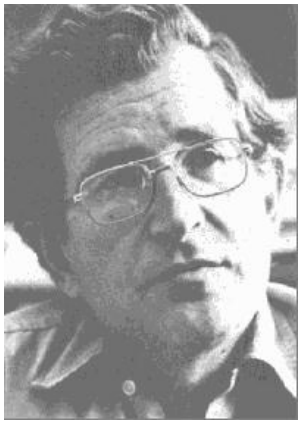

- "Internal mental structures crucial for learning."

  Evidence: universal linguistic rules ("Chomsky grammars"); "self-correction" in language learning, ability to appreciate puns.

1. Brain is "prewired" for language.
2. Must understand mental structures to understand behavior

Presenting:

Your brain

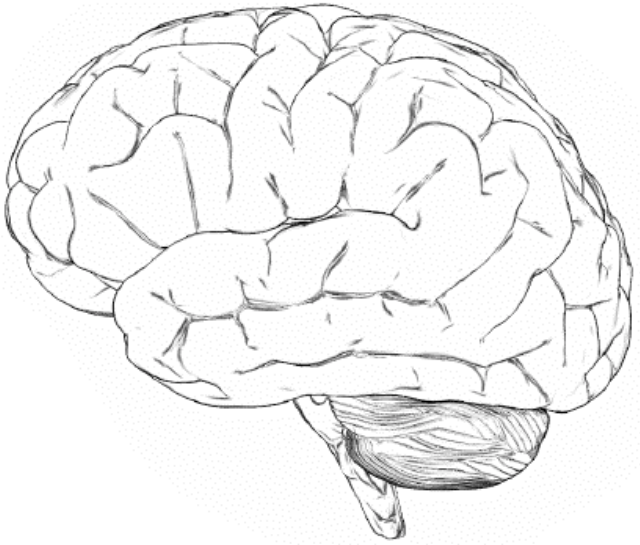

The brain

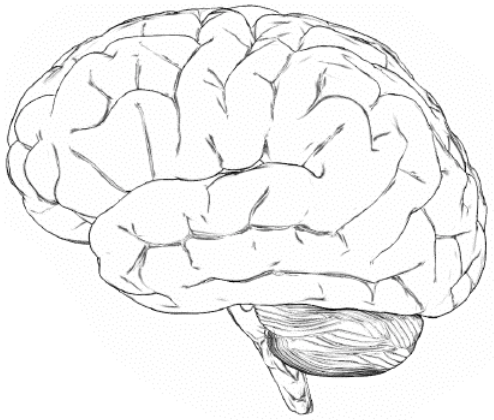

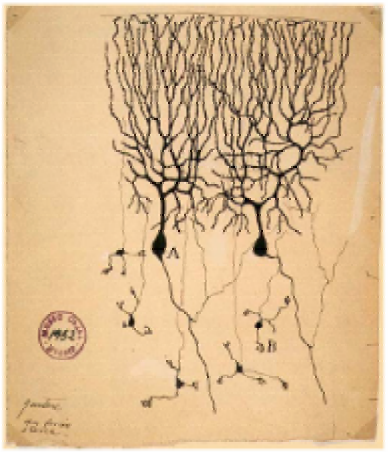

- Network of 100 billion neurons
- Evidence of timing mechanisms (“clock”)
- About 100 firings per second
  - Total of 10^{13} firings (“operations”) per second
  - Number of operations per sec in fast desktop PC: 10^{10}

A comparison

Your brain

10^{11} neurons

Your life on a DVD

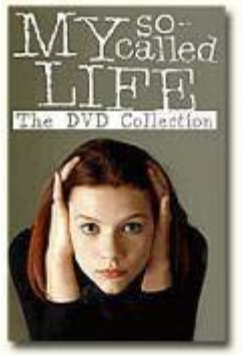

4.3 Gb for 3 hours

> 10^{17} bytes for entire life

Conclusion: Brain must contain structures that compress information and store it in an interconnected way for quick associations and retrieval

A simplistic model of neurons—Neural Net [McCulloch – Pitts 1943]

- Neuron computes "thresholds"

