computer science

illuminated

Artificial Intelligence

Nell Dale & John Lewis

(adaptation by Michael Goldwasser)

Computer vs. Humans?

- A computer can do some things better than a human can
  - Adding a thousand four-digit numbers
  - Drawing complex, 3D images
  - Store and retrieve massive amounts of data

Computer vs. Humans?

Let’s reverse the tables.

- Name some things that a human can do better than a computer.
  -
  -
  -

Computer vs. Human

Figure 13.1 A computer might have trouble identifying the cat in this picture.

- Point out the cat in the picture
  - A computer would have difficulty making that identification

Computer vs. Humans?

- Could the following occupations be performed by computers? If so, should they be?
  - Postman
  - Bookstore Clerk
  - Librarian
  - Doctor
  - Lawyer
  - Judge
  - Professor

Artificial Intelligence

- The field of artificial intelligence (AI) is the study of computer systems that attempt to model and apply the intelligence of the human mind

- Of course, first we have to understand why we use the term “intelligence” regarding humans.
  - What defines “intelligence”?
  - Why is it that we assume humans are intelligent?
  - Are monkeys intelligent? Dogs? Ants? Pine trees?

Early History (1950s)

- In 1950 English mathematician Alan Turing wrote a landmark paper titled “Computing Machinery and Intelligence” that asked the question: “Can machines think?”
- Further work came out of a 1956 workshop at Dartmouth sponsored by John McCarthy. In the proposal for that workshop, he coined the phrase a “study of artificial intelligence”
- <u>how would we know if we’ve succeeded?</u>

Can machines think?

- So Turing asked: “Can machines think?” He felt that such machines would eventually be constructed.
- But he also realized a bigger problem. How would we know if we’ve succeeded?

