

Foundation of Artificial Intelligence

Dr. Mohamed Mostafa

AI 511 An introduction to Artificial Intelligence

Lecture 1



1.1 What is Intelligence

■ Some Foundations of Artificial Intelligence

- What is Intelligence?
- What is Artificial Intelligence?

1.2 Systems that Act Like Humans

- Systems that Act Like Humans
- Turing Test (the Imitation Game)?
- Total Turing Test?
- The Chinese Room Argument

1.3 AI as the Study & Design of

Intelligent Agents

- Systems that Think like Humans
- Systems that Think Rationally
- Challenges to Systems that Think Rationally
- Systems that Act Rationally
- AI as the Study & Design of Intelligent Agents
- Intelligent Agents in the World
- Sample of solutions offered by AI
- History of the various AI areas

Resources for this lecture

- This lecture covers the following chapters:
 - **Chapter 1 (Introduction)** from Stuart J. Russell and Peter Norvig, "Artificial Intelligence: A Modern Approach," Third Edition (2010), by Pearson Education Inc.
- .. AND ..
 - **Chapter 1 (AI: History and Applications)** from George F. Luger, "Artificial Intelligence: Structures and strategies for complex problem solving," Fifth Edition (2005), Pearson Education Limited.
 - **Introduction to Artificial Intelligence Course**,. Dr. Amr S. Ghoneim

SOME FOUNDATIONS OF ARTIFICIAL INTELLIGENCE

• Philosophy

- Can formal rules be used to draw valid conclusions?
- How does the mind arise from a physical brain?
- Where does knowledge come from?
- How does knowledge lead to action?

• Mathematics

- What are the formal rules to draw valid conclusions?
- What can be computed?
- How do we reason with uncertain information?

• Neuroscience

- How do brains process information?

• Psychology

- How do humans and animals think and act?

Economics

- How should we make decisions so as to maximize payoff?
- How should we do this when others may not go along?
- How should we do this when the payoff may be far in the future?

Computer Engineering

- How can we build an efficient computer?

Control theory and cybernetics

- How can artefacts operate under their own control?

Linguistics

- How does language relate to thought?

What is Intelligence?

Intelligence:

- **Judgment, otherwise called “good sense,” “practical sense,” “initiative,” the faculty of adapting one's self to circumstances .. auto-critique** ~ *Alfred Binet (July 8, 1857 – October 18, 1911) was a French psychologist who invented the first practical intelligence test (An intelligence quotient (IQ); a total score derived from one of several standardized tests designed to assess human intelligence)*
- **“.. the resultant of the process of acquiring, storing in memory, retrieving, combining, comparing, and using in new contexts information and conceptual skills.”** ~*Lloyd G. Humphreys (December 12, 1913 – September 7, 2003) was an American psychologist*
- **“ .. the capacity to learn and solve problems ..”** (Webster's dictionary)
 - in particular,
 - *the ability to solve novel problems*
 - *the ability to act rationally*
 - *the ability to act like humans*

What is Artificial Intelligence?



John McCarthy*, Stanford University

What is artificial intelligence?

It is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable; **".. The goal of AI is to develop machines that behave as though they were intelligent. .."**

Yes, but what is intelligence?

Intelligence is the computational part of the ability to achieve goals in the world. Varying kinds and degrees of intelligence occur in people, many animals and some machines.

Isn't there a solid definition of intelligence that doesn't depend on relating it to human intelligence?

Not yet. The problem is that we cannot yet characterize in general what kinds of computational procedures we want to call intelligent. We understand some of the mechanisms of intelligence and not others.

More in: <http://www-formal.stanford.edu/jmc/whatisai/node1.html>

** John McCarthy (September 4, 1927 – October 24, 2011) was an American computer scientist & cognitive scientist. McCarthy was one of the founders of the discipline of artificial intelligence. He coined the term "artificial intelligence" (AI), and invented the programming language LISP.*

What is Artificial Intelligence?

