

TURING TEST AND
NATURAL LANGUAGE
PROCESSING

ANNOUNCEMENT

- **First HW due this Thursday before class! Submit via Canvas as PDF.**
- TAs and TA office hours
 - Ross Dimassimo (cs1060 AT spam.im)
 - William Garnes III (wagarnes AT gmail.com)
 - Vikram Raj (vikram.raj AT utah.edu)
- **Current TA office hours (subject to change, locations TBA):**
 - Monday 11:30 am – 2:30 pm William (2 hours)
 - Tuesday 3:30 pm – 5:30 p.m Ross (2 hours)
 - Wednesday 1:00 p.m – 3:00 p.m Vikram (2 hours)
 - Thursday 3:30 p.m – 5:30 pm Ross (2 hours)
 - Friday 1:00pm – 3:00 p.m. William (2 hours)

If you can't make any of these office hours, please send us email to make appointment to meet!

RECALL THE
TURING TEST

IN A 1950 PAPER, TURING ASKED “CAN MACHINES THINK?”

- We would eventually be able to create a computer that thinks
- BUT, how will we know when we have succeeded?
- Prediction: **by the year 2000, machines would be capable of fooling 30% of human judges after five minutes of questioning.**

Turing Test: A method to empirically determine whether a computer has achieved intelligence

TURING TEST

- A HUMAN interrogator sits in one room and uses a computer to communicate with respondents A and B in another room
 - One of A and B is a human, the other is a computer
 - The interrogator does not know which is a computer
- After conversing with both respondents, the interrogator must decide which is the computer
 - This experiment is repeated with numerous human interrogators
 - If the computer can fool enough humans interrogators, it is considered intelligent

IMITATING THE
IMITATION GAME
THE TOURNAMENT



RULES OF PLAY (AMONG 2+ TEAMS)

- Each has 3 members
- For each team, each member has a laptop with internet access
- Each team will pick randomly and independent to act either as a HUMAN or a CHATBOT
- For 5 minutes, the class will ask each team arbitrary questions
- Each time a different team member will answer the question
- As HUMAN, a team will come up with an answer by typing on their computer
- As CHATBOT, a team will type the question into each of their 3 favorite CHATBOT programs and obtain an answer among them
- All Questions and answers will show up on the **Canvas** chat channel
- Students will vote at the end of the process whether the team is actually HUMAN or CHATBOT
- The team that manage to fool more classmates wins!

The trick: as HUMAN, can you answer in a reasonable way to trick the audience?

As CHATBOT, can you choose the BEST answer that imitate HUMAN?



THIS IS A BONUS CLASS ACTIVITY

- Each team who is participating obtains 2 bonus points
- The team that wins the tournament will obtain additional 3 bonus points

How do you prepare for the tournament?

- Choose among some of the more well-known Chatbots that have won the **Loebner Prize**: Mitsuku, Rose, etc. See https://en.wikipedia.org/wiki/Loebner_Prize
- Or choose among all possible Chatbots found under: <https://www.chatbots.org/>
- You need to design a character for your team: an age group (13+ year old), a gender, a profession, and a background story about your character
- As many teams can sign up, it is encouraged to do so.

A QUICK DEMO


RECALL THE TURING TEST

A method to empirically determine whether a computer
has achieved intelligence

OBJECTIONS TO THE TURING TEST

OBJECTIONS TO THE TURING TEST

What must a computer be good at to pass the Turing test?



SEARLE'S CHINESE ROOM ARGUMENT BY JOHN SEARLE

- Searle can not understand any Chinese
- He is in a room with input and output windows, and a list of **rules** about manipulating Chinese characters
- The characters are all **squiggles** to him
- Chinese scripts and questions come in from the input window
- Following the rules, he manipulates the characters and produces a reply, which he pushes through the output window

THE CHINESE ROOM ARGUMENT



SEARLE'S CHINESE ROOM ARGUMENT

- The Chinese Room argument is one kind of objection to functionalism, specifically to the Turing Test
- Also an attack on **strong AI**
- Searle makes distinction between strong AI and weak AI
 - Strong AI “the appropriately programmed computer really is a mind, in the sense that computers, given the right programs can be literally said to **understand**”
 - **Weak AI**: computers can **simulate** thinking and help us to learn about how humans think
- Searle objects only to strong AI

THE MAIN ARGUMENT

1. The Chinese answers that Searle produces are very good: nobody can tell that he is not a native Chinese speaker
2. Searle's Chinese Room passes the Turing Test
 - Searle has only conducted **symbol manipulation**, which no understanding, yet he passes the Turing Test
 - No one would claim that a room and rule book are intelligent
3. Searle concludes that symbol manipulation alone can never produce understanding
 - Computer programming is only symbol manipulation
 - Computer programming can never produce understanding
4. Strong AI is false and functionalism is wrong

