ARTIFICIAL INTELLIGENCE

In today's era of technological progress and the digital revolution, the term artificial intelligence is increasingly used. This is also because artificial intelligence has experienced great growth and progress in recent years. Artificial intelligence is credited with the potential to change our lives in all areas, from medicine to economics, through **intelligent systems that can quickly analyse vast amounts of data and draw relevant conclusions from it**. Today, artificial intelligence systems usually match the level of humans in solving various tasks, or even surpass them (e.g. image classification, face recognition, content generation, chess tournaments).

Artificial intelligence is often associated with an image of the future, where autonomous robots and intelligent systems take over activities that were once performed by humans. When understanding artificial intelligence, it is not necessary to have a limited view that it is purely generative artificial intelligence in the form of conversational chatbots (e.g. the currently popular ChatGPT), or even consider them as a synonym for artificial intelligence.

Artificial intelligence can have different forms of its implementation and wide use. Even today, we meet artificial intelligence daily, and we may not even realize it. We encounter it, for example, in the following cases:

- The camera in smartphones – improves the quality of photos, can automatically recognize the type of scene, or optimize the camera settings.
- Unlocking your smartphone through face recognition.
- Personalization of content on the web and social networks - based on the content viewed so far, websites show new content that we might like.
- Streaming services (Netflix, Spotify) - suggesting new songs/movies according to songs listened to and movies watched so far, preferred genres, or songs added to playlists.
- Designing the tagging of specific recognized persons on photos on social networks.
- Virtual assistants – Siri, Alexa, Google Home, Cortana.
- Spam filtering.
- Navigation – designing the optimal route considering various circumstances, including the traffic situation.

As we can see, artificial intelligence does not have to be strictly just robots that can move and interact with the real world. It is usually a form of computer program, so artificial intelligence systems can also be operated on a computer. From the point of view of their integration, artificial intelligence systems can be divided into:

- **Software-based artificial intelligence -** Usually found in the form of algorithms and software applications. These are, for example, voice assistants, image, or voice recognition. It can be deployed on servers in the cloud or on local computers or smartphones and is often available as a service for various industries, such as healthcare, the financial sector, e-commerce, or data analysis. Its advantages include the fact that it is flexible, easily updatable and has a wide range of applications. However, it requires a device with sufficient power and may have limited interaction with the physical world.
- **Embedded artificial intelligence in devices –** It is integrated directly into the hardware of devices, such as robots (in factories), drones, smart appliances, self-driving cars. These systems are designed so that they can directly process data and make decisions without the need to connect to external software or the cloud. This enables faster and more efficient real-time data processing. However, it is less flexible, more demanding to update and tends to be more narrowly focused on specific tasks.

Artificial intelligence is undoubtedly an interesting and one of the most dynamic fields in today's technological world. Its influence and range of possibilities are still at the beginning, so it is important that we deal with this topic with understanding and responsibility. Artificial intelligence is a powerful tool with enormous potential. Understanding its functioning and impact on society is key to being able to use it responsibly and shape our future for the better.

1.1 What is artificial intelligence?

Intelligence in general is the ability to adapt to new situations, gain experience, gain knowledge from the process of education (learning) and apply it to solving problems. It includes cognitive processes such as thinking, learning, perception, understanding and communication. Intelligence is manifested in various areas in humans and some animals.

Naturally, over time, the question arose as to whether it is possible for inanimate objects to manifest intelligence with the help of humans. In this case, we can talk about artificial intelligence. We can understand the term artificial intelligence:

- the ability of inanimate objects to exhibit signs of human-like intelligence,
- a scientific discipline that deals with the intelligent behaviour of machines - their simulation of human intelligence.

Artificial intelligence as a capability

Artificial intelligence is the ability of machines to perform activities that would typically require human intelligence and thus be the domain of humans. It is a simulation of human intelligence and thinking demonstrated by machines. The term artificial intelligence is thus used in the context when a machine imitates the cognitive abilities of a person (thinking, learning, solving problems). Thanks to this, it is possible to observe "intelligent" behaviour in machines. Artificial intelligence enables systems and machines to perceive and analyse their environment, as well as make decisions with a certain degree of autonomy to successfully achieve their specific goal. An important element is the autonomy of the system, which represents the ability of the artificial intelligence system to autonomously (independently) perform tasks in the environment without constant control by the user. With artificial intelligence systems, we can also talk about their adaptability - the ability to improve their performance through learning from experience.

