

Artificial Intelligence

One way to define **Artificial Intelligence (AI)** is as a branch of science trying to determine — and formally describe, permitting a computer implementation — the solutions for **hard problems**. Hard problems, in turn, are those which people solve — with more or less intellectual effort — but which solutions they cannot verbalize and describe in detail.

Is it not in any way a precise definition.

Is this a hard problem:



98731269868414316984251684351 × 985316846315968463198643541684?

And how about this: „Darling, please go and buy a nice piece of beef for a roast!“

And now a really hard problem: pour water from a glass to another container. More specifically: having a video camera (or two) connected to a computer, and a mechanical hand with fingers, write a computer program capable of lifting a glass of water from a table, and pour it into a container. Any kind of glass. From any given table. Into any container provided.

What is, and what isn't, the intelligence?

The concept of intelligence, or the lack of it, is often abused. The term **intelligent building** is used to describe a building equipped with a system for controlling the heating.



At the same time people often call **stupid** (taken as: lacking intelligence) a computer program which corrects spelling mistakes made by the (intelligent) human, when such program once in a while suggests a wrong correction.

What is the essence of (natural) intelligence?

We have fast and cheap computers, with large and perfectly reliable memories. They are highly accurate, make no mistakes (let's say), and never tire, maintaining their accuracy after many hours of work. What is the problem then, what is it in human intelligence, that the computers cannot capture?

Partially, the problem is just in this persistent and merciless accuracy!

People solve hard problems by using **abstraction** — multilevel problem analysis and the ability for non-schematic problem **decomposition**, ie. breaking it down into smaller parts. Their thinking is **flexible** — changing point of view and multidirectional reasoning. They are capable of efficient **pattern recognition**, **matching facts**, and **using analogies**.

Computers, in turn, have difficulties with recognizing specific situations, changing their way of reasoning, and adjusting it to changing situations. The existing algorithms for pattern recognition can be highly effective if they are highly specialized, but then they fail when the situation changes.

AI-enthusiasts and AI-skeptics


Artificial intelligence has its supporters and opponents. **The opponents claim that artificial intelligence cannot be built because the nature of (natural) intelligence is not computational, and is a unique feature of human mind.** The existing practical systems cannot be truly intelligent — by definition — since they are based in computer programs, and only perform number crunching.

Artificial intelligence is a moving target. When some tasks and goals set forth in the past have later been solved, the AI opponents simply concluded that apparently solving those particular tasks apparently did not involve intelligence, only their solutions were not previously known.

It would be useful to have an objective test, which could help verify, whether a specific system is intelligent or not.


The Turing Test (ca 1950)

The task is to construct a system enclosed in an isolated room, and connect to an independent operator with a communication terminal (teletype, the only kind existing in Turing's time). The operator is able to communicate with the system in natural language, ask questions, obtain answers, etc. Another terminal connects the operator with another room with a terminal and a human in front of it. If, after an unlimited time of probing both lines, the operator is unable to unequivocally state which line goes to the system and which to the human, then the system can be assumed to be intelligent.



Despite a long time passed since the test was proposed, it is still valid, ie. no system claims to have passed it satisfactorily.

Practical aspects of the Turing test



The practical application of the Turing test depends on the number of human participants: the interrogator, the human contestant, the referee, and the judges. They introduce the subjectivity into the test.

Turing test related contests

The Turing test is an abstraction and does not have strict rules. Nevertheless, attempts are made to implement and pass this test.



In 1990 Hugh Loebner funded a prize of \$100,000 and a gold medal for the first computer whose responses in the procedure constituting a version of the Turing Test, will be indistinguishable from human responses.

The competition rules define the content and scope of communication between the participants and judges. But ultimately it is the judges who decide whether the partner in conversation is a man or a machine. So the outcome may be the result of a mistake (poor intelligence?) of a judge.

For example, in another contest held in 2014 to mark the 60th anniversary of Turing's death, 33% of the judges thought that a Russian program Eugene Goostman impersonating a Ukrainian boy was human. The organizer of the contest announced that the Turing test has been passed, which was criticized.

In 2011 program Watson (IBM) won the \$1M prize over two human contestants in the T.V. game Jeopardy!, where the host gives a general knowledge clue, and the contestants must confirm understanding it by formulating a question.

Computer chess

Chess is a game of intelligence and has always been a challenge and a natural test field for artificial intelligence techniques. One of the first programmers of computer chess was Alan Turing, who could not run his program on a computer and executed it by hand simulation.