THE RESPONSE TO THE OBJECTION: THE SYSTEM REPLY

- Searle is part of a larger system
- Searle does not understand Chinese, but the whole system (Searle + room + rules) does understand Chinese
- The knowledge of Chinese is in the rules contained in the room

SEARLE'S RESPONSE TO THE SYSTEMS REPLY

1. It is absurd to say that the room and the rules can provide understanding
2. What if I memorized all the rules and internalized the whole system. Then there would just be me and I still would not understand Chinese

COUNTER-RESPONSE TO SEARLE'S RESPONSE

If Searle could internalize the rules, part of his brain would understand Chinese. His brain would house two personalities: English-speaking Searle and Chinese-speaking system



THE COMPLEXITY REPLY (A TYPE OF SYSTEM REPLY)

1. Searle's thought experiment is **deceptive**. A room, a man with no understanding of Chinese and a few slips of paper cannot pass for a native Chinese speaker.
2. It would be **incredibly difficult** to simulate a Chinese speaker's conversation.
3. You need to program in **knowledge of the world**, an individual **personality** with simulated **life history** to draw on, and the ability to be **creative** and **flexible** in conversation.
4. Basically you need to be able to simulate the complexity of an adult human brain, which is composed of billions of neurons and trillions of connections between neurons.

- The brain is just a bunch of neurons - do we say it is absurd that

COMPLEXITY CHANGES EVERYTHING

- The brain is just a bunch of neurons – do we say it is absurd that we are intelligent?
- Our intuitions about what a complex system can do are highly unreliable.
- Tiny ants with tiny brains can produce complex ant colonies, called an **emergent** property of the system
- Computers that at the most basic level are just binary switches that flip from 1 to 0 can play chess and beat the world's best human player.
- Maybe **symbol manipulation of sufficient complexity** can produce understanding as an emergent property



Credit: wikipedia

FINAL POINT: MAYBE BOTH ARE RIGHT?

If the computer is organized as a really long list of questions with canned answers, the computer does not have mental states such as belief or desire.

But if the computer is organized like a human mind, with concepts and complex organization, perhaps intelligence can result.

COMPUTATIONAL
THINKING BEHIND THE
CHATBOTS

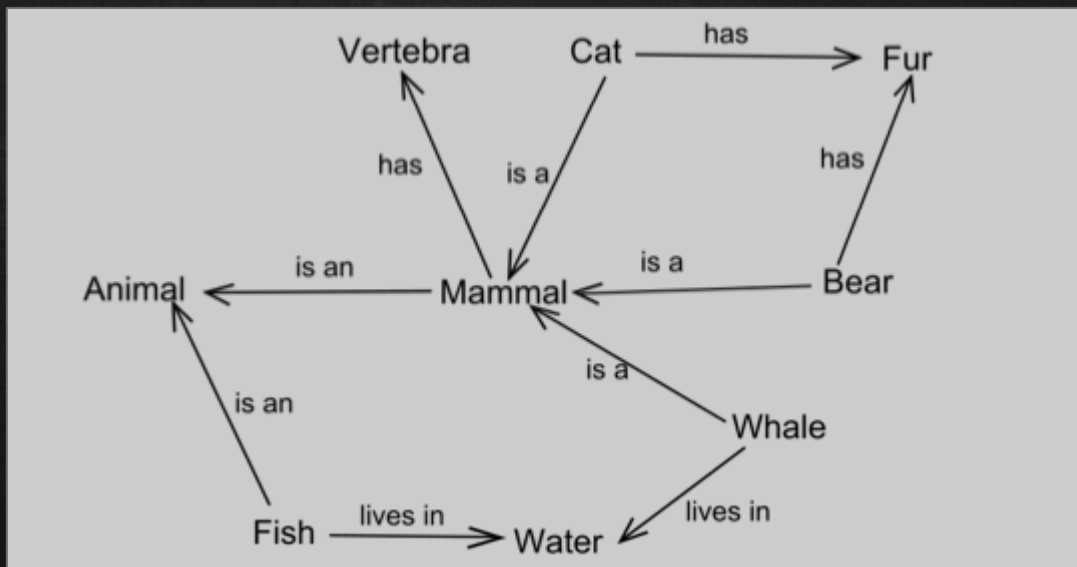
KNOWLEDGE REPRESENTATION IN AI

- Capture facts and relationship: search and process information efficiently
- **Semantic network**: a knowledge representation that focuses on relationships among objects.
- **Search tree**: a knowledge representation that captures alternatives (such as in game playing).

SEMANTIC NETWORKS

A directed graph where nodes represent objects and edges represent relationships among objects.

- What relationships should be represented?
- What questions can be answered easily? With difficulty?