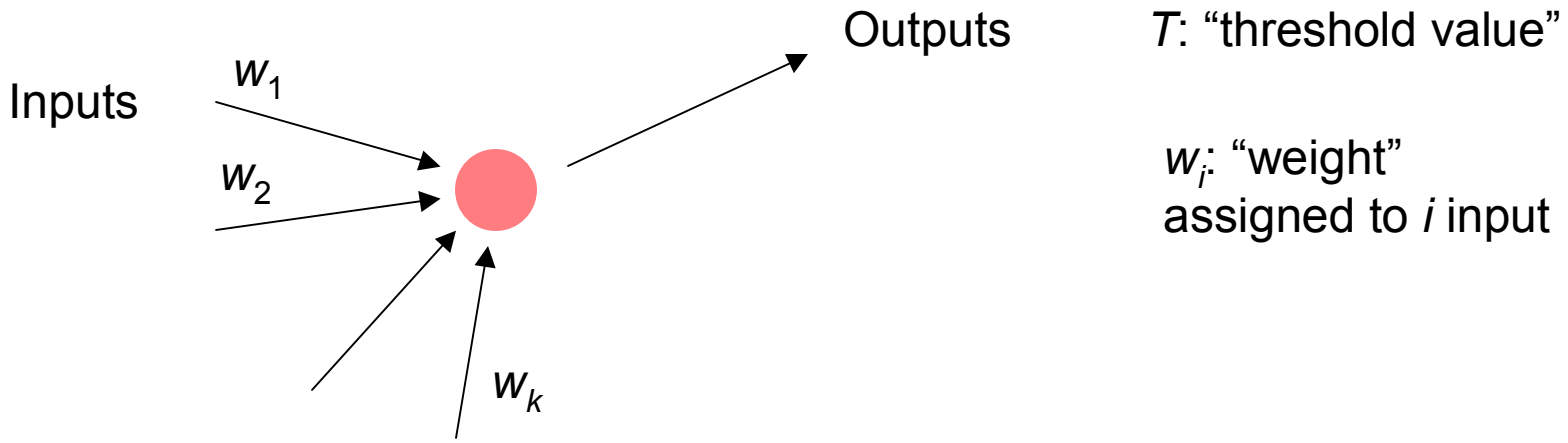

- Take the sum of weights of all neighbors that are firing
- If sum > *T*, fire

Does a neural network model remind you of something??

Why AI is feasible in principle: the simulation argument

- Write a simulation program that simulates all 10^{11} neurons in the brain and their firings.
- For good measure, also simulates underlying chemistry, blood flow, etc.
- Practical difficulty: How to figure out properties (threshold value, w_i's) of each of 10^{10} neurons, the intricate chemistry

Only hope: brain is organized around simpler principles.

Rest of the lecture: Some Principles of machine learning

A machine's "experience" of the world

- *n* sensors, each produces a number
  "experience" = an array of *n* numbers

- Example: video camera: 480 x 640 pixels
  $n = 480 \times 640 = 307200$

- In practice, reduce *n* via compression or preprocessing

Example: Representing wood samples

Brownness scale	1 light	…	10 dark
Texture scale	1 smooth	…	10 rough

(3, 7) = wood that is fairly light brown but kind of on the rough side

A learning task and its mathematical formulation

- Given: 100 samples of oak, maple
- Figure out labeling ("clustering")
- Given a new sample, classify it as oak, maple, or mahogany