The Turing Test

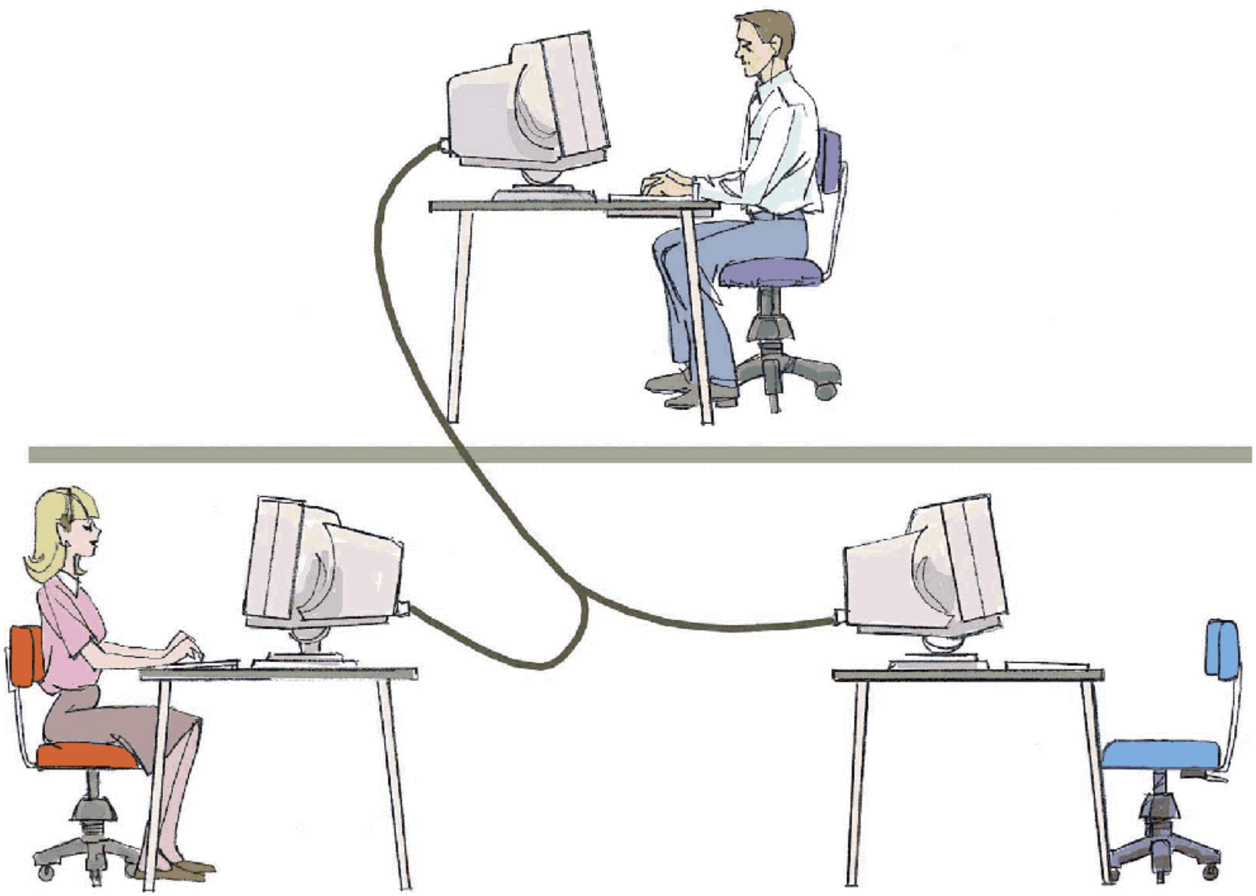

Figure 13.2
In a Turing test, the interrogator must determine which respondent is the computer and which is the human

The Turing Test

- Passing the Turing Test does not truly show that the machine was thinking. It simply shows that it generated behavior consistent with thinking.
- **weak equivalence:** the two systems (human and computer) are equivalent in **results** (output), but they do not necessarily arrive at those results in the same way
- **Strong equivalence:** the two systems use the same internal processes to produce results

Overview of Issues

- We want to compare the way that computers and humans work to see if we can better understand why each have their (computational) strengths.
  - Processing Models
  - Knowledge Representation
  - Reasoning

The Human Brain

- Let’s first look at how a biological neural network works
  - A neuron is a single cell that conducts a chemically-based electronic signal
  - At any point in time a neuron is in either an excited or inhibited state

