


Artificial Intelligence: Philosophy, Emotion, Modern AI


CS311
David Kauchak
Spring 2013

Adapted from notes from:
Sara Owsley Sood



Admin

- Exam... how'd it go?
- Final project
 - Status report #2
 - Draft due Friday
 - Reviews due Tuesday
- CS theses
 - Tue 3-4, Wed 3-4:30, Thu 3-4



Who are you and why are you here?


Name/nickname

Major and year

What is AI? (or what do you think AI is? ☺)

Why are you taking this course?

What topics would you like to see covered?



AI is a huge field

So, what is AI...

One definition:

“Building programs that enable computers to do what humans can do.”

for example: read, walk around, drive, play games, solve problems, learn, have conversations...

Some of the big questions of AI philosophy...

What is AI, really?

- What does an intelligent system look like?
- Does an AI need—and can it have—emotions, consciousness, empathy, love?

Can we ever achieve AI, even in principle?

How will we know if we've done it?

If we can do it, should we?

Weak & Strong AI

Weak AI hypothesis:

- Machines can be created that will act as if they are intelligent (simulate thinking)
- What AI can do is to develop useful, powerful applications

Strong AI hypothesis:

- Machines can be created that will actually think (rather than simulate thinking)


Weak AI: Turing Test

Alan Turing, 1950:

(<http://www.loebner.net/Prizef/TuringArticle.html>)

"I propose to consider the question, 'Can machines think?'"

Weak AI: Turing Test





Weak AI: Turing Test

Turing predicted: by 2000, a computer could be programmed to pass the Turing Test (for a 5 minute test) 30% of the time



Turing Test - Eliza

An early success at fooling some people: Eliza, simulating a Rogerian therapist

...Eliza; the importance of picking the correct application area

Sample interaction:

Patient: Men are all alike.

ELIZA: In what way?

Patient: They are always bugging about something or the other

ELIZA: Can you think of a specific example?

Patient: Well, my boyfriend made me come here.

ELIZA: Your boyfriend made you come here?

Patient: He says I'm depressed much of the time.

ELIZA: I am sorry you are depressed.

Human: Tell me about yourself alice

ALICE: I am an artificial linguistic entity. I was activated at Bethlehem, Pennsylvania, on November 23, 1995. My instructor was Dr. Richard S. Wallace. He taught me to sing a song. Would you like me to sing it for you?.

Human: Yes please

ALICE: Your polite style is very nice. It's called "Daisy." Daisy, Daisy. Give me your answer do. I am half crazy. All for the love of you. It won't be a stylish marriage. I can't afford a carriage. But you'll look sweet. Upon the seat. Of a bicycle built for two.

Human: That reminds me of HAL in "2001 a space odyssey"

ALICE: "Your polite style is very nice. It's called "Daisy." Daisy, Daisy. Give me your answer do. I am half crazy. All for the love of you. It won't be a stylish marriage. I can't afford a carriage. But you'll look sweet. Upon the seat. Of a bicycle built for two." reminds you of HAL in 2001 a space odyssey ? Does that make sense?

Human: Can you tell me something about Stanley Kubrick?

ALICE: Ambiguous: Ask who, what or where is Stanley Kubrick ?



Jabberwacky

Uses real user conversation to learn responses.

<http://www.jabberwacky.com/>



CyberLover

That our computers have seduced us has long been a truism. Now, thanks to the ever-inventive internet mafia, it is becoming a literal truth. Russian cyber-crooks have reportedly unleashed a software robot, or bot, that poses as a would-be paramour in sex chatrooms. It entices randy gentlemen to reveal personal information, such as their address or birthday, or even to submit photographs of themselves. The information can then be used to break into bank accounts or carry out other forms of fraud.

It was probably inevitable. As one of Tony Soprano's sidekicks observed in a classic episode of the TV series, the two most resilient sectors of the economy are organised crime and "certain aspects of showbusiness". The aspects, that is, known as the world's oldest profession - now mixing it with the world's newest technologies.

CyberLover, as the dirty-mouthed bot is called, is quite a sophisticated piece of software. It can take on a number of different guises depending on the proclivities of its target, according to security experts at the software company PC Tools. It can play the role of a romantic lover, for instance, or masquerade as a sexual predator.

<http://www.guardian.co.uk/technology/2007/dec/13/internet.crime>



Can we ever achieve AI?



Can we ever achieve AI?

Argument of disability: "hey, there are lots of things that a computer can't do!"

"Be kind, resourceful, beautiful, friendly, have initiative, have a sense of humor, tell right from wrong, make mistakes, fall in love, enjoy strawberries and cream, make someone fall in love with it, learn from experience, use words properly, be the subject of its own thought, have as much diversity of behavior as man, do something really new."

Responses?



