Introduction

CPSC 383: Explorations in Artificial Intelligence and Machine Learning
Fall 2025

Jonathan Hudson, Ph.D
Associate Professor (Teaching)
Department of Computer Science
University of Calgary

August 27, 2025

UNIVERSITY OF CALGARY

Outline

- What is AI?
- Why AI?
- Limits of AI
- Examples of AI

What is AI?

Think about it ape…

So, what is Artificial Intelligence (AI)

CPSC 383 Definition of AI:

AI deals with the development of systems either displaying a behavior humans associate with intelligence or solving a problem humans think only an intelligent being can solve.

- subjective definition
- changes over time!

Preference to use the term **Knowledge-based Systems** as this describes what systems must deal with and do → knowledge and knowledge processing

*(many consider regular computer science to be **information/data** processing systems, it is therefore implied that in AI **knowledge** is a further layer on top)*

AI: Moving Target

- **Lots of things you take for granted were called AI once.**

- Alexa/Siri/etc., google maps, biometrics, google search, automatic translation, natural language understanding, handwriting recognition, …

- We'll talk about the history of AI,
  - but in short, the **trough of disillusionment** in technology development is historic key

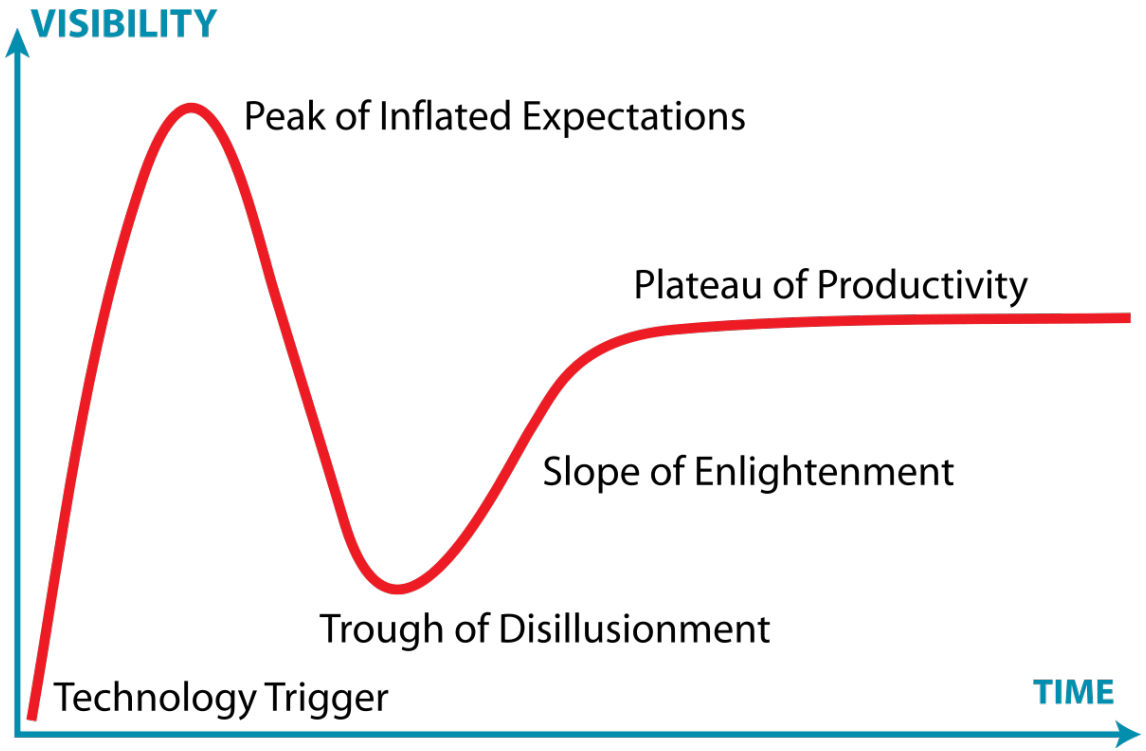

AI vs AGI

- **Artificial general intelligence** (**AGI**) is a **hypothetical** form of artificial intelligence (**AI**) capable of understanding, learning, and applying knowledge.
  - Like human cognitive abilities.
  - AGI is often the layman (non-professional) definition of AI

- Professional definition of AI includes **AGI** as well as the **currently existing specialized AI** systems.

