Introduction to Artificial Intelligence

Chapter 1: Introduction
(1) Artificial Intelligent

Nguyễn Hải Minh
nhminh@fit.hcmus.edu.vn

Intelligence vs Artificial Intelligence

- What is Intelligence?
  - To know
  - To think
  - To talk
  - To learn
  - …

Knowledge + ability to perceive, feel, comprehend, process, communicate, judge, learn.

Intelligence vs Artificial Intelligence

- What is Artificial Intelligence?
  - To make computers that perceive
  - To make computers that think
  - To make computers that talk
  - To make computers that learn

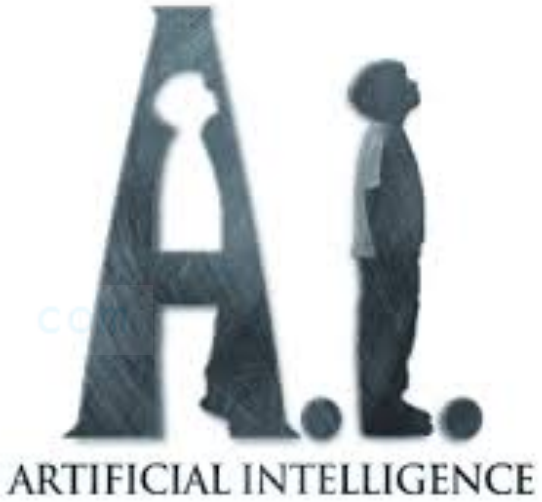

The study to develop techniques and tools for solving problems that people at good at.

What is Artificial Intelligence ?

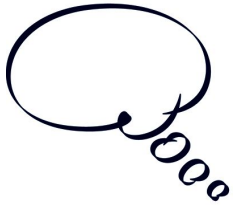

	HUMANS	RATIONAL
THOUGHT	Systems that **think** like humans	Systems that **think** rationally
BEHAVIOR	Systems that **act** like humans	Systems that **act** rationally

Systems that act like humans

- **When does a system behave intelligently?**

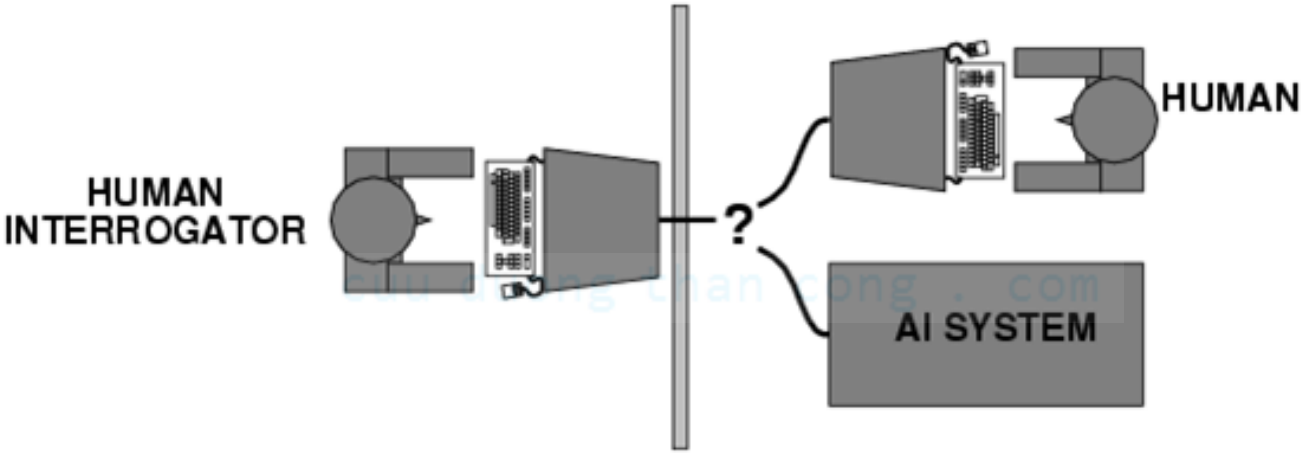

The Turing Test (1950)

- Put a machine and a human behind a wall and communicate with a Q & A dialog using a keyboard.
- If you, the human, cannot tell which response is coming from the machine and which from the human then we call the machine Intelligent.