The Human Brain

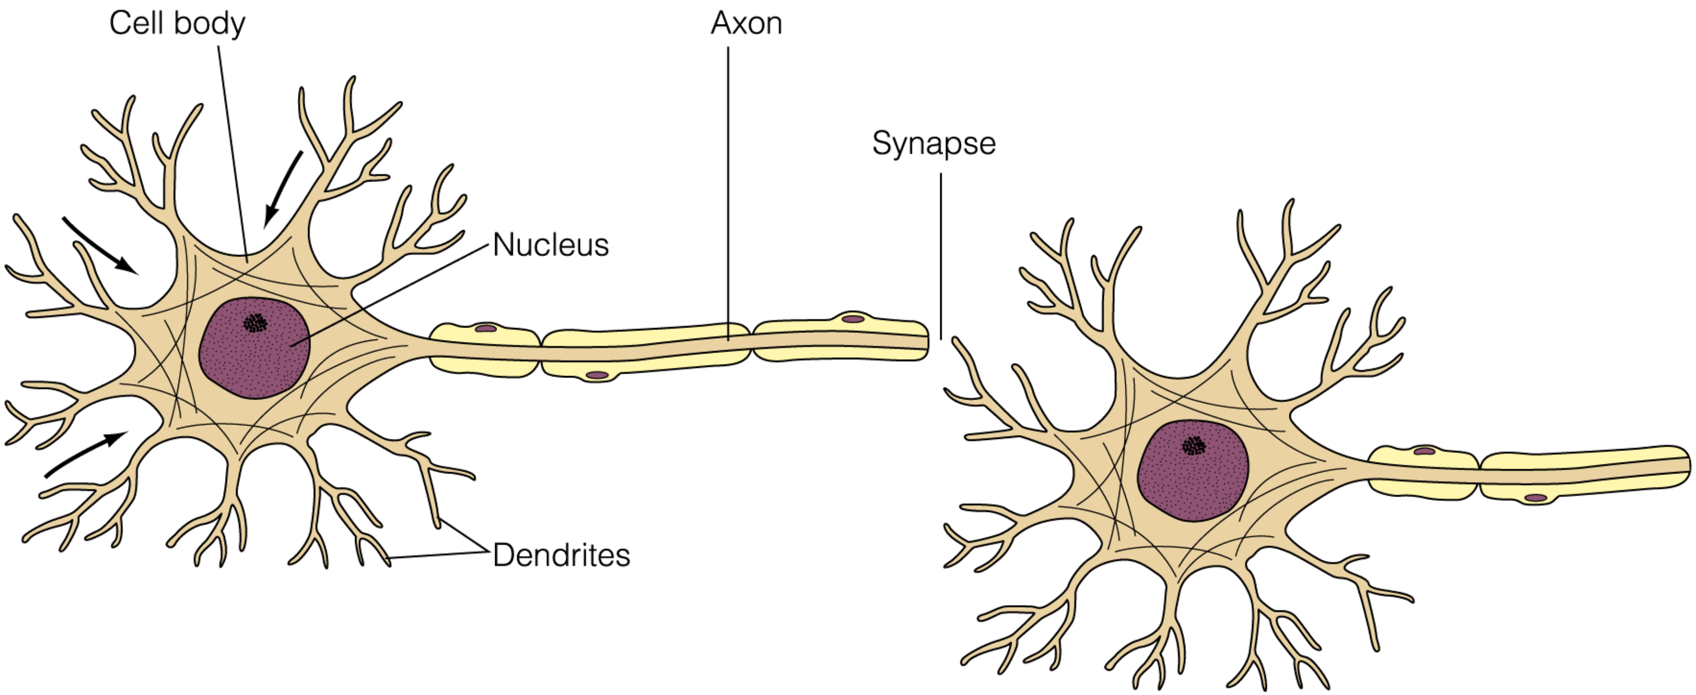

Figure 13.6 A biological neuron

The Human Brain

- A series of connected neurons forms a pathway
- A series of excited neurons creates a strong pathway
- A biological neuron has multiple input tentacles called dendrites and one primary output tentacle called an axon
- The gap between an axon and a dendrite is called a synapse

The Human Brain

- A neuron accepts multiple input signals and then controls the contribution of each signal based on the “importance” the corresponding synapse gives to it
- The pathways along the neural nets are in a constant state of flux
- As we learn new things, new strong neural pathways in our brain are formed

Artificial Neural Networks

Some have tried to use computers to mimic the neural network model of the human brain.

- Each processing element in an artificial neural net is analogous to a biological neuron
  - An element accepts a certain number of input values and produces a single output value of either 0 or 1
  - Associated with each input value is a numeric weight

Artificial Neural Networks

– The **effective weight** of the element is defined to be the sum of the weights multiplied by their respective input values

$v1*w1 + v2*w2 + v3*w3$

– Each element has a numeric threshold value

– If the effective weight exceeds the threshold, the unit produces an output value of 1

– If it does not exceed the threshold, it produces an output value of 0

Artificial Neural Networks

- The process of adjusting the weights and threshold values in a neural net is called **training**
- A neural net can presumably be trained to produce whatever results are required

(Ay, there’s the rub)

Human vs. Computer

	Human Brain	Massively Parallel Computer
speed	Neurotransmitters travel at rate of perhaps 1000 feet/second.	Electrons travel at speed of light (186000 miles/second)
memory	Roughly 100 billion neurons; estimated to represent equivalent of 50 trillion bits.	The top supercomputers today might approach this much memory.
communication	Each neuron connected to roughly 1000 other neurons	Processor perhaps connected to up to 100 other processors

Expert Systems

- A **knowledge-based system** is a software system that embodies and uses a specific set of information from which it extracts and processes particular pieces
- **Expert system:** usually embodies the knowledge of a specialized field, modeling the expertise of a professional in that field
  - An expert system uses a set of rules to guide its processing
  - The inference engine is the part of the software that determines how the rules are followed

Expert Systems

- Example: What type of treatment should I put on my lawn?
  - NONE—apply no treatment at this time
  - TURF—apply a turf-building treatment
  - WEED—apply a weed-killing treatment
  - BUG—apply a bug-killing treatment
  - FEED—apply a basic fertilizer treatment
  - WEEDFEED—apply a weed-killing and fertilizer combination treatment

Expert Systems (cont.)

- Boolean variables
  - BARE—the lawn has large, bare areas
  - SPARSE—the lawn is generally thin
  - WEEDS—the lawn contains many weeds
  - BUGS—the lawn shows evidence of bugs
- Other Variables
  - LAST – the last date of lawn treatment
  - CURRENT – the current date
  - SEASON – the current season

Expert Systems (cont.)

- Some rules
  - if (CURRENT – LAST < 30) then NONE
  - if (SEASON = winter) then not BUGS
  - if (BARE) then TURF
  - if (SPARSE and not WEEDS) then FEED
  - if (BUGS and not SPARSE) then BUG
  - if (WEEDS and not SPARSE) then WEED
  - if (WEEDS and SPARSE) then WEEDFEED

Expert Systems (cont.)

- An execution of our inference engine
  - System: Does the lawn have large, bare areas?
  - User: No
  - System: Does the lawn show evidence of bugs?
  - User: No
  - System: Is the lawn generally thin?
  - User: Yes
  - System: Does the lawn contain significant weeds?
  - User: Yes
  - System: You should apply a weed-killing and fertilizer combination treatment.

