

COMPUTER PROGRAMMING IN JAVA

COLUMBIA UNIVERSITY HIGH SCHOOL SCIENCE HONORS PROGRAM

Intro to AI: Chatbots and Turing Tests

Lecture

2007 Apr 12 Sat

Intro

There is currently no generally accepted formal definition of “artificial intelligence”. Because of this fundamental confusion, AI means different things in different contexts.

AI as a modern research field can be said to have started after World War II, building upon the work done in the decades preceding by Alan Turing, Alonzo Church, John von Neumann, and many others. [<http://www-formal.stanford.edu/jmc/whatisai/node1.html>]

AI spans many topics, from general deduction, reasoning, and problem solving to knowledge representation to decision making to natural language processing to machine learning; some of these have become research fields in their own right, while others have remained under the general umbrella of AI.

For many years one of the main directions of AI was on teaching computers how to play games (chess, checkers, go, Othello, etc). It was believed that much of the cognitive abilities we colloquially call “intelligence” is visible in the process of playing these games, and games form a nice, natural, easy-to-understand and easy-to-formulize environment to experiment with different forms of AI.

Eventually, AI programs were developed that played at or above the level of the best humans in checkers, chess, Othello, and many others. However, it is important to realize *how* these victories were achieved: the greatest strengths of the computing machine (vast amounts of memory and the ability to perform millions to billions of simple calculates per second) were leveraged to the limit. That is, it is difficult to argue that there was any *real* intelligence being exhibited by these AI programs.

The Turing Test

Alan Turing, one of the most important people in computer science, proposed a test that could be applied to a computer program to measure its level of *real* intelligence. This test is known as the “Turing test”.

The test is simple: there is an interrogator and a subject, who cannot see each other, but has a form of remote communication (say, a computer terminal). The interrogator can ask a series of questions, and based on the responses she must decide whether or not the subject is a human or a computer.

As intuitive as this might seem, there are also very many obvious flaws. For a idea of just some of the traditional objections (and defenses) of the Turing test, search for “strong AI versus weak AI” and John Searle’s “Chinese room argument”.

Chatbots

An early attempt at building a machine that could pass the Turing test was the program ELIZA, written by Joseph Weizenbaum [<http://en.wikipedia.org/wiki/ELIZA>] in 1966. It was able to fool some users even though it was relatively simple: its algorithm boiled down to taking the interrogator’s question and rephrasing it into a new question. For example:

ELIZA: Hello, I am ELIZA. How can I help you?

H: Do you have ice cream?

ELIZA: Oh, I have ice cream.

H: Can I have some?

ELIZA: Do you want to be able to have some?

[<http://www-ai.ijs.si/eliza/eliza.html>]

A more recent example is ALICE (Artificial Linguistic Internet Computer Entity), willing prizes as recently as 2004.

[http://en.wikipedia.org/wiki/Artificial_Linguistic_Internet_Computer_Entity]