- AGI, if it is created, would be able to **adapt** and **perform tasks not explicitly programmed** for, like human intelligence.

History: What is AI?

Think back about it ape…

Historically: What is AI?

Historically definitions can be divided dimensionally by defining AI as

1. Being like a human versus perfect rationality
2. Intelligence is achieved internally (thinking) versus externally (behaviour)

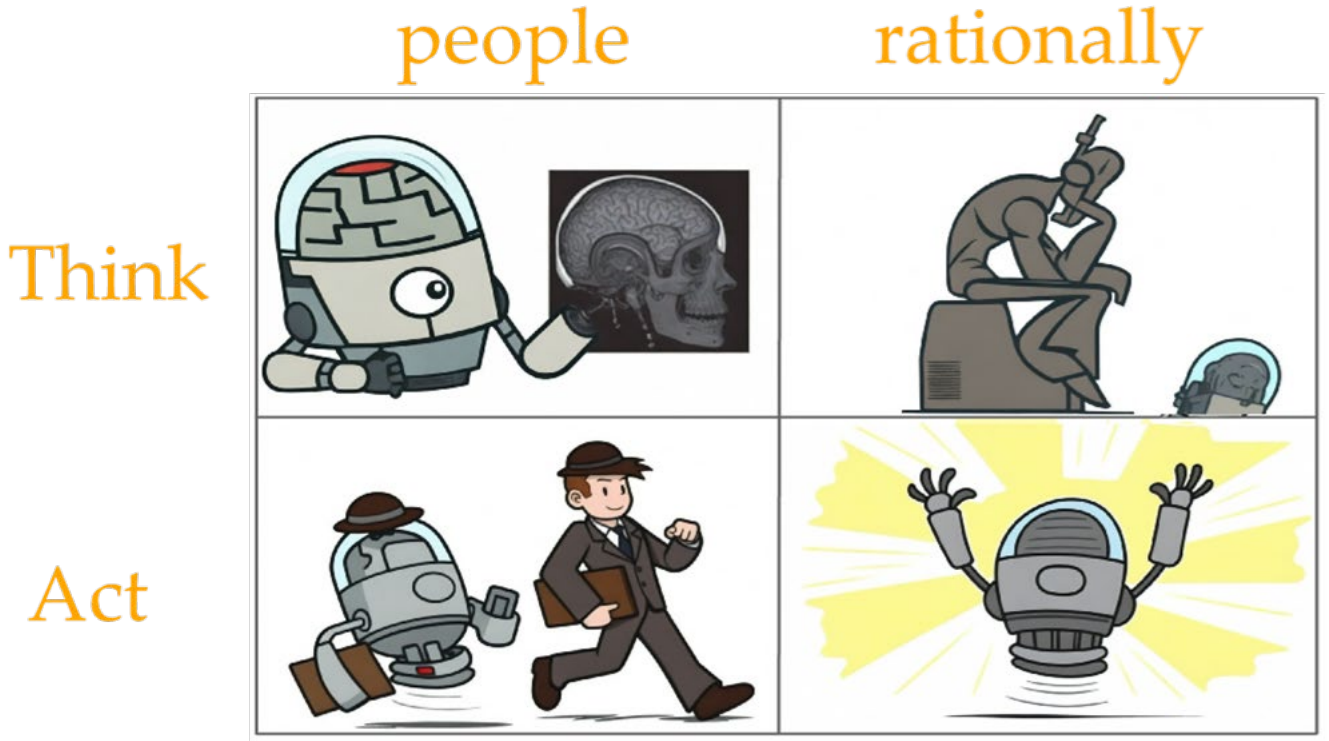

Rationality

Rational Decisions

- We'll use the term **rational** in a very specific, technical way:
  - Rational: maximally achieving pre-defined goals
  - Rationality only concerns what decisions are made
    - (not the thought process behind them)
  - Goals are expressed in terms of the **utility** of outcomes
- Being rational means **maximizing your expected utility**

Acting Human

Acting humanly: *The Turing Test*

- The Turing machine (universal model of complete computer) was developed in the 20th century by British mathematician and WWII codebreaker **Alan Turing** in paper "***Computing Machinery Intelligence***." [1950s]

- **Turing Test** is the oldest and most famous methods for testing consciousness in artificial intelligence.

