



**COMP 200: Elements of Computer Science**  
**Fall 2004**

**Lecture 26: November 1, 2004**

*Artificial Intelligence & the Turing Test*

**On the Board**

Homework 5 due today

2<sup>nd</sup> test: available Friday

**Back to the Last Lecture**

Genetic algorithms were developed as part of a field called *Artificial Intelligence*. Does a GA behave intelligently? Or is it just randomization carried to an extreme level?

- What is intelligence? When is it artificial?
- What properties constitute intelligence or intelligent behavior?
  - Creativity? Judgement? An instinct for self-preservation? Tool use? (viruses or machine-building GAs) Language skills?
- Can an algorithm behave intelligently?
- What would it take to convince you that something is intelligent?
  - Dog? Cat? Tropical fish? Lobster (holding claws)? Antilock brakes?
  - Ability to parallel park a car? Ability to play ping pong? Ability to play soccer (in real time)?
  - Ability to identify sedans driving down Kirby? Ability to recognize a plane (by model) in flight? Ability to scan a complex scene and interpret it?
  - Ability to recognize human faces — as in, I've seen this person before from another angle? Ability to read emotional states from a person's voice or facial expression?
  - Ability to play with blocks — stack three blocks into a step pyramid?

## The Turing Test

Alan Turing, considered by many as one of the founding lights of modern Computer Science, proposed the following test for “intelligence”. Place the tester in a room with two computer terminals. One is connected to a program and the other is connected to a human being. The tester is given a limited amount of time, say thirty minutes, to converse with both terminals. At the end, the tester must declare which is the program and which is the human.

Variations limit the context in which the conversation can take place, the amount of time that the tester has, and so on. A program that can consistently fool testers is declared intelligent. (A single trial is insufficient because the tester could randomly pick the correct assignment of human and program to the two terminals.)

- Is the Turing test a reasonable assessment of intelligence? (or is it biased toward language processing skills)
- What questions would you ask?
  - What is  $3^{15}$ ? (Answer: 14,348,907; the program would need to hesitate before answering.)
  - What did you think about Camacho’s departure from Rice?
  - How about them BoSox?
  - Did Gerard Manley Hopkins really intend us to understand his conception of *inscape* through “The Windhover” or is that just an invention of the critics who have tried to understand his work?

Examples of systems that were (potentially) intelligent and how they behave.

- SHRDLU — used a camera to “see” a scene in blocks-world and to follow simple directions. Could describe a scene in near human terms. Eventually, could stack up blocks in response to directions. SHRDLU required a huge amount of contextual knowledge, including solid geometry, the laws of gravity, notions of structure, the ability to infer hidden lines (the back of a block), and a basic understanding of the directions that it was given.

At the time, SHRDLU was a huge accomplishment. It relied on a carefully designed and implemented knowledge representation, along with a significant amount of pre-programmed context.

- ELIZA — Eliza was an interactive program that simulated a Rogerian psychologist. Eliza “listens” to what you type and responds with terse questions that (sometimes) appear to emulate a psychologist’s queries.

For example, see <http://www.manifestation.com/neurotoys/eliza.php3>; downloadable versions are available at <http://www.spaceports.com/~sjlaven/eliza.htm> and at <http://ecceliza.cjb.net>. (The latter claims to be the best...)

I have seen better versions of Eliza. The original was written by Professor Joseph Weizenbaum and named after Eliza Doolittle (in *My Fair Lady*).

## Unfair Standards

In some sense, AI has suffered from rising expectations. Over the last forty years, challenges have been laid down in the form: “If a program can do  $x$ , then that program will exhibit intelligent behavior.” When a clever implementor showed that  $x$  was capable of algorithmic solution in a tractable way, the community declared that  $x$  did not require intelligence. Examples include

- Game playing: Tic-tac-toe is tractable (Game tree has  $9!$  Moves). Checkers is harder. Chess was considered to require intelligence. Over the last 20 years, both checkers and chess have fallen to programs. Now, skeptics say that the games simply require a lot of horsepower and some good heuristics. (What if that is how Karpov plays?)
- Conversation: Eliza fooled some people back in the 1960s. Once you take her outside of her area, however, she is revealed as a clever pattern matching program. Does adding context, as in “What do you think about Proposition 2’s potential impact on the pension-fund crisis in Houston?”, change the task or simply increase the knowledge base required? How many of you can answer that question and sound convincingly intelligent.

Another argument against some of the techniques of AI is that they simply rely on probability. GAs role the dice repeatedly and keep the winners. Is that how we work? Do we know? If we use probability in QuickSort, are we cheating? (pick a pivot element at random). If we throw darts to compute PI, is that an algorithm?

Is a virus intelligent? It has motivation. It reproduces. Some use mutation. Some have self-preservation tactics. Hmmm.