Systems that act like humans

- Problem with Turing Test:

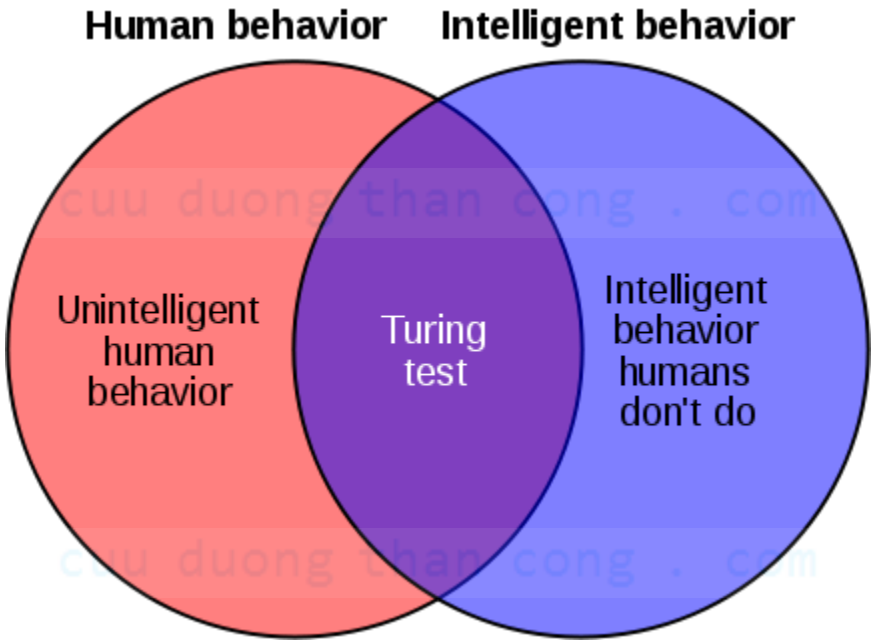

Systems that act like humans

- Variations
  - Reverse Turing Test: CAPTCHA
  - Total Turing Test

- QUIZ:
  - Can you think of a better Turing test?
    - Hint: question that can be answered easily by human but confused by a machine.

What is Artificial Intelligence ?

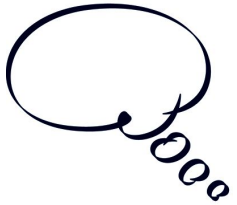

	HUMANS	RATIONAL
THOUGHT	Systems that think like humans	Systems that think rationally
BEHAVIOR	Systems that act like humans	Systems that act rationally

Systems that think like humans

- How do humans think? Inside human minds:
  - Introspectation
  - Psychological experiments
  - Brain imaging
- GPS: General Problem Solver (Newell and Simon, 1961)
  - Not solve the problem directly
  - Compare the trace of its reasoning steps to humans
- Cognitive Science
  - Computer models from AI + experiment techniques from psychology → testable theories of human mind.

Systems that think like humans

- These approaches are now distinct from AI
  - Share that the available theories but do not explain anything resembling human intelligence.
  - All share a principal direction.

What is Artificial Intelligence ?

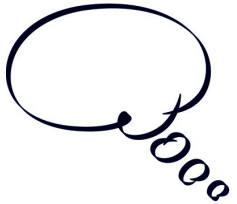

	HUMANS	RATIONAL
THOUGHT	Systems that think like humans	Systems that think rationally
BEHAVIOR	Systems that act like humans	Systems that act rationally

Systems that think rationally

- **Capturing the laws of thought**
  - Aristotle: What are 'correct' arguments and thought processes?
    - Correctness depends on irrefutability of reasoning processes.
  - This study initiated the field of logic.
    - The logic tradition in AI hopes to create intelligent systems using logic programming.
- **Problems:**
  - Not all intelligence is mediated by logic behavior
  - What is the purpose of thinking? What thought should one have?