In 1957, Herb Simon, one of the pioneers of artificial intelligence, predicted that within 10 years, a digital computer would be the world's chess champion.

He was off significantly. After many years of effort on the construction of algorithms, programs, and specialized computers to play chess, it was only in 1997 that the chess computer Deep Blue first beat world champion Gary Kasparov in a single match. This did not mean, however, that the competition was finally won by the computers over the people of computers in chess. Over the next 10 years, a number of further programs competed with the best chess players with variable success. In 2006, the program Deep Fritz defeated the world champion Vladimir Kramnik in a tournament. Since then, the interest in the competition of programs with the people began to fall, which in some way signifies the end of this fight, with computers winning.

For a number of years, the game "go" remained the stronghold for humans, because it is so complex, that fast searching ability gave the computers no advantage over humans. Up until **2016**. In March 2016 the program AlphaGo developed by Google beat the Korean grandmaster Lee Sedol.

Other games

One can say that games are to artificial intelligence what grand prix racing is to the automotive industry. They have always been viewed as a challenge for researchers and programmers. When one game has been worked out — either theoretically, or by brute force search combined with clever techniques — the focus shifted to other games.

In checkers a program first defeated the world champion in 1994. Later, checkers have been solved theoretically. If both parties are playing their optimal strategy, the game ends in a draw.

In Othello the best programs dominate the people and the competition does not make sense. Conversely, in go (1000 BC), the number of possible moves is so large that a reasonable strategy must be based on logical analysis. Brute force power advantage disappears, and the best programs play at the amateur level.

An interesting result was achieved in Backgammon, where the program TDGammon (1992) reached the master level, thanks to the ability to learn, and not by the implementation of the best known strategies. Strategies discovered by the program were later adopted by humans.

Weak and strong artificial intelligence

With the prospect of constructing artificial intelligence, two levels of achieving this have been formulated.

- The **strong artificial intelligence** hypothesis claims that it is possible to construct a truly intelligent system, capable of thinking and having a mind.
- The **weak artificial intelligence** hypothesis leads to constructing systems which could act and solve problems as if they could think and have a mind.

This distinction has mainly philosophical and ethical significance.

The goals of AI

In practice, the goals of research and engineering work in artificial intelligence are:

→ working out a computational (algorithmic) theory of intelligence, human brain functioning, memory, emotions, instincts, etc.

In this sense, artificial intelligence is connected with biology, psychology, philosophy, as well as with mathematics and computer science, and other sciences.

→ building intelligent computer systems for effectively solving hard problems, and capable of functioning in the real physical world

In this sense artificial intelligence should collaborate, in addition to computer science, with robotics, mechanical science, mechatronics, and some engineering domains.

Tasks to be solved

Among other, artificial intelligence must face and solve the following tasks:

- **knowledge representation** to accept and express incoming information about the world, understand it, confront with already possessed knowledge, and use it efficiently
- **reasoning** to draw correct and necessary conclusions from acquired data, and make decisions about further actions
- **learning** to draw on existing knowledge and capabilities to adjust to new circumstances, unexpected by the system constructors, comprehend new phenomena, etc.
- **natural language understanding** is a practical requirement so that the AI system can communicate with humans
- **utilize visual data** from video cameras in order to independently acquire information about the environment
- **robotics** implies constructing practical systems capable of moving and executing tasks in the real world

Knowledge representation

The problem of knowledge representation is central to all areas and techniques of artificial intelligence. The problem is to choose, or create, a language for expressing facts, relations, dependencies, actions, their properties, meaning, consequences, and other information about the problem domain and environment which may be significant for its solving.

The problem is to use a **good** language for knowledge representation. **Using a proper representation language usually enables efficient reasoning, and makes it easier to find the simplest and/or optimal solution**, while using a wrong representation can make the problem much harder or impossible to solve.

A very simple, albeit distant, example might be the image representation for processing: vector or raster. Vector representation is good for object detection, while operations such as blurring work much better with raster formats.

A good knowledge representation is also very important for effective and efficient problem solving by humans. A good representation language permits humans a good understanding, eg. professionals from different domains collaborating to solve a problem.

Machine learning

An artificially intelligent agent is **learning if she improves her performance on future tasks** after making observations of her environment and previous achievements.

So we allow artificially intelligent agents who can learn, and others, who cannot. This can raise two sorts of doubts.

First, the human intelligence seems inherently possessing the ability to learn. The **human reasoning processes appear inseparable from the learning processes**. We would not consider intelligent a person who hasn't learned from her experience, at least in the simplest ways. So why is it separate in artificial intelligence?