- **The Imitation Game**
  - "Can machines think?" requires test of
    - → "Can machines behave intelligently?"

Turing Test

Turing Test

- The test requires a program to have a conversation (via typed messages) with an interrogator for five minutes
  - ELIZA program early example
  - Eugene Goostman program fooled 33% of the untrained amateur judges in Turing Test
- Problem: Turing test is not reproducible, constructive, or amenable to mathematical analysis
- Today we consider the Turing test as an **operational test of behaviour** and not a test of consciousness.

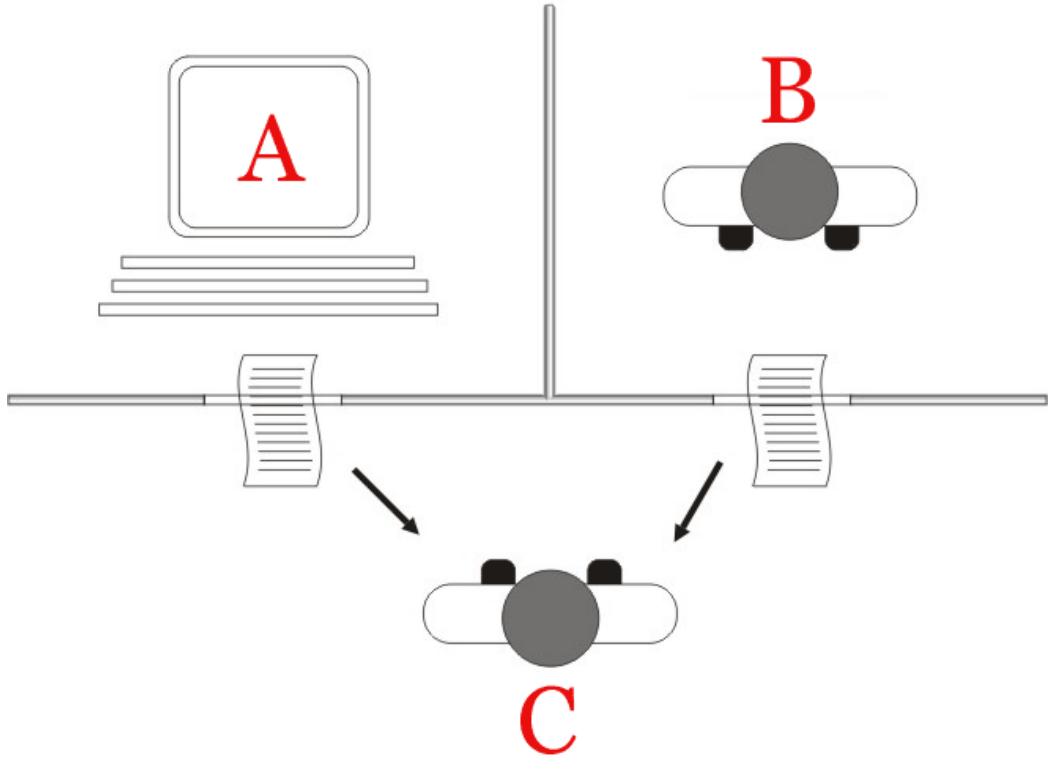

By Juan Alberto Sánchez Margallo - File:Test_de_Turing.jpg, CC BY 2.5, https://commons.wikimedia.org/w/index.php?curid=57298943

Thinking Human

Thinking humanly: *Cognitive Science*

- 1960s “cognitive revolution”:
  - information-processing psychology replaced behaviourism

- 1. Cognitive Science and
- 2. Cognitive Neuroscience now distinct from AI

- Like CPSC AI, their available theories do not yet explain anything resembling human-level general intelligence

Thinking Rationally

Thinking rationally: *Laws of Thought*

- Aristotle: what are correct arguments/thought processes?