EXPERT SYSTEM

- A software system based on the knowledge of human experts.
 - A user consults an expert system when confronted with a particular problem, and the system advises the user how to proceed.
 - The system uses a set of rules to guide its processing
- Information is gathered from the user by the asking of questions.
 - Example (human expert): a doctor, how does he/she diagnose?

EXAMPLE A

An expert system that embodies the knowledge of a gardener for when to apply treatments.

Rules:

if (DATE - LASTAPPLIED < 30) then NONE

if (SEASON = winter) then not BUGS

if (BARE) then TURF

if (SPARSE and not WEEDS) then TURF

if (WEEDS and not SPARSE) then WEED

if (WEEDS and SPARSE) then WEEDFEED

Diagnosis:

- Does the lawn have large, bare areas? No
- Is the lawn generally thin? Yes
- Does the lawn contain significant weeds? Yes
- Has it been a month since the last treatment?
Yes

Result:

You should apply a weed-killing and fertilizer combination treatment.

NATURAL LANGUAGE PROCESSING (NLP)

- 3 basic types of processing during human computer voice interaction:
 - **Voice recognition**: recognizing human words
 - **Natural language comprehension (NLC)**: interpreting human communication
 - **Voice synthesis**: recreating human speech
- Suppose that a computer can recognize the words spoken by a human, what is the meaning of those words?
 - NLC is the most challenging part of NLP
 - Natural language is inherently ambiguous. In what ways?

NATURAL LANGUAGE COMPREHENSION (NLC)

Lexical ambiguity: A single word can have multiple definitions and represent multiple parts of speech.

Time flies like an arrow

Syntactic ambiguity: Phrases can be put together in multiple ways

I saw the Grand Canyon flying to New York

Referential ambiguity: The use of pronouns is non-specific

The brick fell on the computer but it is not broken



https://www.youtube.com/watch?v=_429UizN1JM&feature=relmfu

TIME FLIES LIKE AN ARROW

- "time" moves quickly in a straight path
- measure the speed of flying insects like you would measure that of an arrow (thus interpreted as an imperative) – i.e. (You should) time flies as you would (time) an arrow
- measure the speed of flying insects that are like arrows – i.e. (You should) time (those) flies (that are) like an arrow
- all of a type of flying insect, "time-flies," collectively enjoy a single arrow
- the magazine *Time*, when thrown, moves in a similar manner to that of an arrow.

AI APPLICATIONS

- ❑ Medical assistants
- ❑ Shopping advisors
- ❑ Stock market trading
- ❑ **Machine learning** is the name for **new AI** based on statistical properties
- ❑ Robotics
- ❑ And much more!

AI IN MEDICINE

- Data mining, treatment / drug recommendations, identify high risk patients (e.g. liver failure, etc.)
- Personal medicine
- **Assist, not replace** doctor to reduce mortality rate
- **Fast and accurate diagnostics** of cancer using artificial neural networks
- Therapeutic robots, e.g. help Alzheimer's patients
- Eliminating some human errors due to tiredness
- Surgical robotics, minimally invasive surgery, CyberKnife (radiation)
- Virtual doctor, online medical data submission, reduce cost on office visit

SEARCH ENGINE AND WEB SERVICES

- Tailored search results
- Virtual personal assistant: Siri and smart home (voice-activated AI)
- Revolutionize how we interact with our environment
- It is kind of already here: Siri in iOS 6...
- Human Computer Interaction (HCI)

CHALLENGES AHEAD

- Privacy
- Information bank v.s. diagnostic tool (subject to slow FDA approval)
- Diagnostic with probability

THINKING OUTSIDE
THE BOX

WHAT ARE THE POTENTIAL BUSINESS OPPORTUNITIES FOR AI-TYPE COMPUTATIONAL THINKING?

- Chatbot in online customer services (have you ever thought of the person you chat with might not be human?): specific targeted questions and users.
- CHECK OUT Inteliwise, Virtuoz (Paypal, H&R block), Codebaby, MyCyberTwin
- Mitsuku deployed to Kik messenger (80% Kik users are under 22), dating advice, anyone?
- Advertiser's new frontier: talking to a Chatbot “Chatvertising” (WSJ article)
- Chat with a Brand
- WeChat: 400 million users and growing -- large user base!

WHAT'S NEXT?
THINK LIKE A
COMPUTER SCIENTIST!

PROBLEM SOLVING AND
LIGHT PROGRAMMING
FOR PEDESTRIANS



THANKS!

Any questions?

You can find me at
beiwang@sci.utah.edu

<http://www.sci.utah.edu/~beiwang/teaching/cs1060.html>

CREDITS

Special thanks to all the people who made and released these awesome resources for free:

- Presentation template by [SlidesCarnival](#)
- Photographs by [Unsplash](#)