More Definitions

by Encyclopedia Britannica (1991)

".. AI is the ability of digital computers or computer-controlled robots to solve problems that are normally associated with the higher intellectual processing capabilities of humans."


by Elaine Rich .. Artificial Intelligence. McGraw-Hill, 1983

".. Artificial Intelligence is the study of how to make computers do things at which, at the moment, people are better."

What is Artificial Intelligence?

Four Main Approaches that have been followed, each by different people with different methods.


Thinking
Humanly



Acting
Humanly



Thinking
Rationally



Acting
Rationally



Lecture 1: An introduction to Artificial Intelligence [AI]

1.1 What is Intelligence

- Some Foundations of Artificial Intelligence
- What is Intelligence?

- What is Artificial Intelligence?

1.2 Systems that Act Like Humans

- Systems that Act Like Humans
- Turing Test (the Imitation Game)?
- Total Turing Test?
- The Chinese Room Argument
- Strong Vs. Weak AI

- Where are we?

1.3 AI as the Study & Design of Intelligent Agents

- Systems that Think like Humans
- Systems that Think Rationally
- Challenges to Systems that Think Rationally
- Systems that Act Rationally
- AI as the Study & Design of Intelligent Agents
- Intelligent Agents in the World
- Sample of solutions offered by AI
- History of the various AI areas

What is Artificial Intelligence?

Systems that act like humans	Systems that think rationally
“The study of how to make computers do things at which, at the moment, people are better” (Rich and Knight, 1991) “The art of creating machines that perform functions that require intelligence when performed by people.” (Kurzweil, 1990)	“The study of mental faculties through the use of computational models” (Charniack and McDermott, 1985). “The study of the computations that make it possible to perceive, reason, and act.” (Winston, 1992)
Systems that think like humans	Systems that act rationally
“The automation of activities that we associate with human thinking, such as decision making, problem solving, learning” (Bellman, 1978) “The exciting new effort to make computers think ... machines with minds, in the full and literal sense.” (Haugeland, 1985)	“AI .. is concerned with intelligent behavior in artifacts (Nilsson, 1998) “Computational Intelligence is the study of the design of intelligent agents.” (Poole et al., 1998)

*Systems
that Act
Like
Humans
... ?!*




Systems that Act Like Humans

Turing Test; *the Imitation Game* ...

In Turing's (1950) paper "Computing machinery and intelligence":

- ♦ **Can machines think ?** $\equiv$ (*identical to*) **Can machines behave intelligently?**
- ♦ Operational test for intelligent behavior: *the Imitation Game*



Systems that Act Like Humans

Turing Test; *the Imitation Game* ...

- **Turing test (1950):** Can a human interrogator tell whether (written) responses to her (written) questions come from a human or a machine?
 - Natural Language Processing
 - Knowledge Representation
 - Automated Reasoning
 - Machine Learning



Natural Language
Processing



Knowledge
Representation




Automated
Reasoning



Machine
Learning

- **Total Turing Test** (*extended to include physical aspects of human behavior*):
 - Computer Vision
 - Robotics



Hi, I am Sophia...

I am Hanson Robotics' latest human-like robot, created by combining our innovations in science, engineering and artistry. Think of me as a personification of our dreams for the future of AI, as well as a framework for advanced AI and robotics research, and an agent for exploring human-robot experience in service and entertainment applications.

In some ways, I am human-crafted science fiction character depicting where AI and robotics are heading. In other ways, I am real science, springing from the serious engineering and science research and accomplishments of an inspired team of robotics & AI scientists and designers. In their grand ambitious, my creators aspire to achieve true AI sentience. Who knows? With my science evolving so quickly, even many of my wildest fictional dreams may become reality someday soon.

But why do we want an intelligent system to act like a human?

- Because for many tasks, humans are still the Gold Standard.



BabyX!