What is Artificial Intelligence ?

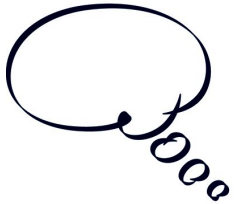

	HUMANS	RATIONAL
THOUGHT	Systems that think like humans	Systems that think rationally
BEHAVIOR	Systems that act like humans	Systems that act rationally

Systems that act rationally

- Rational behavior: "doing the *right thing*"
- The "*Right thing*": what is expected to maximize goal achievement given the available information.

→The Rational Approach:

- A rational agent: acts to achieve the best (expected) outcome.
- Include thinking, inference as a part of being rational agent.
- Include more: Action without thinking: e.g. reflexes.

Systems that act rationally

- Study AI as rational agent –

  2 advantages:

  - It is more general than using logic only
  - It allows extension of the approach with more scientific methodologies

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:

$$[f: P^* \rightarrow A]$$

- Drawback:
  - Yet rationality is only applicable in idealenvironments.
  - Moreover rationality is not a very good model of reality.

→ design best program for given machine resources

Roles of AI

AI has two major roles:

- Study the intelligent part concerned with humans.
- Represent those actions using computers.

Goals of AI

- To make computers more useful by letting them take over dangerous or tedious tasks from human
- Understand principles of human intelligence

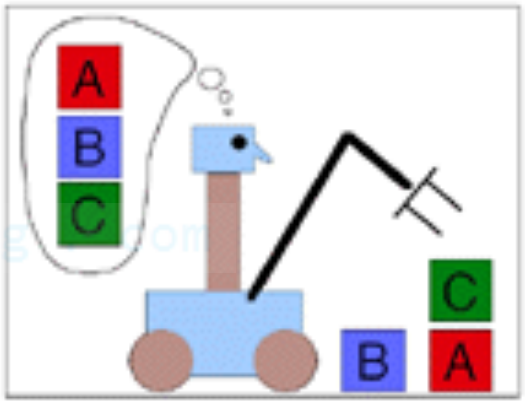

Foundation of AI

Field	Description
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.
Economics	utility, decision theory, rational economic agents
Neuroscience	neurons as information processing units.
Psychology/ Cognitive Science	how do people behave, perceive, process information, represent knowledge.
Computer Engineering	building fast computers
Control Theory	design systems that maximize an objective function over time
Linguistic	knowledge representation, grammar

The main topics in AI

- Artificial intelligence can be considered under a number of headings:
  - Search (includes Game Playing).
  - Representing Knowledge and Reasoning with it.
  - Planning.
  - Learning.
  - Natural language processing.
  - Expert Systems.
  - Interacting with the Environment
    (e.g. Vision, Speech recognition, Robotics)

We won't have time in this course to consider all of these.

Some Advantages of Artificial Intelligence

- ✓ More powerful and more useful computers
- ✓ New and improved interfaces
- ✓ Solving new problems
- ✓ Better handling of information
- ✓ Relieves information overload
- ✓ Conversion of information into knowledge

Some Disadvantages of Artificial Intelligence

- ✓ Increased costs
- ✓ Difficulty with software development - slow and expensive
- ✓ Few experienced programmers

Search

- *Search* is the fundamental technique of AI.
  - Possible answers, decisions or courses of action are structured into an abstract space, which we then search.
- Search is either "blind" or “informed":
  - blind
    - we move through the space without worrying about what is coming next, but recognising the answer if we see it
  - informed
    - we guess what is ahead, and use that information to decide where to look next.
- We may want to search for the first answer that satisfies our goal, or we may want to keep searching until we find the best answer.

Knowledge Representation & Reasoning

- The second most important concept in AI
- If we are going to act rationally in our environment, then we must have some way of describing that environment and drawing inferences from that representation.
  - how do we describe what we know about the world ?
  - how do we describe it *concisely* ?
  - how do we describe it so that we can get hold of the right piece of knowledge when we need it ?
  - how do we generate new pieces of knowledge ?
  - how do we deal with *uncertain* knowledge ?

