

CMSC 471

Introduction

Tim Finin, finin@umbc.edu

What is AI?

- Q. What is artificial intelligence?
- A. It is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable.

<http://www-formal.stanford.edu/jmc/whatisai/>

Ok, so what is intelligence?

- Q. Yes, but what is intelligence?
- A. Intelligence is the computational part of the ability to **achieve goals** in the world. Varying kinds and degrees of intelligence occur in people, many animals and some machines.

<http://www-formal.stanford.edu/jmc/whatisai/>

Big questions

- Can machines think? Can they learn from their experience?
- If so, how?
- If not, why not?
- What does this say about human beings?
- What does this say about the mind?

A little bit of AI history

Ada Lovelace



- Babbage thought his machine was just a number cruncher
- Ada Lovelace saw that numbers can represent other entities, enabling machines to reason about anything
- But, she wrote: *“The Analytical Engine has no pretensions whatever to originate anything. It can do whatever we know how to order it to perform.”*

AI prehistory and early years

- George Boole invented propositional logic (1847)
- Karel Capek coined term **robot** in play R.U.R. (1921)
- John von Neumann: **minimax** (1928)
- Norbert Wiener founded field of **cybernetics** (1940s)
- **Neural networks**, 40s & 50s, among the earliest theories of how we might reproduce intelligence
- Isaac Asimov *I, Robot* (1950) Laws of Robotics
- **Turing test**, proposed in 1950 & debated ever since
- Early work on **Chess** By Alan Turing, 1950

AI prehistory and early years

- Logic Theorist and GPS, 1950s, early symbolic AI; focus on search, learning, knowledge representation
- [Marvin Minsky](#): neural nets (1951), AI founder, blocks world, Society of Mind
- [John McCarthy](#) Lisp (1958), coined AI (1957)
- Allen Newell, Herbert Simon: GPS (1957), AI founders
- Noam Chomsky: analytical approach to language (1950s)
- Dartmouth summer conference (1956) established AI as a discipline

1956 Dartmouth AI Project



Five of the attendees of the 1956 Dartmouth Summer Research Project on AI reunited in 2006: Trenchard More, [John McCarthy](#), [Marvin Minsky](#), [Oliver Selfridge](#), and [Ray Solomonoff](#). Missing were: [Arthur Samuel](#), [Herbert Simon](#), [Allen Newell](#), [Nathaniel Rochester](#) and [Claude Shannon](#).

1956 Dartmouth AI Project

“We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.”

<http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>

1956 Dartmouth AI Project

“We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.”

<http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>

Recent AI History

- AI has had it's ups and downs
 - 50-60 up, 70s down, 80s up, 90s down, 00s up, 10s up, ...
 - Like the stock market, the overall trend is up
- Hot topics today?
 - Neural networks again: deep learning
 - Machine learning, datamining,
 - Exploiting big data
 - Autonomous vehicles, robotics
 - Text mining, natural language technology, speech
 - Computer vision

Why AI?

Engineering: get machines to do useful things

e.g., understand spoken natural language, recognize individual people in visual scenes, find the best travel plan for your vacation, etc.

Cognitive Science: model and understand how natural minds and mental phenomena work

e.g., visual perception, memory, learning, language, decision making, etc.

Philosophy: explore basic, interesting and important philosophical questions

e.g., the mind body problem, what is consciousness, do we have free will, etc.