Artificial intelligence as a scientific discipline

The term artificial intelligence also refers to the scientific discipline dealing with systems showing signs of intelligent behaviour. Artificial intelligence, as an interdisciplinary field of science, tries to investigate and understand intelligent systems. It also focuses on the creation of machines or systems that can imitate human cognitive functions and with the help of which it is possible to solve tasks with procedures that, if solved by a person, would be considered a manifestation of intelligent behaviour. The goal of artificial intelligence is to create intelligent machines that can function autonomously in a complex environment and perform tasks that would be difficult or impossible for humans.

1.2 History

Since **ancient times**, people have thought about creating an artificial being with intelligence. For centuries, the idea of a **man-made being that could think and act like a human** has existed. This is also evidenced by various legends and literary and artistic works:

- In science fiction literature and films, he often appears as an android character, that is, a humanoid robot.
- In Greek mythology, the sculptor Pygmalion is mentioned, who carves a statue of a woman, falls in love with her, and the goddess Aphrodite breathes life into the statue.
- Dr. Frankenstein creates a living being by combining different parts of human bodies.
- In the fairy tale written by Carlo Collodi, there is a wooden puppet Pinocchio who comes to life and longs to become a real boy.

Philosophy plays and plays an important role even before the emergence of artificial intelligence. Already in the **17th century**, several philosophers, such as René Descartes, Thomas Hobbes, or Blaise Pascal, dealt with questions about the human mind and thinking, what it means to be intelligent, how to make decisions, and people's behaviour. **They dealt with ideas on a philosophical level.**

In the 18th century, Wolfgang Kempelen, a native of Bratislava, came up with the idea of constructing a chess automaton, i.e. a machine that would act independently to achieve a set goal.

From the middle of the **19th century** and especially in the 20th century, in accordance with the significant development of science and technology, people began to try to construct machines that would be able to perform certain tasks instead of humans. In 1843, the Viennese professor Jozef Maximilián Petzval (born in Spišská Bela in 1807) introduced the term "**thinking machine**" as a thought abstraction of clever, intelligent machines.

In the **20th century**, the process of modern artificial intelligence began. During this period, the British mathematician and cryptanalyst Alan Turing studied various ways in which machines could be designed to demonstrate intelligent behaviour. His work contributed to the field of computer science and artificial life. In the years 1936 and 1947, **Alan Turing** dealt with the idea of a **universal Turing machine** – it was the concept of an abstract device that can simulate any other machine or computer system. Turing is also known for his Turing Test, which he proposed in the article "Computing Machinery and Intelligence" in 1950. This test is also called the Imitation Game. Its goal was to determine whether a computer exhibits intelligent behaviour comparable to that of a human. According to him, a machine can be characterized as thinking when its behaviour is indistinguishable from human behaviour. The Turing Test consists of an interaction between a human evaluator and a machine, where the evaluator does not know whether he is communicating with a computer or another human. The computer tries to act (communicate) like a person. If he can convince the evaluator that he is human, he is considered capable of displaying intelligent behaviour.

With the advancement in technology, the development of artificial intelligence could also advance. Neural networks, which are currently one of the most popular artificial intelligence algorithms, have their origins in the 1940s. In 1943, Warren McCulloch and Walter Pitts proposed a model of an **artificial neuron**.
In 1956, John McCarthy organized the conference The Dartmouth summer research project on artificial intelligence (also known as the Dartmount conference) for those interested in machine intelligence. At this conference, **the term "artificial intelligence**" was introduced for the first time.

In 1957, the American psychologist Frank Rosenblatt invented **the perceptron** - a simple model of a neural network, which had the task of modelling decision-making processes in the brain. He introduced the Perceptron in his publication "The Perceptron: A Perceiving and Recognizing Automaton".

One of the manifestations of an intelligent system is the ability to perceive and communicate with its surroundings. Between 1964 and 1966, Joseph Weizenbaum created **the first chatbot ELIZA** at the Massachusetts Institute of Technology (MIT). ELIZA was programmed to simulate a simple psychotherapy session. It answered the user's questions using simple predefined patterns and reactions. Even though ELIZA had no real intelligence or understanding, people were surprised at how authentic her interaction appeared.

In 1966, **the Shakey robot** was developed at Stanford University. Shakey was the first general-purpose mobile robot capable of making decisions about its actions. It was one of the first examples of a real-time planning and decision-making program that allowed a robot to navigate a dynamic environment and perform tasks as instructed.

The 1970s saw the spirit of **expert systems** that use knowledge and rules to solve complex problems and simulate human expertise.