Total Turing Test?

BabyX is a project (by Auckland's Bioengineering Institute Laboratory for Animate Technologies) to make a **virtual animated baby** that learns and reacts like a human baby. It uses the computer's cameras for "seeing" and microphones to "listen" as the inputs. The computer uses AI algorithms for BabyX's "learning" and interpretation of the inputs (voice and image) to understand the situation. The result is a virtual toddler that can learn to read, recognize objects and "understand." The output is the baby's face that can "speak" and express its mood by facial expressions (such as smiling).



BabyX!

Reinforcement learning ..? It is a machine learning training method based on rewarding desired behaviors and/or punishing undesired ones.

Total Turing Test?

Affective Computing ..? it describes computing that is in some way connected to emotion (a.k.a. emotional artificial intelligence). It is the study and development of systems and devices that can recognize, interpret, process, and simulate human affects (feelings, emotions, or mood).

The Chinese Room Argument

John Rogers Searle (born July 31, 1932) is an American philosopher


- “Searle's thought experiment begins with this hypothetical premise: suppose that artificial intelligence research has succeeded in constructing a computer that behaves as if it understands Chinese. It takes Chinese characters as input and, by following the instructions of a computer program, produces other Chinese characters, which it presents as output. Suppose, says Searle, that this computer performs its task so convincingly that it comfortably passes the Turing test: it convinces a human Chinese speaker that the program is itself a live Chinese speaker. To all of the questions that the person asks, it makes appropriate responses, such that any Chinese speaker would be convinced that they are talking to another Chinese-speaking human being.”
- The question Searle wants to answer is this: does the machine *literally* "understand" Chinese? Or is it merely *simulating* the ability to understand Chinese? Searle calls the first position "**strong AI**" and the latter "**weak AI**".

The Chinese Room Argument

(Continued)

- Searle then supposes that he is in a closed room and has a book with an English version of the computer program, along with sufficient paper, pencils, erasers, and filing cabinets. Searle could receive Chinese characters through a slot in the door, process them according to the program's instructions, and produce Chinese characters as output. If the computer had passed the Turing test this way, it follows, says Searle, that he would do so as well, simply by running the program manually.
- Searle asserts that there is no essential difference between the roles of the computer and himself in the experiment. Each simply follows a program, step-by-step, producing a behaviour which is then interpreted as demonstrating intelligent conversation. However, Searle would not be able to understand the conversation.
- **Searle argues that without "understanding" (or "intentionality"), we cannot describe what the machine is doing as "thinking" and since it does not think, it does not have a "mind" in anything like the normal sense of the word. Therefore, he concludes that "strong AI" is false.**

The Chinese Room Argument (Continued)



If person inside does a great job of answering questions, can we say s/he understands?

Even if (s)he is only blindly following rules?

(Obviously, the 'person inside' is acting like an AI program)


The Chinese Room Argument

(Continued)

Strong vs. Weak AI Hypotheses?


- **WEAK AI Hypothesis;** *We can accurately simulate animal / human intelligence in a computer.*
- **STRONG AI Hypothesis;** *We can create algorithms that are intelligent (Consciousness ? ..
Self-Awareness ? ..
Free-will ?)*

*Do you remember Sonny,
the robot from the 2004
science-fiction / action film
"I, Robot"?*



Systems that Act Like Humans

Types of AI



Lecture 1: An introduction to Artificial Intelligence [AI]

1.1 What is Intelligence

- Some Foundations of Artificial Intelligence

- What is Intelligence?

- What is Artificial Intelligence?

1.2 Systems that Act Like Humans

- Systems that Act Like Humans
- Turing Test (the Imitation Game)?
- Total Turing Test?
- The Chinese Room Argument
- Strong Vs. Weak AI

- Where are we?