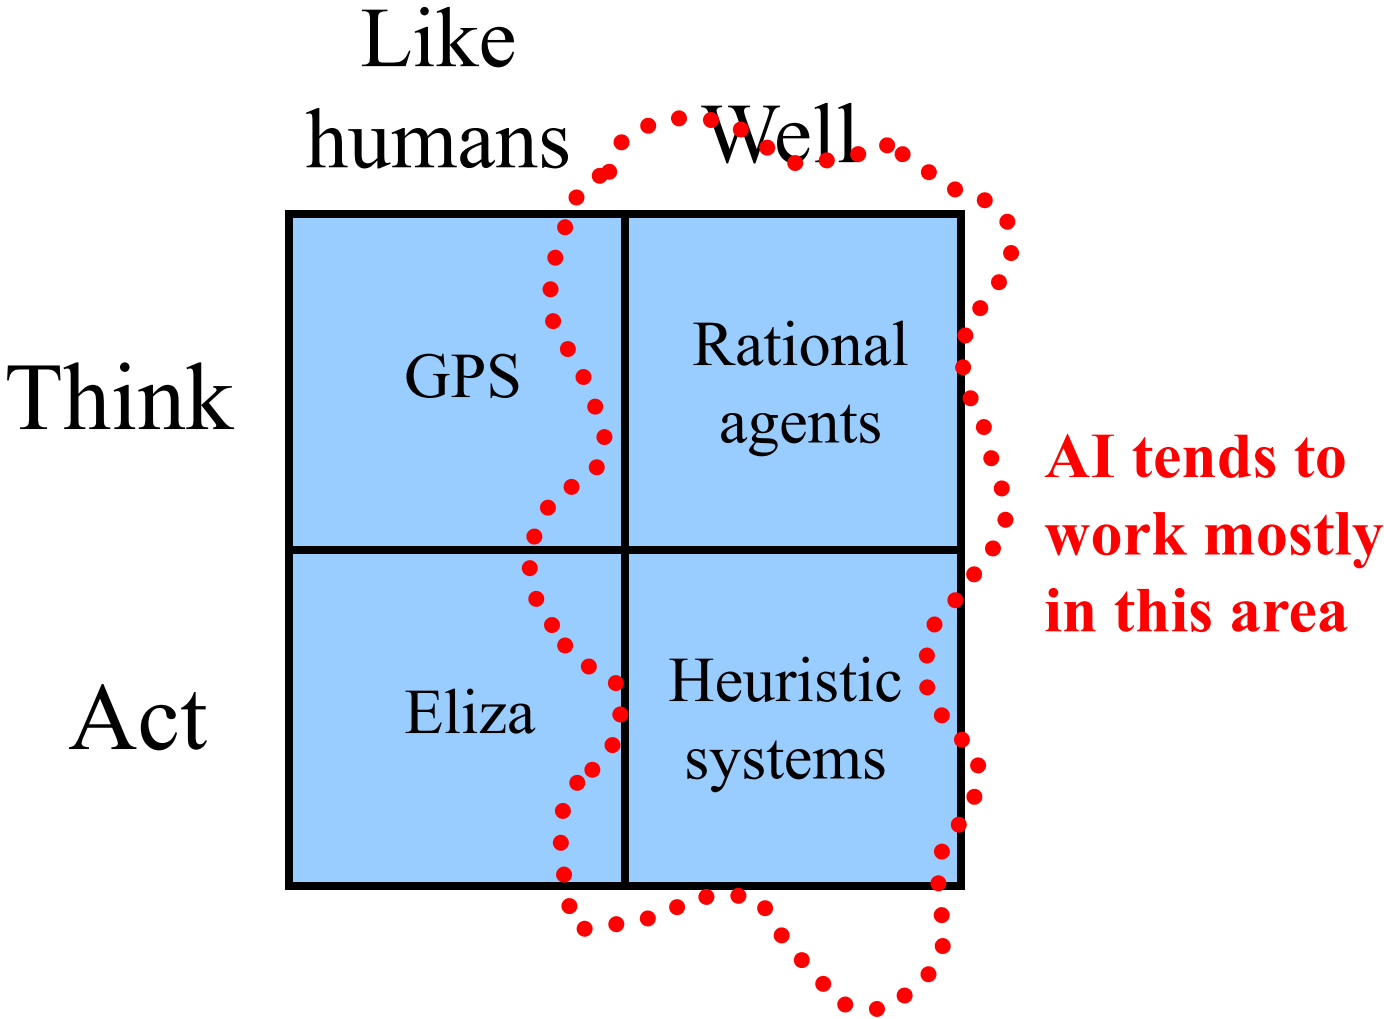
Possible AI approaches

	Like humans	Well
Think		
Act		

Possible approaches


	Like humans	Well
Think	GPS	Rational agents
Act	Eliza	Heuristic systems

Possible approaches



Think well

	Like humans	Well
Think	GPS	Rational agents
Act	Eliza	Heuristic systems




- Develop formal models of knowledge representation, reasoning, learning, memory, problem solving, that can be rendered in algorithms
- Often an emphasis on a systems that are provably correct, and guarantee finding an optimal solution