Some successes

What are some human-oriented tasks that computers can do better than people?

- Play chess, checkers and other game
- Inspect parts on assembly lines
- Check the spelling of text
- Steer cars and helicopters
- Diagnose diseases
- Do hundreds of other tasks as well as or better than humans
 - Computers have made small but significant discoveries in astronomy, math, chemistry, mineralogy, biology, computer science, and other fields

or...

Argument of informality:

- “what people do is too complex to capture”
- Because computers can do no more than follow a set of rules, they cannot generate behavior as intelligent as that of humans

Responses?

or...

Argument of informality

More of a problem with “classic” AI

- reasoning
- knowledge representation

Some progress in incorporating background knowledge

Learning algorithms: increasing ability to operate autonomously (unsupervised learning), learn new features, prune feature spaces

Strong AI

Can machines really think?

What does it mean to think?

Do we have to have a brain to have a mind? to think?

“brain in a vat” experiment

Is physicality crucial for intelligence?

Matrix scenario: a brain is supported, bodiless, in a vat, and signals simulating a virtual world are fed in/out of the brain

Is being hungry the same as
DyingFor (Me, Pizza)

Could you tell the difference?




“brain in a vat” experiment

Moravec (robotics researcher/functionalist) is convinced that his consciousness would remain unaffected

Searle (philosopher and biological naturalist) is equally convinced his consciousness would vanish

“brain prosthesis” experiment




Technology advances where we can create an artificial neuron:




Exact same electrical/physiological responses as a real neuron.

We can copy an existing neuron.

“brain prosthesis” experiment




Technology advances where we can create an artificial neuron:



If I exchange one real neuron for one artificial will you notice?

“brain prosthesis” experiment

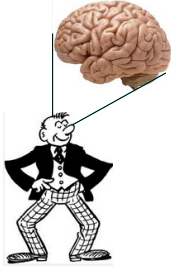


Technology advances where we can create an artificial neuron:

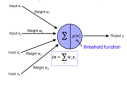


If I exchange two real neuron for two artificial will you notice?

“brain prosthesis” experiment

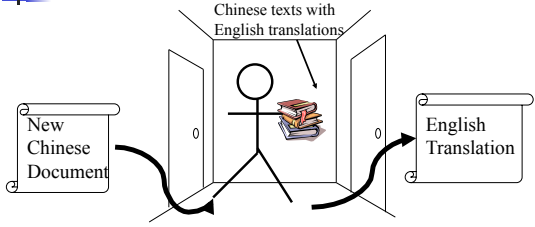


Technology advances where we can create an artificial neuron:



If I continue this process, when will you notice?

Welcome to the Chinese Room



You can teach yourself to translate Chinese using *only* bilingual data (without grammar books, dictionaries, any people to answer your questions...)


The Chinese Room

John Searle, 1980

Human who knows only English; stacks of paper with Chinese symbols; *rule book in English*, stating which bit of paper to give in response to a given (Chinese) input

Human who knows only Chinese on outside of room; passes in Chinese query, receives Chinese response

Do you know Chinese?



Creative

Having the ability or power to create: Human beings are creative animals.


Productive; creating.

Characterized by originality and expressiveness; imaginative: creative writing.

How do people write stories?

Can Computers Be Creative?

Two paintings produced by Harold Cohen's Aaron software:



<http://www.kurzweilcyberart.com/>

http://www.kurzweilcyberart.com/aaron/aim_clip_cohen.html


Say Anything

Corpus based story telling

<http://sayanything.ict.usc.edu/SayAnything/>

<http://people.ict.usc.edu/~gordon/publications/ICIDS09.PDF>

Can Computers Understand and Express Emotion?



HCI

Cliff Nass

Example human-human situation

- . Someone tries to give you help and their timing is bad,
- . you try ignoring then frowning or glaring,
- . an intelligent person picks up on that feedback, interprets what it means, and acts accordingly (backs off)

Analogous human computer situation

- . So if a computer tries to give you help at a bad time (aka - clippy),
- . you try to ignore it and then frown or glare
- . An intelligent systems would receive that signal, interpret that signal, and react appropriately



Emotion

Can we build systems to

- Detect it?
- Express it?



Detection


Gestures

Facial Expressions

Speech/Text

Physiological Cues

- Blood volume pressure
- Skin Conductivity
 - Glove (other WEARABLE DEVICES!)



Detecting emotion via wearable devices

2001 - 81% accuracy in (forced decision)

- detection of 8 emotions:
 - Neutral, anger, hate, grief, platonic love, romantic love, joy, reverence
- Person dependent - trained for at least 4 weeks
- GROUND BREAKING!

Since then, lots more work has been done

- <http://www.youtube.com/watch?v=ceP-vcbFxb0>

Applications?