
Artificial Intelligence

A Very Brief Overview of a Big Field

Notes for CSC 100 - The Beauty and Joy of Computing
The University of North Carolina at Greensboro

Reminders

Blown to Bits

Chapter 5 or 6: Contribute to discussion by next Monday

Project


Get solid, workable code ready by Friday!
Demos and final presentations in lab this Friday

What is Artificial Intelligence (AI)?

Definition by John McCarthy:

Getting a computer to do things which, when done by people, are said to involve intelligence.

John McCarthy (September 4, 1927 – October 24, 2011) was an American computer scientist and cognitive scientist. He invented the term "artificial intelligence" (AI), developed the LISP programming language family, significantly influenced the design of the ALGOL programming language, popularized timesharing, and was very influential in the early development of AI.



McCarthy received many accolades and honors, such as the Turing Award for his contributions to the topic of AI, the United States National Medal of Science, and the Kyoto Prize.

Some Aspects of AI

Natural Language Processing (NLP)

Communicating the way people do

A form of Human-Computer Interaction (HCI)

Machine Learning

Performing better with experience: supervised, unsupervised, reinforcement

Sensory Input

Hearing (processing spoken English), Seeing (image recognition, writing recognition, ...)

Planning

Deciding on a future plan of action

Robotics

Includes HCI, NLP, machine vision, planning, ...

Computer creativity

Computer generated music, art, ...

Progress in AI

A lot of hype in the early days (late 50's and throughout the 60's). For example:

In 1957, Allen Newell said: "Within ten years a digital computer will be the world's chess champion" (wasn't until 1997 that a computer beat the human world champion in a match).

Hasn't progressed like people thought, but in past 10 years a lot of progress:

- Handwriting recognition on envelopes (zip codes), check amounts (ATM deposits), ...
- Face recognition in images (auto face tagging in Facebook or Picasa)
- Software personal assistants: Siri, Amazon Alexa, Google Now, ...
- "Intelligent" games (checkers is solved, chess is past human players, and Jeopardy...)

Can we test for Artificial Intelligence?

The Imitation Game....

Most widespread notion of testing for intelligence is the Turing Test.

Avoids answering what it means for a "machine to think."

Basic idea: An "interrogator" is communicating with either a person or a computer - can they tell the difference?

Some basic points:

- If conversing with a computer is *indistinguishable* from conversing with a human, is the computer intelligent?
- Oddly enough, computer might have to dumb-down some responses: "What is 1591356 times 672435?"
- Is this different from a computer "understanding" a conversation?



Note: Annual "Loebner Prize" competition built on the Turing Test.

"The Singularity"

Idea promoted by futurist Ray Kurzweil:


- When technological "superintelligence" exceeds intelligence of people, and can build smarter and smarter machines.
- "Intelligence Explosion" - "superintelligences design successive generations of increasingly powerful minds, might occur very quickly and might not stop until the agent's cognitive abilities greatly surpass that of any human"

Note: Ray Kurzweil is the current person popularly associated with this idea, but the idea itself has roots in the writing of I.J. Good and Vernor Vinge

A famous view from science fiction

Isaac Asimov's short story "Runaround" (in the "I, Robot" compilation) set out "Three Laws of Robotics":

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2. A robot must obey the orders given to it by human beings, except where such orders would conflict with the First Law.
3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.



Became a very common meme in science fiction and in tech culture.

Flipping AI

Taking advantage of lack of progress in AI!

Problem: We would like to ban automated account registrations that sign up for free accounts to send spam, etc. We want humans to succeed, but computers to fail - this is a Turing Test!

Solution: CAPTCHAs - Completely Automated Public Turing test to tell Computers and Humans Apart


There are some things that humans are still much better at than computers - such as being able to read distorted images of text, so...



Why do I have to do this?
This trap helps protect your account from spam and other abuse, and helps to ensure that your email will be delivered.

What is next in AI?

Andrew Ng on "The Future of Robotics and Artificial Intelligence"



http://www.youtube.com/watch?v=AY4ajbu_G3k (link on class web page)
