



مجموعة جامعة
الجميع الذكية
للتعليم القابضة



Master of Science in Artificial intelligent

Lecture 1

AI511

Foundations of Artificial Intelligence

Prepared by
Prof. Lamiaa Fattouh Ibrahim

www.midocean.ae



An Introduction to Artificial Intelligence

This lecture covers the following

chapters:

Resources for this lecture

- **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.




Outline

What is Intelligence

- **Some Foundations of Artificial Intelligence**
- **What is Intelligence?**
- **What is Artificial Intelligence?**
 - **Systems that act like humans**
 - **Turing Test; the Imitation Game ...**
 - **The Chinese room argument**
 - **Strong vs. Weak AI Hypotheses**
 - **Systems That Think Like Humans**
 - **Systems that Think Rationally**
 - **Systems that Act Rationally**


ARTIFICIAL INTELLIGENCE:INTRODUCTION

- Artificial Intelligence does not only mean robots or Sci-Fi movies!
- Artificial Intelligence applications are everywhere!
- Google search engine, amazon recommender systems, Facebook facial recognition (tagging), Siri





MACHINE LEARNING AND ARTIFICIAL INTELLIGENCE: PROCESS

- Machine learning is the study of algorithms that teach computers to learn from experience.
- Through experience (more training data), computers can continuously improve their performance.



MACHINE LEARNING: BIG PICTURE





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? **Logic**
- What can be computed? **Algorithm**
- How do we reason with uncertain information? **Probability**

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*

INTELLIGENCE

- **“the capacity to learn and solve problems”** (Websters dictionary)
in particular
 - *the ability to solve novel problems*
 - *the ability to act rationally*
 - *the ability to act like humans*

Knowledge Transfer

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. ..”**

Lisp Language=> List processing Language

What is Artificial Intelligence?

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**



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)

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)

Exercises

What have we learned?

- Read the following: [Chapter 1 (Introduction) from Stuart J. Russell and Peter Norvig, "Artificial Intelligence: A Modern Approach."] then describe briefly in a point or two how the following disciplines contribute/d to AI:
 - Psychology
 - Computer Engineering
 - Neuroscience
 - Economics
- What are the four Approaches to AI according to Russell & Norvig?

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**





Figure 1.1 The Turing test.


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


Context



- Natural Language Processing (NLP)
 - ✓ Natural Language Understanding
 - ✓ Natural Language Synthesis or generation
- Knowledge Representation
 - ✓ to store knowledge




- Automated Reasoning: to use the stored information and answer questions and to draw new conclusion.
- Machine Learning : to adapt to new circumstances and to detect and extrapolate patterns.



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

Turing's Test

- If the human cannot tell whether the responses from the other side of a wall are coming from a human or computer, then the computer is intelligent. **Machine pass the Turing's Test.**





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.


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



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




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 he or she understands?

Even if she or 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 (conscious?)*

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



Strong Vs. Weak AI .. Where are we?


Systems that Act Like Humans

Strong Vs. Weak AI .. Where are we?


Stage 3: Machine Consciousness



Stage 2: Machine Intelligence



Stage 1: Machine Learning



Three Stages of AI

@AmitPaka

Sources: <https://www.scoro.com/blog/artificial-intelligence-everything-you-want-to-know/>
<https://www.ubs.com/microsites/artificial-intelligence/en/new-dawn.html>

15

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)

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 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

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 Rationally

- 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).

Exercises

What have we learned?

- Describe briefly the Turing Test "imitation game". (Illustrate through drawing)
- Describe briefly the Total Turing Test.
- Describe briefly the "Weak AI Hypothesis" versus the "Strong AI Hypothesis".
- Criticize Turing's criteria for computer software being "intelligent"; What is Searle's thought experiment (the Chinese Room Argument)?
- Describe briefly Reinforcement Learning.
- Describe briefly Affective Computing.

Additional Resources

- The Turing test: Can a computer pass for a human? - Alex Gendler
<https://www.youtube.com/watch?v=3wLqsRLvV-c>
- The Chinese Room Argument:
<https://www.youtube.com/watch?v=18SXA-G2peY>
- The Chinese Room - 60-Second Adventures in Thought:
<https://www.youtube.com/watch?v=TryOC83PH1g>
- We Talked To Sophia – The AI Robot That Once Said It Would ‘Destroy Humans’:
<https://www.youtube.com/watch?v=78-1MlkxyqI>
- This Freaky Baby Could Be the Future of AI. Watch It in Action:
<https://www.youtube.com/watch?v=yzFW4-dvFDA>
- Two robots debate the future of humanity:
<https://www.youtube.com/watch?v=1y3XdwTa1cA>



مجموعة جامعة
الجميع الذكية
للتعليم القابضة

MIDOCEAN
UNIVERSITY
BY EVERYONE'S SMART UNIVERSITY



Thank You