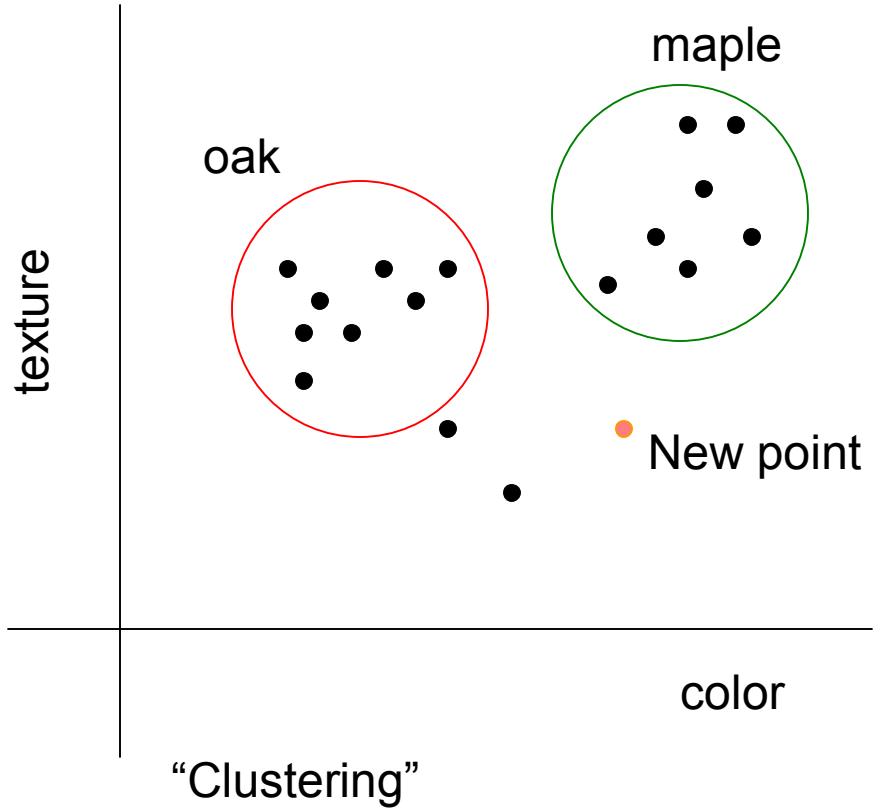

An algorithm to produce 2 clusters

- Some notions:
  - Mean of k points $(x_1, y_1), (x_2, y_2), \ldots, (x_k, y_k)$ is $\left(\frac{x_1 + x_2 + \ldots + x_k}{k}, \frac{y_1 + y_2 + \ldots + y_k}{k}\right)$

    ("center of gravity") 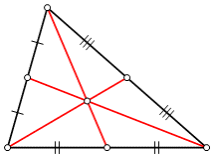
  - Distance between points $(x_1, y_1), (x_2, y_2)$ is $(x_1 - x_2)^2 + (y_1 - y_2)^2$

2-means Algorithm (cont.)

Start by randomly breaking points into 2 clusters

Repeat many times:

{

- Compute means of the current two clusters, say $(a, b), (c, d)$
- Reassign each point to the cluster whose mean is closest to it; this changes the clustering

}

What about learning a more dynamic object?

- Speech?

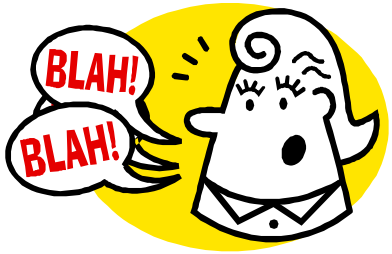

- Motion?

Similar data representation

- Handwriting?

One major idea: modeling uncertainty using probabilities

- Example: Did I just hear "Ice cream" or "I scream"?

- Assign probability ½ to each
- Listen for subsequent phoneme
  - If "is", use knowledge of usage patterns to increase probability of "Ice cream" to 0.9

Probabilities + states = Markovian models

- Markov decision process
- Hidden Markov models

Are “learnt” by machine after extensive training.
(Condensed representation of data corpus)

Rough overview of speech recognition

- Markovian model of language (machine's idea of how language is produced)

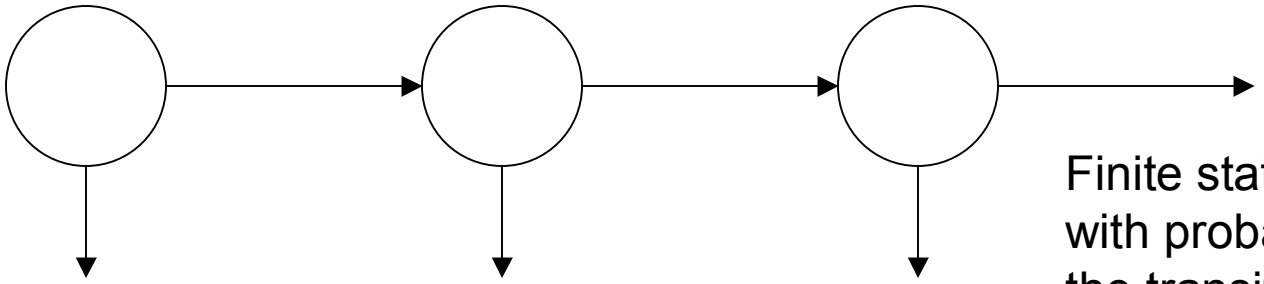

Finite state machine with probabilities on the transitions

- Estimate model parameters using data corpus + user training

Next lecture: Turing test

- Turinghub.com
- Randomly assigns you to chat with a machine or a human
- Note: Machine cannot possibly store answers to all possible 5-minute conversations!