- Several Greek schools developed various forms of logic:
  - notation and rules of derivation for thoughts;
  - may or may not have proceeded to the idea of mechanization

Propositional logic

- A Porsche is a black car.
- Black cars are fast cars.
- Bad cars are slow cars.

- $porsche \ \wedge black$
- $black \rightarrow fast$
- $bad \rightarrow \neg\, fast$

- A Porsche is a good car.
- $porsche \wedge \neg bad$
- ? $porsche \wedge good$

Logic Dangers

- But be careful: “tertium non datur” (no third possibility is given)
  - The car is green =: p
  - The car is red =: q
  - We need in addition:
    $q \leftrightarrow \neg p$

- What if there are more colours?
- What if there are more cars?
- What if answer is not True/False?
- What if we can’t tell if there’s an answer?

Acting Rationally

Acting rationally: *Rational Agents*

- **Rational behaviour**: doing the right thing
- The right thing?
  - that which is expected to maximize goal achievement
    - given the available information

Why AI?

Think about it harder ape…

“First solve AI, then use AI to solve everything else.”

Demis Hassabis, CEO of Google DeepMind

Short Term: Benefits/Risks

- **Benefits:**
  - Decrease repetitive work
  - Increase production of goods and services
  - Accelerate scientific research (disease cures, climate change and resource shortages solutions)
- **Risks:**
  - Lethal autonomous weapons, surveillance and persuasion
  - Biased decision making
  - Intellectual property abuse
  - Fake reality generation
  - Impact on employment
  - Safety-critical applications, cybersecurity threats

Long Term: Science Fiction Concerns

- Development of an artificial superintelligence that surpasses human intelligence may pose a significant risk

- Analogous to the "Gorilla problem"
- Humans and gorillas evolved from the same species, but humans have more control than other primates.

- Thus, we should design AI systems in such a way that they do not end up taking control in the way that Turing suggests they might.

Limits of AI?

Not so fast buckaroo

Limits of AI

- Philosopher John Searle (1980):
  - **weak AI**:
    - the idea that machines could act as if they were intelligent
  - **strong AI**:
    - the assertion that machines that do so are actually consciously thinking (not just simulating thinking)

Arguments

- **The argument from informality**
  - Turing's "argument from informality of behaviour" says that human behaviour is far too complex to be captured by any formal set of rules

- **The argument from disability**
  - The "argument from disability" makes the claim that "a machine can never do X."

Mathematical Objection

- **The mathematical objection**
- Turing (1936) and Gödel (1931) proved that certain mathematical questions are in principle unanswerable by particular formal systems.
- Philosophers such as J. R. Lucas (1961) have claimed that this theorem shows that machines are **mentally inferior to humans**

- **Problems with Lucas' claim:**
  - No entity—human or machine—can prove things that are impossible to prove
  - incompleteness theorem technically applies only to formal systems that are powerful enough to do arithmetic.

Conventions

- Some philosophers claim that a machine that acts intelligently would not be actually thinking, but would be only a simulation of thinking

- Turing ‘argued’ the **polite convention** that everyone and machine think.
  - **Turing test only asked for Weak AI**

- John Searle rejects the **polite convention**
  - **Declaration of Strong AI different from Weak AI**

The Chinese room

- **The Chinese room**
  - A human, who understands only English
  - Inside a room that contains a rule book, written in English
  - Pieces of paper containing indecipherable symbols are slipped under the door to the room
  - The human follows the instructions in the rule book to convert these symbols to other symbols
  - At some point the rule book has human return symbols back under the door

  - It is given that the human does not understand Chinese
    - The ‘indecipherable symbols’ of the example

  - Computers are in essence doing the same thing
    - Therefore computers generate no understanding

Next… history and definitions

Jonathan Hudson, Ph.D.
jwhudson@ucalgary.ca
https://cspages.ucalgary.ca/~jwhudson/

UNIVERSITY OF CALGARY