There is no clear answer to this. Most developed and widely used artificial intelligence paradigms in their basic form perform reasoning without learning. **The ability to learn has to be added.**

Why machine learning?

Another doubt about machine learning might be: if it is not inherently obvious, or obligatory, then **why is it needed**, or is it really? If we can program reasoning processes, and are able to tune them to perfection, then perhaps we can obtain an artificially intelligent agent, who does not have to learn, or cannot learn anything more.

There is an answer to this question, and there are several good reasons.

Why machine learning? (2)

First, the designers of AI **cannot anticipate all possible situations** that the agent might find herself in. For example, a robot navigating a maze must learn the layout of each new maze she encounters.

Second, the designers **cannot anticipate all changes over time**. A program designed to predict tomorrow's stock market prices must learn to adapt when conditions change in an unpredictable way.

Third, sometimes human programmers have **no idea how to program a solution** themselves. For example, most people are good at recognizing the faces of their acquaintances, but even the best programmers are unable to program the computer to do this, except by using learning algorithms.

Applications — visual perception

- recognizing objects, symbols
- image segmentation
- 3D reconstruction
- image classification

Applications — natural language communication

Natural language processing technologies:

- text “understanding”, converting text to a formal representation
- machine translation
- information extraction
- question answering
- text classification, spam filtering, etc.

Speech processing technologies:

- automatic speech recognition (ASR)
- speech synthesis (text-to-speech, TTS)
- dialog systems

Applications — robotics

The robotics combines elements of mechanics and electronics (mechatronics), and artificial intelligence.


When we try to build robots and to deploy them for acting in the real world, we encounter problems far exceeding the developed theory.

Issues, existing technologies, applications:

- action planning
- controlling vehicles (walking, driving, flying)
- rescue systems
- social robots — care-taking for humans in need

AI history — the 50-ties of 20th century

- 19th century (and earlier) ideas: philosophy, logic, probability, human brain functioning research
- the 50-ties of 20th century: foundations of AI is connected with the emergence of computer science, the LISP programming language (McCarthy)
- puzzle and games, classical problems: "monkey and bananas", "cannibals and missionaries", and others
- early systems: GPS (Newell, Shaw, Simon), the checkers program (Samuel)
- theoretical models: the perceptron (Minsky)




AI history — the last quarter of 20th century

- practical applications, most importantly commercial
- theorem proving and symbolic computations
- natural language understanding, automatic translation of texts, speech understanding
- automatic programming: construction and verification of programs
- visual information analysis and autonomic robot (vehicle) control
- expert systems for many domains: medicine, geology, engineering design, economy, financing, etc.
- learning methods



AI history — the third quarter of 20th century

- appearance of formal methods based on mathematical logic
- connection with the development of robotics: perception methods, action planning, learning
- after an initial eruption of enthusiasm about the many promising methods that appeared and hopes for quick achievement of AI goals, there came understanding of computational complexity and combinatorial barrier
- inadequacy of classical logic: the need for approximate reasoning and making assumptions under lack of perfect information
- common sense-based reasoning



AI History — 21st century

This is a tentative summary, the 21st century isn't quite over yet ...

- Strong development of numerical methods, eg. for CSP problems, surprising spectacular result, solved some difficult issues in polynomial or even linear time, eg. the GSAT algorithm.
In this context, the algorithms amenable to parallelizing gain importance.
- Strong development of statistical methods, such as natural language processing based on corpora, and other applications.
- Development of methods based on probabilistic models, Markov processes, reinforcement learning, etc.
- The connections with economy (an intelligent agent must act rationally and economically), game theory, etc.
- Knowledge representation based on ontologies is revived and practical technologies are developed in the context of the Internet, the so-called *Semantic Web Initiative*.
- Agent approaches continue to be popular in different contexts.

- There are new fields of application, such as the social robotics.
- Further fields are separated from artificial intelligence and begin a life of their own, such as data mining.
- Artificial intelligence transfers to the practical life in different forms, from intelligent assistants for software packages, intelligent service systems, systems assisting in operating various system, eg. driving automobiles, to systems designed to function on the battlefield.
- More and more applications of artificial intelligence rise dilemmas of when and to what extent the decision processes can be transferred to intelligent machines.

Short quiz

1. What is a hard problem?
2. What is knowledge representation?
3. Define two main goals of AI.
4. What is the difference between strong and weak AI?
5. Can one claim the Turing test has been passed, at least partially?
6. Why is machine learning ability developed in addition to basic AI?