Natural Language Processing

- There are three basic types of processing going on during human/computer voice interaction
  - Voice recognition—recognizing human words
  - Natural language comprehension—interpreting human communication
  - Voice synthesis—recreating human speech
- Common to all of these problems is the fact that we are using a natural language, which can be any language that humans use to communicate

Voice Synthesis

- There are two basic approaches to the solution
  - Dynamic voice generation
  - Recorded speech
- To generate voice output using dynamic voice generation, a computer examines the letters that make up a word and produces the sequence of sounds that correspond to those letters in an attempt to vocalize the word
- Human speech has been categorized into specific sound units called phonemes

Voice Synthesis

Consonants			
Symbols	Examples	Symbols	Examples
p	pipe	k	kick, cat
b	babe	g	get
m	maim	ŋ	sing
f	fee, phone, rough	š	shoe, ash, sugar
v	vie, love	ž	measure
θ	thin, bath	č	chat, batch
ð	the, bathe	ǰ	jaw, judge, gin
t	tea, beat	d	day, bad
n	nine	ʔ	uh uh
l	law, ball	s	see, less, city
r	run, bar	z	zoo, booze

Vowels	
Symbols	Examples
i	eel, sea, see
I	ill, bill
e	ale, aim, day
ɛ	elk, bet, bear
æ	at, mat
u	due, new, zoo
ʊ	book, sugar
o	own, no, know
ɔ	aw, crawl, law, dog
a	hot, bar, dart
ə	sir, nerd, bird
ʌ	cut, bun

Semi Vowels	
w	we
h	he
j	you, beyond

Dipthongs	
aj	bite, fight
aw	out, cow
ɔj	boy, boil

Figure 13.7 Phonemes for American English

Voice Synthesis

- The other approach to voice synthesis is to play digital recordings of a human voice saying specific words
  - Telephone voice mail systems often use this approach: “Press 1 to leave a message for Alex Wakefield”

Voice Synthesis (cont.)

- Each word or phrase needed must be recorded separately
- Furthermore, since words are pronounced differently in different contexts, some words may have to be recorded multiple times
  - For example, a word at the end of a question rises in pitch compared to its use in the middle of a sentence

Voice Recognition

- The sounds that each person makes when speaking are unique
- We each have a unique shape to our mouth, tongue, throat, and nasal cavities that affect the pitch and resonance of our spoken voice
- Speech impediments, mumbling, volume, regional accents, and the health of the speaker further complicate this problem

Voice Recognition (cont.)

- Furthermore, humans speak in a continuous, flowing manner
  - Words are strung together into sentences
  - Sometimes it's difficult to distinguish between phrases like "ice cream" and "I scream"
  - Also, homonyms such as "I" and "eye" or "see" and "sea"
- Humans can often clarify these situations by the context of the sentence, but that processing requires another level of comprehension
- Modern voice-recognition systems still do not do well with continuous, conversational speech

Natural Language Comprehension

- Even if a computer recognizes the words that are spoken, it is another task entirely to understand the meaning of those words
- Natural language is often ambiguous, for a variety of reasons. Let's look at several classes of ambiguity (though admittedly there is some overlap in such a classification)

Lexical Ambiguity

- A single word can have two meanings.
  - *“The bat slipped from his hand”*
  - *“Cinderella had a ball”*
  - *“Ron lies asleep in his bed”*
- Worse yet, those meanings may even constitute different parts of speech.
  - *“Time flies like an arrow”*
  - *“They are racing horses”*
  - *“Stampeding cattle can be dangerous”*

Syntactic Ambiguity

- Even if all words have a clear meaning, ambiguity may exist because the phrases can be combined in several ways when parsing.
  - *“I saw the Grand Canyon flying to New York”*
  - *“The clams are ready to eat”*
  - *“I saw the man in the park with the telescope”*

Referential Ambiguity

- Pronouns may cause ambiguity when it is not clear which noun is being referenced.

  - *"The brick fell on the computer but it is not broken"*

  - *"Jon met Bill before he went to the store"*

Rules of Conversation

– *"Do you know what time it is?"*
  - Presumably, a correct response is "Yes."
  - Does this sentence's meaning change if it is said by your boss when you walk into a meeting 30 minutes late?

– *"Do you know you have a flat tire?"*

– *"I'd like to ask everyone to raise their hand for two seconds."*
  - How did you respond?

Need for “real world” knowledge

– *“Norman Rockwell painted people”*
  - Did he do tattoos? Face-painting?