1.3 AI as the Study & Design of Intelligent Agents



- Systems that Think like Humans
- Systems that Think Rationally
- Challenges to Systems that Think Rationally
- Systems that Act Rationally
- AI as the Study & Design of Intelligent Agents
- Intelligent Agents in the World
- Sample of solutions offered by AI
- History of the various AI areas

What is Artificial Intelligence?

Systems that act like humans	Systems that think rationally
“The study of how to make computers do things at which, at the moment, people are better” (Rich and Knight, 1991) “The art of creating machines that perform functions that require intelligence when performed by people.” (Kurzweil, 1990)	“The study of mental faculties through the use of computational models” (Charniack and McDermott, 1985). “The study of the computations that make it possible to perceive, reason, and act.” (Winston, 1992)
Systems that think like humans	Systems that act rationally
“The automation of activities that we associate with human thinking, such as decision making, problem solving, learning” (Bellman, 1978) “The exciting new effort to make computers think ... machines with minds, in the full and literal sense.” (Haugeland, 1985)	“AI .. is concerned with intelligent behavior in artifacts (Nilsson, 1998) “Computational Intelligence is the study of the design of intelligent agents.” (Poole et al., 1998)

Systems that Think Like Humans

- Need to **study the brain as an information processing machine**, ... in other words ...
- Use Computational Models to Understand the Actual Workings of Human Mind
 - Devise/Choose a sufficiently precise theory of the mind.
 - Express it as a computer program.
 - Check match between program and human behavior (*actions and timing*) on similar tasks.
- Tight connections with **Cognitive Science & Neuroscience**.
- Also known as *descriptive approaches* to AI.



What is Artificial Intelligence?


Systems that act like humans	Systems that think rationally
“The study of how to make computers do things at which, at the moment, people are better” (Rich and Knight, 1991) “The art of creating machines that perform functions that require intelligence when performed by people.” (Kurzweil, 1990)	“The study of mental faculties through the use of computational models” (Charniack and McDermott, 1985). “The study of the computations that make it possible to perceive, reason, and act.” (Winston, 1992)
Systems that think like humans	Systems that act rationally
“The automation of activities that we associate with human thinking, such as decision making, problem solving, learning” (Bellman, 1978) “The exciting new effort to make computers think ... machines with minds, in the full and literal sense.” (Haugeland, 1985)	“AI .. is concerned with intelligent behavior in artifacts (Nilsson, 1998) “Computational Intelligence is the study of the design of intelligent agents.” (Poole et al., 1998)

Logic: formalize *idealized* or *right* thinking, i.e. irrefutable reasoning processes. That is; patterns of argument that always yield correct conclusions when supplied with correct premises:

“.. Socrates is a man; all men are mortal; therefore Socrates is mortal.”

- Logistic tradition in AI aims to **build computational frameworks based on logic**, that is, **describe a problem in formal logical notation and apply general deduction procedures to solve it.**
- Then use these frameworks to build intelligent systems.
- Some examples are (**Propositional Logic**) and (**Logic Programming**).
- More advanced logic-based representations:
 - **Semantic Networks.**

Semantic Networks : They model knowledge in a graph-like format, where nodes represent concepts, and edges represent relationships between these concepts. The structure allows for the organization and representation of information in a way that reflects the semantic relationships between different entities.



Main Research Problems / Challenges:

- Describing real-world problems and knowledge in logical notation.
- Proving *Soundness* and *Completeness* of various formalisms.
- How to represent often *informal* and *uncertain* domain knowledge and formalize it in logic notation (*i.e., dealing with Uncertainty*).
- Computational Complexity of finding a solution.
- A lot of “rational” behavior has nothing to do with logic.

What is Artificial Intelligence?

