CS:4420 Artificial Intelligence

Spring 2018

Introduction

Cesare Tinelli

The University of Iowa

Copyright 2004–18, Cesare Tinelli and Stuart Russell [a]

[a] These notes were originally developed by Stuart Russell and are used with permission. They are copyrighted material and may not be used in other course settings outside of the University of Iowa in their current or modified form without the express written consent of the copyright holders.

Readings for This Lecture

- Chap. 1 of [Russell and Norvig, 2012]

What is Artificial Intelligence?

A scientific and engineering discipline devoted to:

- *understanding principles that make intelligent behavior possible in natural or artificial systems*
- *developing methods for the design and implementation of useful, intelligent artifacts [Poole, Mackworth, Goebel]*

What is intelligence then?

What is intelligence then?

- Fast thinking?
- Knowledge?
- Ability to pass as a human?
- Ability to reason logically?
- Ability to learn?
- Ability to perceive and act upon one's environment?
- Ability to play chess at grand-master's level?
- ...

Acting Humanly: The Turing test

Turing (1950) "Computing machinery and intelligence":

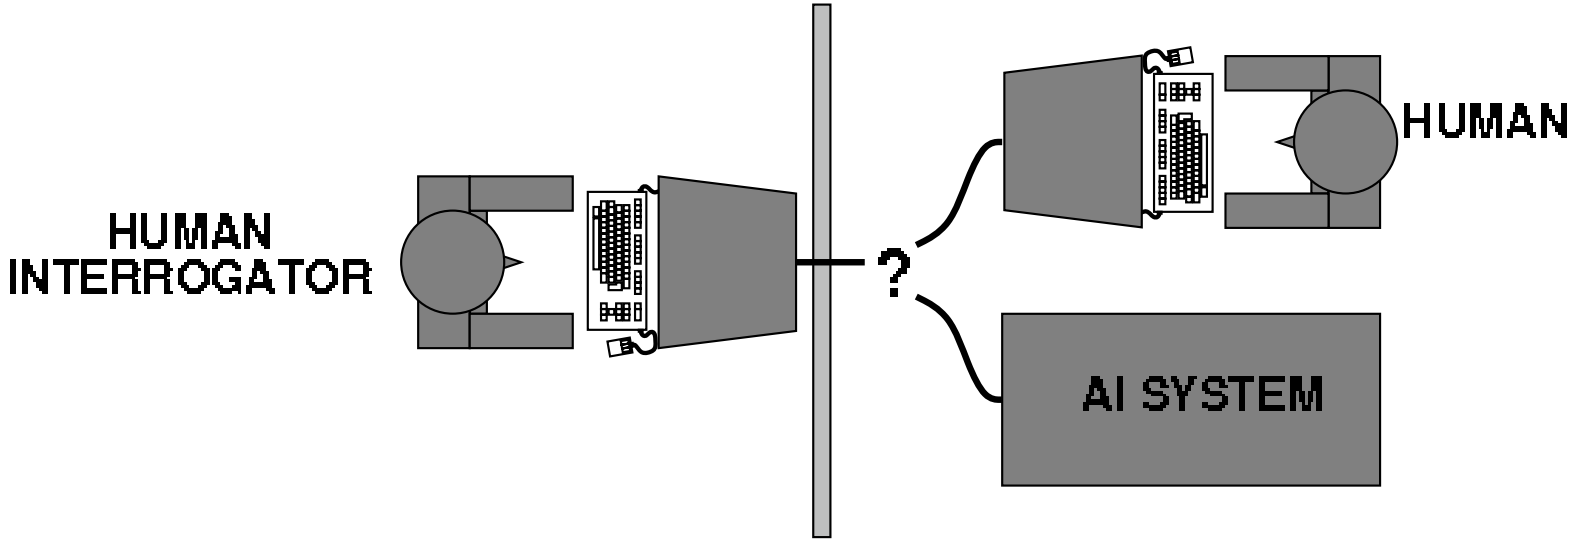

Acting Humanly: The Turing test

- "Can machines think?" $\longrightarrow$ "Can machines behave intelligently?"
- Operational test for intelligent behavior: the Imitation Game
- Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
- Anticipated all major arguments against AI in following 50 years
- Suggested major components of AI:

  knowledge, reasoning, language understanding, learning

Acting Humanly: The Turing test

- "Can machines think?" $\longrightarrow$ "Can machines behave intelligently?"
- Operational test for intelligent behavior: the Imitation Game
- Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
- Anticipated all major arguments against AI in following 50 years
- Suggested major components of AI:

  knowledge, reasoning, language understanding, learning

Problem: Turing test is not reproducible, constructive, or amenable to mathematical analysis

Thinking Humanly: Cognitive Science

- 1960s "cognitive revolution": information-processing psychology replaced prevailing orthodoxy of behaviorism
- Require scientific theories of internal activities of the brain
  - What level of abstraction? "Knowledge" or "circuits"?
  - How to validate? It requires
    1. Predicting and testing behavior of human subjects (top-down, Cognitive Science)
    2. Direct identification from neurological data (bottom-up, Cognitive Neuroscience)

Thinking Rationally: Laws of Thought

Several Greek schools at the time of Aristotle developed various forms of logic:

- Notation and rules of derivation for thoughts.
  They may or may not have proceeded to the idea of mechanization
- Direct line through mathematics and philosophy to modern AI
- Problems:
  1. Not all intelligent behavior is mediated by logical deliberation
  2. What is the purpose of thinking? What thoughts should I have?

Acting Rationally

- Rational behavior: doing the *right thing*, that which is expected to maximize goal achievement given the available information

- Doesn't necessarily involve thinking—e.g., blinking reflex—but thinking should be in the service of rational action

- Aristotle:
  Every art and every inquiry, and similarly every action and pursuit, is thought to aim at some good

Rational Agents

- An *agent* is an entity that perceives and acts
- This course is about designing rational agents
- Abstractly, an agent is a function from percept histories to actions:

  $$f : \mathcal{P}^* \rightarrow \mathcal{A}$$

  For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance
- Caveat: *computational limitations make perfect rationality unachievable*
- Approach: design best program for given machine resources

Summary of Experts' View of AI

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

Systems that *think* like humans

Systems that *think* rationally

Systems that *act* like humans

Systems that *act* rationally

Operational Definition of AI

Systems that *act* like humans

Turing test

Systems that *think* like humans

Cognitive Science

Systems that *think* rationally

Logic-based AI

Systems that *act* rationally

Rational Agents

AI's Foundations

Philosophy	logic, methods of reasoning, mind as physical system, foundations of learning, language, rationality
Mathematics	formal representation and proof, algorithms, computation, (un)decidability, (in)tractability, probability
Psychology	adaptation, perception and motor control, experimental techniques
Economics	formal theory of rational decisions
Linguistics	knowledge representation, grammar
Neuroscience	plastic physical substrate for mental activity
Control theory	homeostatic systems, stability, simple optimal agent designs

State of the art

Which of the following can be done at present?

- Play a decent game of table tennis
- Drive along a curving mountain road
- Drive in the center of Cairo
- Buy a week's worth of groceries at the supermarket
- Buy a week's worth of groceries on the web
- Win at Go
- Play a decent game of bridge
- Discover and prove a new mathematical theorem
- Write an intentionally funny story
- Give competent legal advice in a specialized area of law
- Translate spoken English into spoken Swedish in real time
- Perform a complex surgical operation

Why Study AI?

AI helps:

1. computer scientists and engineers build more useful and user-friendly computers
2. psychologists, linguists, and philosophers understand the principles that constitute what we call intelligence

AI is an interdisciplinary field of study

Many ideas and techniques now standard in CS (symbolic computation, time sharing, objects, declarative programming, . . .) were pioneered by AI-related research

Building AI systems is pretty hard

I went to the grocery store, I saw the milk on the shelf and I bought it

What did I buy?

- The milk?
- The shelf?
- The store?

An awful lot of knowledge of the world is needed to answer simple questions like this one