Act well

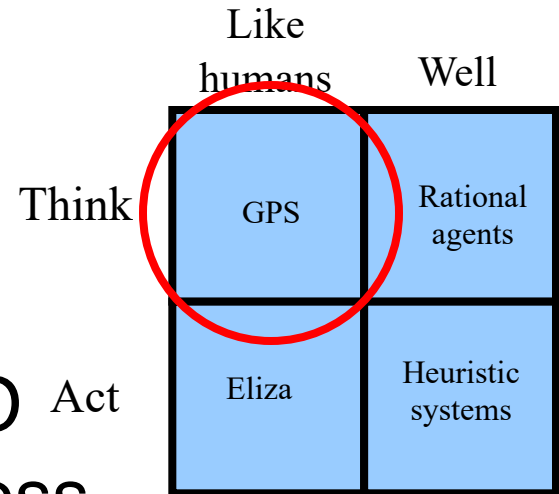
- For a given set of inputs, generate output that's not necessarily correct but gets job done
- A **heuristic** (heuristic rule, heuristic method) is a rule of thumb, strategy, trick or simplification which drastically limits search for solutions in large problem spaces
- Heuristics don't guarantee optimal solutions or even any solution at all: **all that can be said for a useful heuristic is that it offers solutions which are good enough most of the time**
 - Feigenbaum and Feldman, 1963, p. 6

	Like humans	Well
Think	GPS	Rational agents
Act	Eliza	Heuristic systems



Think like humans

- Cognitive science approach
- Focus not just on behavior & I/O but also look at reasoning process
- Computational model should reflect “how” results were obtained
- Provides new language for expressing cognitive theories & new mechanisms for evaluating them
- GPS (**General Problem Solver**): Goal not just to produce humanlike behavior, but to produce a sequence of steps of reasoning process that was similar to those followed by a person




Act like humans

- Behaviorist approach
- Not interested in how you get results, just similarity to what human results are
- Exemplified by the Turing Test (Alan Turing, 1950)
- Has applications in interactive entertainment (e.g., computer games, CGI), virtual worlds and in modeling human intentions

	Like humans	Well
Think	GPS	Rational agents
Act	Eliza	Heuristic systems

Turing Test




- Three rooms: person, computer, interrogator
- Interrogator communicates with others via a ‘teleprinter’
- Interrogator tries to distinguish person and machine
- Machine tries to fool the interrogator into believing that it is the person
- If machine succeeds, we conclude it can think

Eliza

- ELIZA: program simulating a psychotherapist interacting with a patient
 - Joe Weizenbaum (MIT), ~1966, FORTRAN
- First script was DOCTOR
 - Simple collection of syntactic patterns not unlike regular expressions
 - Patterns had associated replies that might include bits of the input after simple transforms (my → your)
- Weizenbaum shocked at reactions:
 - Psychiatrists thought it had potential
 - People unequivocally anthropomorphized
 - Many thought it solved the NL problem

It lives!

Amazingly, a 1966 version was recently found and restored, see <http://elizagen.org/>



Eliza

This is a transcript of a conversation with an Eliza-like program found in Emacs

Try typing

*Meta-x
doctor*

in your favorite Emacs

- I am the psychotherapist. Please, describe your problems. Each time you are finished talking, type RET twice.
- *i'm having a hard time with my students*
- Why do you say that?
- well, they keep asking me for money
- Maybe your plans have something to do with this.
- *what plans?*
- Is it because of your life that you say what plans?
- *No, it's because I didn't know what you were talking about*
- Is it because you didn't know what I was talking about that *you came to me?*
- *no, it's because i wanted to see if you could pass the turing test*
- Does it bother you that you wanted to see if I could really pass the turing test?
-
- How do you reconcile problems at school?
- *This is a waste of time. Good bye*
- My secretary will send you a bill.

The Loebner contest




- Modern version of Turing Test, held annually with a \$100K cash prize
- Hugh Loebner was once director of UMBC Academic Computing Services (née UCS)
- Restricted topic (removed in '95) & limited time
- Participants: humans, computers and judges
- Scoring rank from least to most human
 - Highest median rank wins \$2000
 - If better than a human, win \$100K (Not yet)

Results: [2016](#), [2017](#), [2018](#), [2019](#)

Is it relevant to AI?

- Most AI researchers believe that the Loebner contest is not relevant to AI
- However, techniques used by some are similar to those that are driving many useful chat bot applications
 - These are supported, however, by sophisticated speech to text and text to speech systems
 - And varying degrees of language understanding and machine learning



What's easy and what's hard?

- Easy: many high-level tasks usually associated with “intelligence” in people
 - e.g., symbolic integration, proving theorems, playing chess, medical diagnosis
- Hard: tasks many animals can do
 - walking around without running into things
 - catching prey and avoiding predators
 - Interpreting sensory info. (e.g., visual, aural, ...)
 - modeling internal states of other from behavior
 - working as a team (e.g., with pack animals)
- Is there a fundamental difference between these?

What can AI systems do?

