

01-0: Syllabus

- Office Hours
- Course Text
- Prerequisites
- Test Dates & Testing Policies
 - Check dates now!
- Grading Policies

01-1: Topics

- Topics we will cover:
 - Uninformed (blind) search
 - Heuristic Search
 - Adversarial Search (Game Theory)
 - Probabilistic Reasoning
 - Machine Learning
 - Natural Language Processing
 - Information Retrieval
 - Logic and Inference

01-2: Non-AI Topics

- Topics we will cover:
 - Python
 - HTTP
 - HTML, XML
 - ... and a few other things along the way

01-3: AI Non-Topics

- Cool AI topics that we are not likely to cover, due to time constraints:
 - Vision, language recognition, and other signal processing tasks
 - Robotics
 - .. and a whole lot more. AI is a huge area, we will just scratch the surface

01-4: How to Succeed

- Come to class. Pay attention. Ask questions.
 - A question as vague as “I don’t get it” is perfectly acceptable.
 - If you’re confused, *at least* 4 other people are, too.
- Come by my office

- I am *very* available to students.
- Start the homework assignments early
- Read the textbook. It's actually pretty good
 - Ask questions if you don't understand the textbook!

01-5: **What is AI?**

- What is Artificial Intelligence?

01-6: **What is AI?**

- What is Artificial Intelligence?
 - Machine that can "Act Intelligently"
 - Machine that can think
- Let's break this down a bit more ...

01-7: **What is a Machine?**

- Newell & Simon: Physical Symbol System
 - Collection of symbols, and the rules to combine them into new symbols
 - Claim: *necessary* and *sufficient* means for intelligent action
 - Controversial!
- Turing-equivalent machines
 - For this class, standard computers
- Other Models for a "machine"?

01-8: **What is Thinking?**

- How can we determine if a machine is thinking / intelligent?
- For that matter, how do I know that **you** are thinking / intelligent?
 - Acting like a person
 - Thinking like a person
 - Thinking "rationally"
 - Acting "rationally"

01-9: **Acting like a person**

- Turing Test
 - Engage in an IM session with another entity
 - Try to determine if the entity on the other side is a person
- Is passing the turing test necessary and sufficient for "Intelligence"

01-10: Turing Test

- Necessary?
 - What about non-symbolic tasks?
 - Computer vision – recognizing your grandmother
 - Autonomous navigation – drive from 3rd and Market to Hollywood and Vine
- Sufficient
 - It rather depends on who is asking the questions!
 - First Loebner Prize (1991) bronze was won by a modified version of Eliza
- For a reasonable interrogator, the Turing test is considered to be the Gold Standard – really, **really** hard!

01-11: Turing Test

Sample dialog from Turing, that would show machine intelligence:

HUMAN: In the first line of your sonnet which reads "shall I compare thee to a summer's day", would not a "spring day" do as well or better?
MACHINE: It wouldn't scan
HUMAN: How about a "winter's day." That would scan all right
MACHINE: Yes, but nobody wants to be compared to a winter's day
HUMAN: Would you say that Mr. Pickwick reminded you of Christmas?
MACHINE: In a way
HUMAN: Yet Christmas is a winter's day, and I do not think that Mr. Pickwick would mind the comparison
MACHINE: I don't think you're serious. By a winter's day one means a typical winter's day, rather than a special one like Christmas

01-12: Searle's Chinese Room

- Person in a room with a large stack of instructions that describe how to manipulate indiscipherable symbols, and a large quantity of blank paper
- A paper with symbols are passed in, based on these symbols, the rules, and things written down previously, the person writes notes for later, and then passes out a paper with other symbols on it
- Person has no idea what any of the symbols mean
- From the outside, it appears that there is a conversation happening in Chinese
- Since the person doesn't know Chinese, and paper can't know anything, the room doesn't know Chinese

01-13: Thinking Like a Person

- If a system uses that same thought processes as a person, it is intelligent
- How do people think?
- Cognitive Science researchers build computational models of human problem solving
 - Great for a better understanding of the mind
 - Not so good for solving problems – very hard
- People are actually quite bad at many tasks (estimation, probability, etc) – thinking "like a human" is not always a good thing!

01-14: Thinking Rationally

- Follow sound reasoning processes that always lead to correct outcomes
- Logic & formal reasoning
- Issues:

- Formalizing common-sense knowledge (Cyc project)
- Dealing with uncertainty
- Computational Complexity

01-15: **Acting Rationally**

- System is intelligent if it acts rationally
- That is, acts so as to achieve the best possible (expected) outcome, given its knowledge and ability
- Work on smaller problems that require “intelligence”
- Avoids many philosophical questions

01-16: **A Bit of History**

- AI has traditionally been seen as a “failed discipline”
- Unfortunate history of overpromising and underdelivering

“It is not my aim to surprize or shock you – but the simpliest way I can summerize is to say that there are neow in the world machines that think, that learn and that create. Morover, their ability to do these things is going to increase rapidly until, in a visible future – the range of problems they can handle will be coextensive with the range to which the human mind has been applied.”

Can you guess the date of the quote?

01-17: **A Bit of History**

- AI has traditionally been seen as a “failed discipline”
- Unfortunate history of overpromising and underdelivering

“It is not my aim to surprize or shock you – but the simpliest way I can summerize is to say that there are neow in the world machines that think, that learn and that create. Morover, their ability to do these things is going to increase rapidly until, in a visible future – the range of problems they can handle will be coextensive with the range to which the human mind has been applied.”

[Herbert Simon, 1957]

01-18: **AI == Unsolved**

- As soon as a problem in AI is solved, it is “No longer an AI problem”
 - Voice Recognition
 - Face Recognition
 - Chess
 - Optimal Scheduling
 - Diagnostics (medical, automotive, etc)
 - Automated Translation
 - Although there is still have a long way to go ...

01-19: **Strong AI vs. Weak AI**

- “Strong” AI
 - Machines that actually think

- “Weak” AI
 - Machines that act intelligently (Functional Approach)
 - If it walks like a duck ...

01-20: **End of Philosophy**

- Now done with the philosophical questions for this course
- Concentrate on solving specific problems
- Use an Agent-Based approach
 - Nice way to unify disparate fields of AI