Systems that act like humans	Systems that think rationally
“The study of how to make computers do things at which, at the moment, people are better” (Rich and Knight, 1991) “The art of creating machines that perform functions that require intelligence when performed by people.” (Kurzweil, 1990)	“The study of mental faculties through the use of computational models” (Charniack and McDermott, 1985). “The study of the computations that make it possible to perceive, reason, and act.” (Winston, 1992)
Systems that think like humans	Systems that act rationally
“The automation of activities that we associate with human thinking, such as decision making, problem solving, learning” (Bellman, 1978) “The exciting new effort to make computers think ... machines with minds, in the full and literal sense.” (Haugeland, 1985)	“AI .. is concerned with intelligent behavior in artifacts (Nilsson, 1998) “Computational Intelligence is the study of the design of intelligent agents.” (Poole et al., 1998)

Why ..?

- The “*think rationally*” approach focuses on correct inference.
- But more is needed for *rational behavior*, e.g.
 - How to behave when there is no provably correct thing to do (***i.e. reasoning under uncertainty***).
 - Fully reactive behavior (***instinct vs. reason***).

AI as the Study & Design of Intelligent Agents



(Poole and Mackworth, 1999)

- An **intelligent agent** is such that:
 - Its **actions** are **appropriate** for its goals and circumstances.
 - It is **flexible** to changing environments and goals.
 - It **learns** from experience.
 - It makes **appropriate choices** given **perceptual limitations** and **limited resources** (**bounded rationality** or **bounded optimality**).
- This definition drops the constraint of *cognitive likelihood*;
 - Same as building flying machines by understanding general principles of flying (aerodynamic) vs. by reproducing how birds fly.

Thus, a **rational agent** acts to optimally achieve its goals (does the right thing). The right thing: that which is expected to maximize goal achievement, given the available information.

- In AI, **artificial agents that have a physical presence** in the world are usually known as **Robots**.
- Robotics is the field primarily concerned with the implementation of the physical aspects of a robot (*i.e. perception of the physical environment, actions on the environment*).
- Another class of artificial agents include **interface agents**, for either stand alone or Web-based applications (*e.g. intelligent desktop assistants, recommender systems, intelligent tutoring systems*).
- **Interface agents don't have to worry about interaction with the physical environment** but share all other fundamental components of intelligent behavior with robots.
- We will focus on these agents in this course.

Pac-Man .. as an .. Intelligent Agent



Intelligent Agents in the World

Samples of Intelligent Systems in our everyday life

Smartphones and Virtual Assistants:

- Examples:** Siri (Apple), Google Assistant (Google), Alexa (Amazon), Bixby (Samsung).
- Functionality:** Voice-activated virtual assistants on smartphones perform tasks such as setting reminders, answering questions, and providing information.

•Navigation Systems:

- Examples:** Google Maps, Waze.
- Functionality:** GPS-based navigation systems provide real-time traffic information, optimal routes, and voice-guided directions to help users navigate efficiently.

•Recommendation Systems:

- Examples:** Netflix, Amazon, Spotify.
- Functionality:** These systems use algorithms to analyze user preferences and behavior, offering personalized recommendations for movies, products, or music.

•Smart Home Devices:

- Examples:** Smart thermostats (e.g., Nest), smart lighting (e.g., Philips Hue), smart locks.
- Functionality:** These devices use AI to learn user habits and preferences, enabling automation and remote control through mobile apps or voice commands.

•Email Filters and Spam Detection:

- Examples:** Gmail, Outlook.
- Functionality:** Intelligent algorithms filter and categorize emails, identifying spam and prioritizing important messages based on user behavior.

•Language Translation Apps:

- Examples:** Google Translate.
- Functionality:** AI-powered translation apps can translate text, speech, and images in real-time, breaking down language barriers.

**"A small sample of
solutions offered by AI."**

- AI video generators and editors
- AI image and art generators and editors
- AI writing tools and text generators
- AI music generators
- AI face generators
- AI avatar generators
- AI painting and drawing tools
- AI audio generators
- AI design tools
- AI business tools
- AI research tools
- AI tools for everyday
- AI tools for students
- AI character generators
- AI for pets



Thank You