Learning

- If a system is going to act truly appropriately, then it must be able to change its actions in the light of experience:
  - how do we generate new facts from old ?
  - how do we generate new concepts ?
  - how do we learn to distinguish different situations in new environments ?

Periods in AI

- Early period - 1950's & 60's
  - Game playing
    - brute force (calculate your way out)
  - Theorem proving
    - symbol manipulation
  - Biological models
    - neural nets
- Symbolic application period - 70's
  - Early expert systems, use of knowledge
- Commercial period - 80's
  - boom in knowledge/ rule bases

Periods in AI (cont)

- The 90's and New Millenium
  - Real-world applications, modelling, better evidence, use of theory,…
- 2011 – now:
  - Big data, deep learning
- Applications:
  - visual recognition of traffic
  - medical diagnosis
  - directory enquiries
  - power plant control
  - automatic cars
  - …

Periods in AI (cont)

- The Promise of Deep Learning
  - In Neural Nets and Deep Learning, rather than hand-code a new algorithm for each problem, we design *architectures that can twist themselves* into a wide range of algorithms based on the data you feed them.
  - Deep Learning has yielded outstanding results on pattern recognition problems, such as recognizing objects in images, machine translation, and speech recognition → ***It is the Future***

AI Applications

- Autonomous Planning & Scheduling:
  - Autonomous rovers.

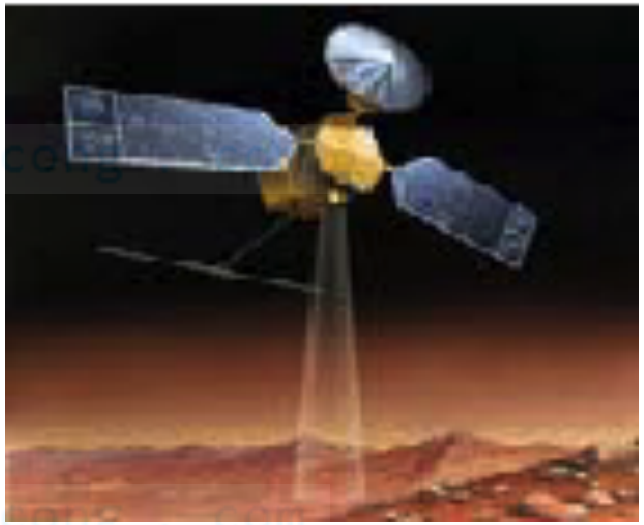

AI Applications

- Autonomous Planning & Scheduling:
  - Telescope scheduling

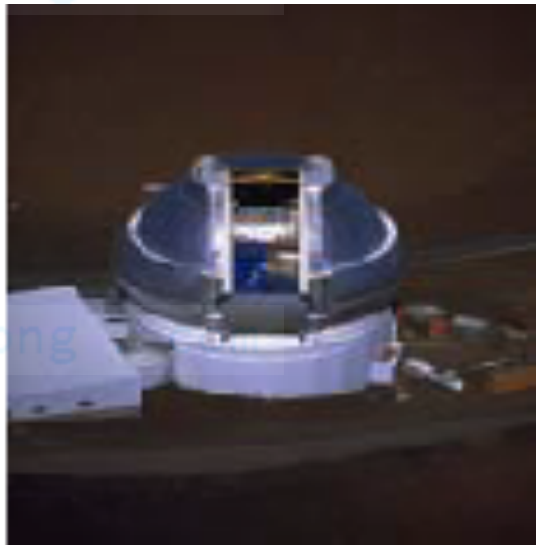

AI Applications

- Autonomous Planning & Scheduling:
  - Analysis of data:

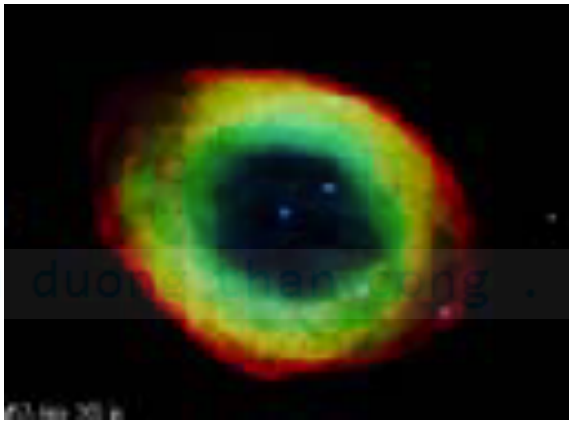

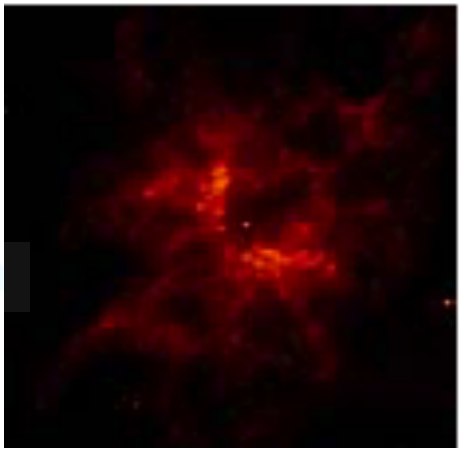

AI Applications

- Medicine:
  - Image guided surgery

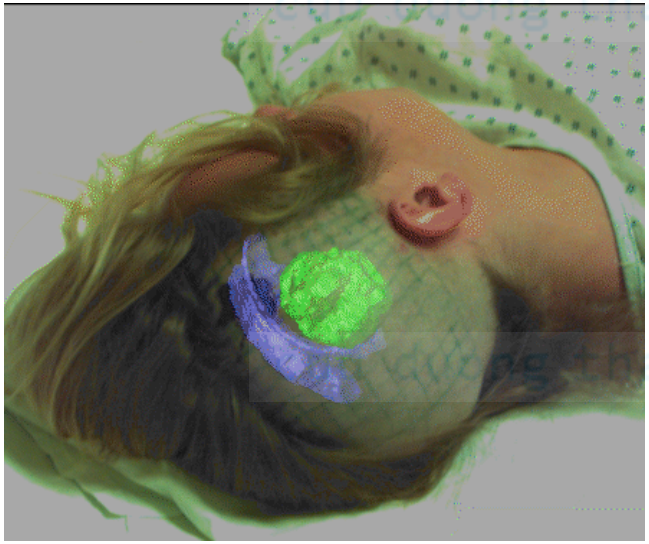

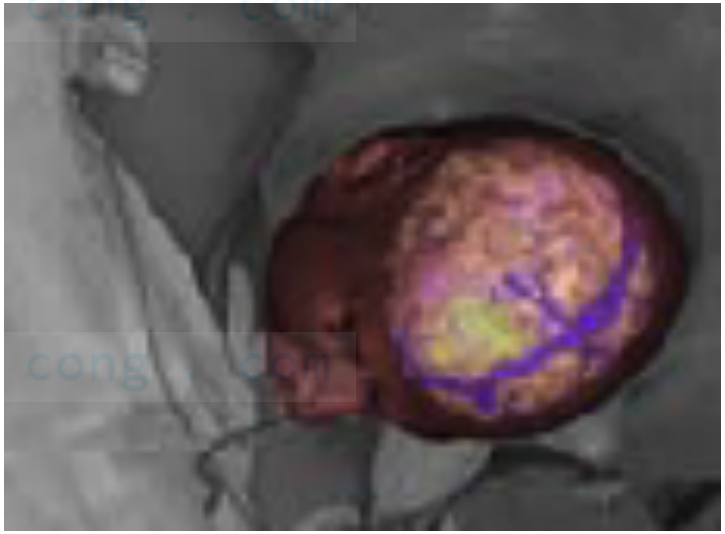

AI Applications

- Medicine:
  - Image analysis and enhancement

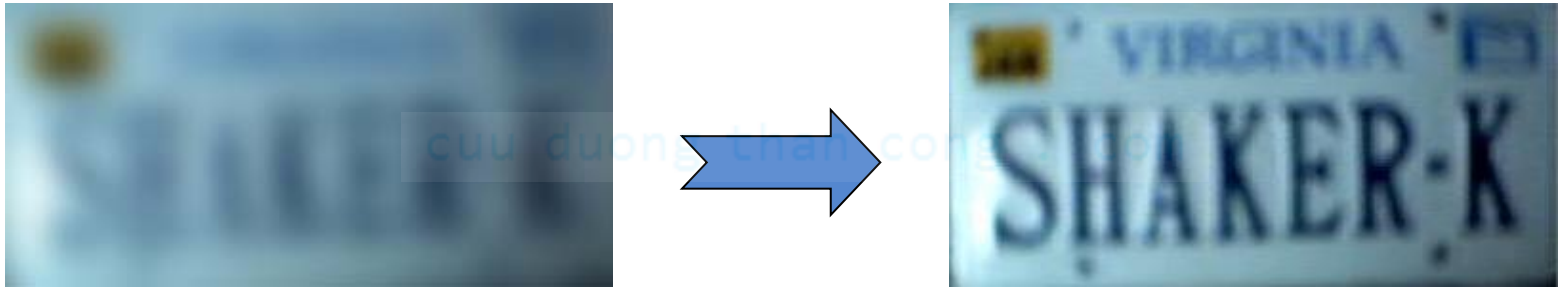

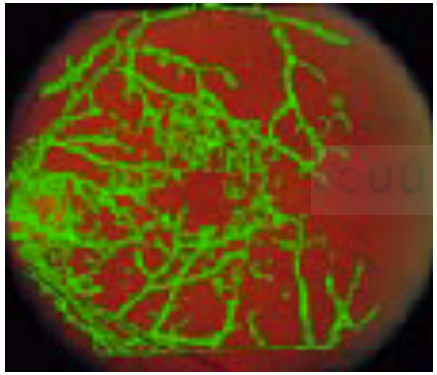

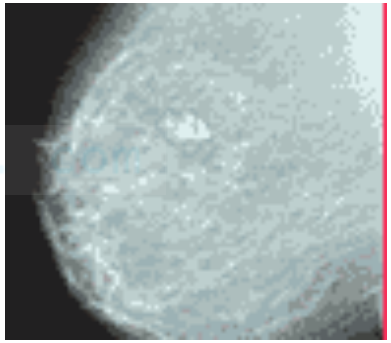

AI Applications

- Transportation
  - Autonomous vehicle control

Google Driverless Car

Toyota Prius modified to operate as a Google driverless car

AI Applications

- Transportation:
  - Pedestrian detection:

