

Objectives:

1. To look at different views as to the meanings of “intelligence” and “artificial intelligence”
2. To distinguish between weak (applied) AI and strong AI (Artificial General Intelligence)

Materials:

1. Projectable of Definitions of AI grid from Russell and Norvig, 2015, p. 2
2. “Tasks Requiring Intelligence” Handout
3. Syllabus
4. Study Guide for Turing Article
5. Copy of this lecture on course web site - both to show how to access and to demonstrate Google translation
6. Projectable of Psalm 8

I. Preliminaries

Have students introduce themselves and say why they're interested in this course.

II. What is AI?

A. At the start of a course like this, it's fairly common for the professor to begin by talking about “what this course is all about”. But I'd like to start differently by asking you what you think this course is all about.

1. Writing Exercise
2. Share Responses

B. One reason I began the lecture this way is that there is a lot of divergence of opinion - even among those working in the field - as to what the term “artificial intelligence” means.

1. For example, the author of one widely used AI textbook speaks “the paradoxical notion of a field of study whose major goals include its own definition” (Luger 2009 p. 2)

2. This paradox should not be surprising, given that even the terms “intelligence” or “intelligent” themselves are not ones that we use in a consistent way.

a) Consider the following statements

(1)“I wonder if there is intelligent life on other planets”.

(2)“Suzy is such an intelligent person”.

(3)“I think German shepherds are the most intelligent kind of dog.”

b) How is the word “intelligent” being used in each of these statements?

ASK

(1)When we wonder if there is “intelligent” life on other planets, we are using the term in an absolute sense. Among earthly creatures, intelligence, in this sense, is an attribute possessed only by humans. (Though, as Christians, we recognize that there are non-earthly intelligent beings, including - at least - God and the angels.)

(2)On the other hand, we use “intelligent” in a relative sense, recognizing degrees of intelligence among humans and also some degree of intelligence in sub-human creatures.

(Though if a space probe discovered a German Shepherd on Mars we probably would not regard that as the discovery of intelligent life.)

3. Related to this issue is a distinction some draw between "intelligence" and "consciousness".
 - a) Is it possible to think without being having consciousness (a "self")? (Evidently, we ourselves do this - sometimes we do things that require intelligence without being conscious that we are doing them.)
 - b) Does consciousness emerge from thought, or is it something distinct?
 - c) Perhaps this distinction is related to two ways in which we use the word "intelligent" - e.g. perhaps the absolute use of "intelligent" is associated with the idea of self-consciousness. Humans are conscious of themselves; it is not clear that dogs, say, are.
4. Here are some statements that appear in the introduction to various textbooks that have been used over the years. (Many of these are from Russell and Norvig, 2015 - I've added several others and rearranged chronologically).
 - a) "[The automation of activities that we associate with human thinking, activities such as decision-making, problem-solving, learning." (Bellman, 1978)
 - b) "The exciting new effort to make computers think ... *machines with minds*, in the full and literal sense" (Haugeland, 1985)
 - c) "The study of mental faculties through the use of computational models" (Charniak and McDermott, 1985)

- d) “The art of creating machines that perform functions that require intelligence when performed by people” (Russell and Norvig attribute this to Kurzweil, 1990 - but he is actually quoting a “classic definition” that he then proceeds to take issue with)
- e) “A field of study that seeks to explain and emulate intelligent behavior in terms of computational processes” (Schalkoff, 1990)
- f) “The study of how to make computers do things at which, at the moment, people are better.” (Rich and Knight, 1991).
- g) “The study of the computations that make it possible to perceive, reason, and act.” (Winston, 1992)
- h) “Artificial intelligence is the enterprise of constructing an intelligent artifact” (Ginsberg,1993)
- i) “The branch of computer science that is concerned with the automation of intelligent behavior” (Luger and Stubblefield, 1993)
- j) “[Artificial Intelligence] is concerned with getting computers to do tasks that require human intelligence.” (Cawsey, 1998)
- k) “Artificial Intelligence, broadly ... defined is concerned with intelligent behavior in artifacts. Intelligent behavior, in turn, involves perception, reasoning, learning, communicating, and acting in complex environments.” (Nilsson, 1998)
- l) “The central scientific goal of computational intelligence is to understand the principles that make intelligent behavior possible, in natural or artificial systems.” (Poole, 1998)
- m) Russell and Norvig (from which I’ve drawn most of these statements) suggest that they can be grouped into four broad categories based on two classifications:

PROJECT: Definitions of AI grid

(1) Whether they view “intelligence” as something distinctly human or as something (rationality) that transcends humans.

(2) Whether they see the goal in terms of internal thought processes (thinking) or external behavior (acting).

