

Artificial Intelligence.
Professor Carolina Ruiz
Department of Computer Science
Worcester Polytechnic Institute

Lecture 1. (Part 1 of 3) INTRODUCTION

The objectives of Part 1 of this lecture are:

- To define Artificial Intelligence (AI).
- To describe the topics to be covered during this term.
- To go over some organizational details.

WHAT IS AI?

- AI is a relatively new field
- It started at the end of the 1940s
- Its name was coined by John McCarthy in 1956
- There are many definitions of Artificial Intelligence. Two of them are:
 - “AI as an attempt to understand intelligent entities and to build them” (Russell and Norvig, 1995)
 - “AI is the design and study of computer programs that behave intelligently” (Dean, Allen, and Aloimonos, 1995)
- What is an “intelligent entity” or what does it mean to “behave intelligently”?
 - “Intelligence is the degree of accomplishment exhibited by a system when performing a task” (Allen. AAAI97 invited lecture)

OTHER DEFINITIONS OF AI

(Adapted from Russell and Norvig's book)

Systems that think like humans	Systems that think rationally
<p>“The exciting new effort to make computers think ... machines with minds, in the full and literal sense” (Haugeland, 1985)</p> <p>“[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning ...” (Bellman, 1978)</p>	<p>“The study of mental faculties through the use of computational models” (Charniak and McDermott, 1985)</p> <p>“The study of the computations that make it possible to perceive, reason, and act” (Winston, 1992)</p>
Systems that act like humans	Systems that act rationally
<p>“The art of creating machines that perform functions that require intelligence when performed by people” (Kurzweil, 1990)</p> <p>“The study of how to make computers do things at which, at the moment, people are better” (Rich and Knight, 1991)</p>	<p>“A field of study that seeks to explain and emulate intelligent behavior in terms of computational processes” (Schalkoff, 1990)</p> <p>“The branch of computer science that is concerned with the automation of intelligent behavior” (Luger and Stubblefield, 1993)</p> <p>“AI is the design and study of computer programs that behave intelligently” (Dean, Allen, and Aloimonos, 1995)</p>

THE TURING TEST

For the “acting humanly” approach

- It was proposed by Alan Turing (1950).
- This test is an operational definition of intelligence:
It defines intelligent behavior as the ability to achieve human-level performance in all cognitive tasks, sufficient to consistently fool human interrogators.
- Test:
A computer is interrogated by a human through a tty terminal and passes the test if the interrogator cannot tell if there is a computer or a human at the other end.
- To pass the Turing test a machine will need to:
 1. represent knowledge
 2. reason automatically
 3. learn
 4. process natural language
- For the TOTAL Turing test (which includes also a video signal so that the interrogator can test the subject’s perceptual abilities) the machine will also need to:
 - 5 “see” (computer vision)
 - 6 “move” (robotics)

There has NOT been a big effort to try to pass the Turing test.

WHAT IS AI? (Cont.)

- AI is at the intersection of
 - philosophy,
 - mathematics,
 - psychology,
 - computer engineering,
 - linguistics,
 - cognitive science, and
 - computer science.
- It differs from philosophy and psychology (which are also concerned with intelligence) in which AI strives to BUILD intelligent entities as well as to understand them.
- It differs from other subareas of computer science and engineering, in its emphasis on perception, reason, and action.

WHAT IS AI? (Cont.)

- AI can be seen as an ensemble of ideas about
 - representing knowledge
 - using knowledge to solve problems
- with two goals:
 - Engineering Goal:
To solve real-world problems using AI
 - Scientific Goal:
To explain various sorts of intelligence.

TOPICS TO BE COVERED IN THIS COURSE

- Core AI:
 - Knowledge Representation Techniques:
Semantic Nets, Rules, Propositional Logic, 1st Order Logic, Probability, ...
 - Problem Solving Strategies:
Blind Search, Heuristic Search, Optimal Search, Tree and Adversarial Search (Game Playing), Constraint Satisfaction, Logical Inference, Planning, Probabilistic Reasoning, ...
- AI Applications:
 - Machine Learning,
 - Machine Vision, and
 - Natural Language Processing.

SUCCESSFUL STORIES IN AI:

- Computer Games: example

Chess: Deep Blue, developed at IBM.

- Robot Explorers: example

Mars exploration rovers designed at the Jet Propulsion Lab.

- Autonomous Vehicles: examples

Google's driverless cars. Also an earlier prototype developed at CMU.

- Expert Systems for Medical Diagnosis: examples

e.g., MYCIN (diagnoses blood infections. It performs as well as human experts and considerably better than junior doctors) developed at Stanford Univ.

IBM's Watson Computer for cancer diagnosis and treatment suggestions.

- Expert Systems for Financial Applications.

- Language Translation Systems.

- Automated Personal Assistants.

- Robots for Hazardous Conditions.