- **Computer vision:** face recognition from a large set
- **Robotics:** autonomous (mostly) automobile
- **Natural language processing:** useful machine translation and simple fact extraction
- **Expert systems:** medical diagnosis in narrow domains
- **Spoken language systems:** e.g., Google Now, Siri, Cortana
- **Planning and scheduling:** Hubble Telescope experiments
- **Learning:** text categorization into ~1000 topics
- **User modeling:** Bayesian reasoning in Windows help (the infamous paper clip...)
- **Games:** Grand Master level in chess (world champion), checkers, etc.

What can't AI systems do yet?

- Understand natural language robustly (e.g., read and understand articles in a newspaper)
- Surf the web and find interesting knowledge
- Interpret an arbitrary visual scene
- Learn a natural language
- ~~Play Go well~~
- Construct plans in dynamic real-time domains
- Refocus attention in complex environments
- Perform life-long learning

Exhibit true autonomy and intelligence!

T.T.T

Put up in a place
where it's easy to see
the cryptic admonishment


T. T. T.

When you feel how depressingly
slowly you climb,
it's well to remember that

Things Take Time.

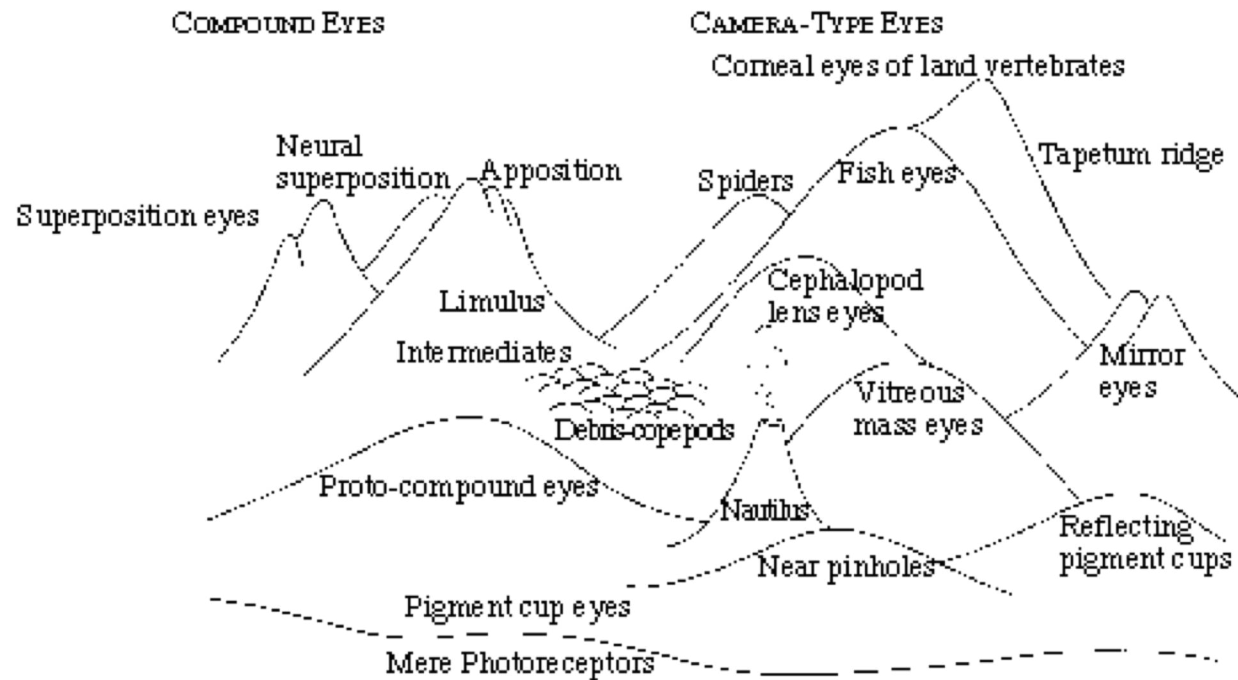
-- [Piet Hein](#)

T.T.T: things take time



- Prior to the 1890's, papers were held together with straight pens.
- The development of “spring steel” allowed the invention of the paper clip in 1899.
- It took about **25 years (!)** for the evolution of the modern “gem paperclip”, considered to be optimal for general use.

Climbing Mount Improbable



“The sheer height of the peak doesn't matter, so long as you don't try to scale it in a single bound. Locate the mildly sloping path and, if you have unlimited time, the ascent is only as formidable as the next step”

-- Richard Dawkins, *Climbing Mount Improbable*, Penguin Books, 1996.