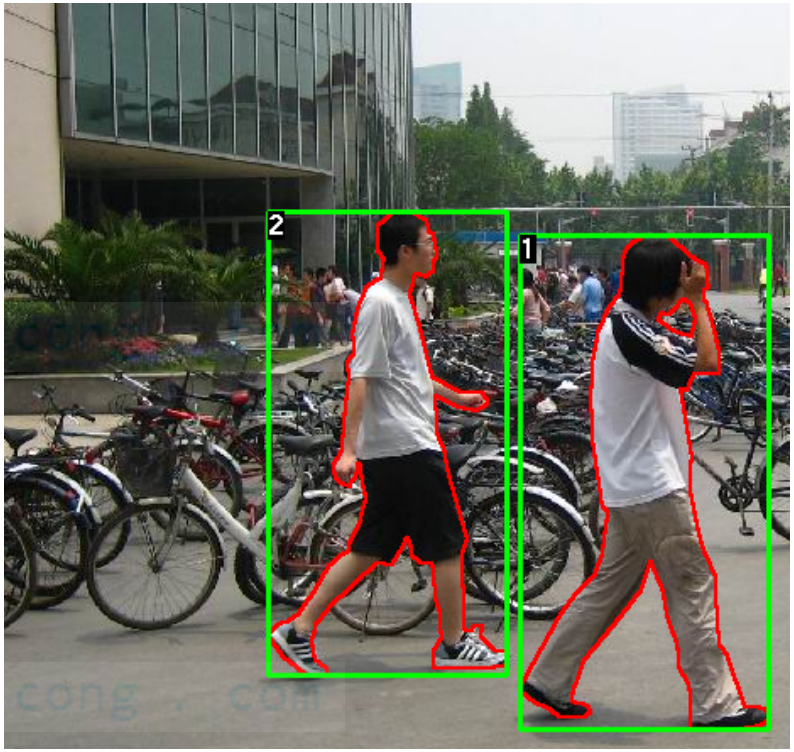

AI Applications

- Games

AI Applications

- Games

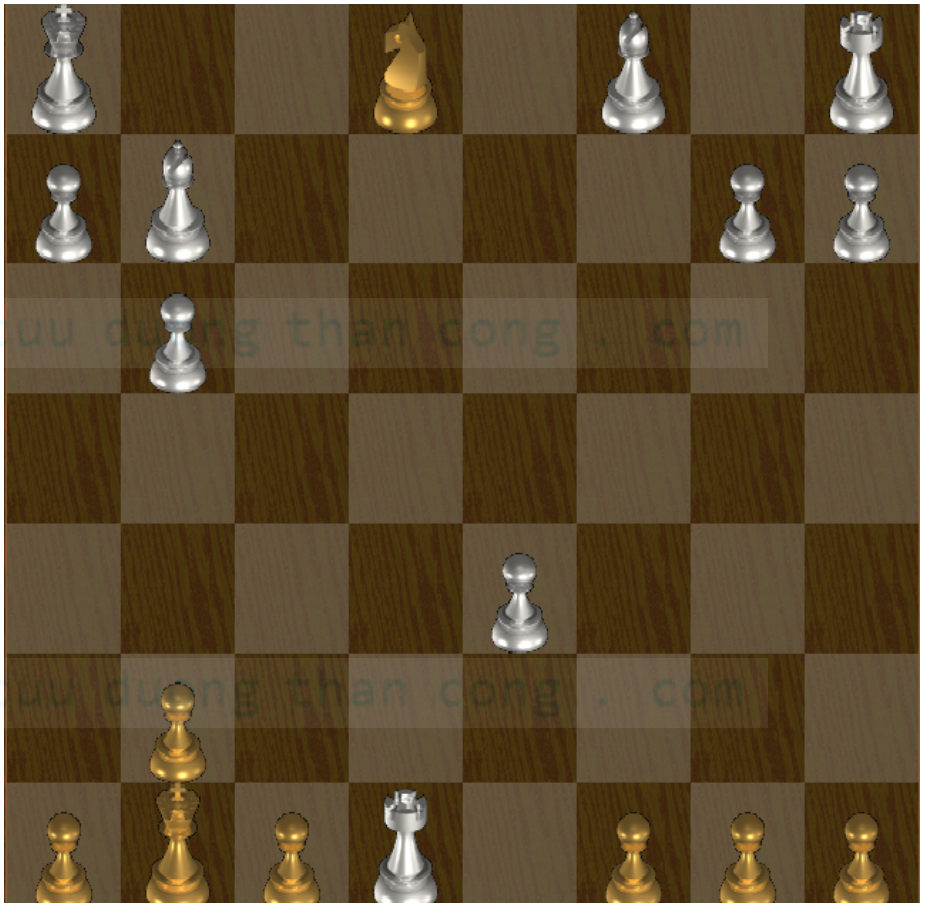

AI Applications

- Robotic:

AI Applications

- Other application areas:
  - Bioinformatics:
    - Gene expression data analysis
    - Prediction of protein structure
  - Text classification, document sorting:
    - Web pages, e-mails
    - Articles in the news
  - Video, image processing
  - Music composition, picture drawing
  - Natural Language Processing
    - Machine Translation
    - Semantic Analysis
  - etc,.

Group Discussion

Suppose you design a machine to pass the Turing test. What are the capabilities such a machine must have?

Homework

- Read Chapter 1 in your textbook (from page 1 to 29)