C. A distinction I find especially helpful is to distinguish between what has sometimes been called “strong AI” and “weak AI”.

1. Strong AI, also called Artificial General Intelligence (AGI) is concerned with creating an artifact that deserves to be called intelligent in exactly the same way that you and I are called intelligent - an artificial person.

a) Two of the definitions of AI we looked at previously fall squarely in this category

(1) “The exciting new effort to make computers think ... *machines with minds*, in the full and literal sense”.

(2) “Artificial intelligence is the enterprise of constructing an intelligent artifact”

b) This is also the view that is behind many popular works of fiction.

Examples?

ASK

- (1) Commander Data of *Star Trek: the Next Generation*
- (2) R2D2 and C3PO of the *Star Wars* movies
- (3) Andrew Martin of *The Bicentennial Man*
- (4) David of *Artificial Intelligence*
- (5) Numerous characters in various science fiction novels

c) An often-cited article that actually predates the term “artificial intelligence” is one by Alan Turing in which he describes what has come to be called “the Turing test”. We will discuss this article next week.

DISTRIBUTE TURING ARTICLE STUDY GUIDE

A view held by many (but by no means all) who regard AI this way is that, ultimately, intelligence is computation: One definition we looked at that espouses this viewpoint rather clearly is “The central scientific goal of computational intelligence is to understand the principles that make intelligent behavior possible, in natural or artificial systems.” (Poole, 1998) (Notice that this author preferred the term “computational intelligence” rather than “artificial intelligence”).

2. Weak AI - also called applied AI - is concerned with making computers behave “intelligently” in specific areas.

a) Two of the definitions of AI we looked at previously fall squarely in this category:

(1) “The art of creating machines that perform functions that require intelligence when performed by people”

(2) “The study of how to make computers do things at which, at the moment, people are better.”

- b) Our textbook's definition is also in this category: "[Artificial Intelligence] is concerned with getting computers to do tasks that require human intelligence."
- c) A problem that can arise with this view of AI, of course, is that it involves a moving target. If a machine can do it, then it no longer requires human intelligence - and if it doesn't require human intelligence, then it is no longer artificial intelligence!

The first of the definitions I cited sidestepped this a bit by talking about "functions that require intelligence when performed by people" - but that begs the question of what human functions "require intelligence".

(1) Distribute "Tasks that require intelligence" handout and ask class to fill out.

(2) Tabulate answers

(3) Discuss in terms of degree of difficulty in automating.

(a) Several of these tasks have been successfully automated

i) What many regard as "the first AI program" - Logic Theorist (1955) successfully proved 38 of the first 52 theorems in Russell and Whitehead's *Principia Mathematica* and found novel proofs for some.

ii) In 1997 a chess-playing program - Deep Blue - beat the world chess champion, Gary Kasparov, in a six game match.

iii) We will look at a computer program that solves Sudoku puzzles later in the course (and will even learn a bit about how it does it!)

iv) Google provides a facility for translating documents from one language to another.

DEMO: Translate on-line version of this lecture - let students choose target language.

translate.google.com - enter URL:
www.cs.gordon.edu/courses/cps331/
lectures-2014/Course%20Introduction.pdf

v) IBM's Watson that won Jeopardy in 2011 and is currently being used for numerous applications,

Others task have not been automated to this day

i) Given a picture showing a familiar animal, identify what type of animal it is (dog, cat, horse ...)

ii) Recognize one's mother

iii) Tie one's shoes

(4) Do you find the relative difficulty of automating these tasks at all surprising? Why?

ASK

3. As a practical matter, of course, no one today knows how to achieve strong AI

a) Many believe, for a variety of reasons, that this is not even possible.

b) Others, however, hold the view that continuing developments in weak AI will ultimately lead to strong AI.

- c) Still others are content to work on practical problems in weak AI, with no expectation that this will lead to anything beyond increasingly useful systems.

III. This Course

A. So what is this course about?

1. Most of the time, we will be talking about strategies for imitating human intelligence, as opposed to considering theories about human intelligence. (I am a computer scientist, not a psychologist!)
2. For this most part, this course is about weak or applied AI.
 - a) We will look at some foundational issues - particularly knowledge representation and search - which are key to understanding most work in AI.
 - b) We will look at a number of practical applications of AI, and something of how they are accomplished. (Recognizing that we can only scratch the surface here)
 - c) We will look at some frontier areas where successful work is being done outside the paradigms of traditional AI.
3. We will also think a bit about Strong AI.
4. Hopefully, one thing that will emerge from our study is a deepened appreciation of what a marvelous and awesome thing our own intelligence is - as the Scripture says “we are fearfully and wonderfully made”.

B. Distribute/go over Syllabus

C. Psalm 8 - PROJECT and read