In the eighties of the 20th century, attention was again directed to **neural networks.** Neural networks were inspired by biological neural networks. The research was directed towards better learning algorithms, the rediscovery of the backpropagation algorithm, which enabled neural networks to learn from data.

Due to its complexity, chess has been referred to as one of the problems where it is possible to compare the human intellect with that of an artificial machine. In 1997, for the first time, the computer **Deep Blue** defeated world chess champion Garry Kasparov in chess. Artificial intelligence has become more popular among the general population.

In 1998, Cynthia Brazeal at MIT developed the robot, **Kismet**. Kismet can simulate people's social interactions and emotions. Kismet was able to recognize and interpret non-verbal communication such as facial expressions, vocal intonations and respond to them in an appropriate manner.
In 1999, Sony introduced **AIBO** robotic home companions. These robots were modeled as interactive pets and had the ability to recognize faces, voice commands and learn from user interactions.

In the nineties of the 20th century, there was a rapid development in the field of personal computers, and at the beginning of **the 21st century**, a new communication era of the Internet began. **The Internet** connected many devices and the services provided by it (websites, social networks) began to produce a lot of data. Such progress in technique, technologies, computing, and memory capacity has once again enabled fundamental progress in the field of artificial intelligence.

In 2005-2006, IBM created the **Watson** communication system. Watson combines machine learning, natural language processing and big data analytics to understand human speech and produce answers to questions. He gained popularity in 2011 when he defeated Jeopardy! He beat two of the best human players of the time.

In 2011, when introducing the iPhone 4S smartphone, Apple also introduced the **Siri** voice assistant. Siri enables control and interaction with the device purely by voice. This represented a breakthrough in speech processing – Siri processed voice and not just written text. At the same time, it represents one of the first cases of mass expansion of artificial intelligence to users around the world.

In 2014, the **Eugene Goostman** chatbot became popular thanks to the fact that it managed to pass the Turing test - one third of the evaluators believed that it was human.

In 2014, Amazon launched an intelligent voice assistant called **Alexa**.

In 2017, the **AlphaGo** artificial intelligence system developed by DeepMind (a subsidiary of Google) defeated professional players of Go (one of the most demanding strategic games requiring understanding and complex decision-making).

In 2018, the artificial intelligence research and development organization OpenAI released the first version of the generative artificial intelligence chatbot **ChatGPT**. This version represented a significant advance in the field of generative models and machine learning, as it was able to generate fluent and authentic text based on input information and context.

1.3 Machine learning

By programming, people can create applications with different number of functions. However, these applications cannot do anything else that the programmer does not include in their code. As we stated in chapter 1.1, we can consider intelligent behaviour as the ability to learn and apply appropriate methods to solve problems in order to achieve a specific goal. Return et al. (2015) states that learning in this context can be understood as changes in the system that are adaptive in the sense that they allow the

system to perform the same task or tasks from the same class of tasks a second time more efficiently and effectively. Even people get better at solving many tasks just by doing them more than once (e.g. playing tennis, playing the piano). Machines can hardly be considered intelligent unless they can learn new knowledge and adapt to new situations. Fully pre-programmed software can perform various actions, but not in a way that we can consider its actions to be intelligent.

Artificial intelligence systems showing signs of intelligent behaviour work on the basis of machine learning, which does not require explicit programming. Machine learning makes it possible to improve the artificial intelligence of a system or machine through its experience, user interaction, and data collection and processing. This makes artificial intelligence smarter and can make better and more accurate decisions.

Training

At first, people thought that the most effective way to create artificial intelligence would be to go step by step and instruct it in detail in every aspect of its operation. For simple tasks, this method has proven to be effective. However, in the real world, the difficulty of solving some tasks can increase significantly, even for tasks that at first glance seem simple to us humans (recognizing whether there is a dog in the picture is a simple task for a human, but not for a computer). We also cannot fully prepare a small child for every situation that may arise. We give him instructions as much as we can, but learning also takes place independently through the child's experiences. When he does something well, we praise him, when he does something wrong, we abhor or punish him, which means that he must not do it again. When we want to teach a child to recognize various objects and elements in the environment around him, for example a dog, we point to a dog or dogs and say that it is a dog.

Machine learning works on a similar principle. First, we need to teach the artificial intelligence what we will demand from it - this process is called **training**. The training consists of a **training** and a **test phase**. **During training**, a mathematical model is created, which we train for a specific task on the training data set in the training phase. Computers are given access to a vast source of data from which machines learn to identify patterns and draw conclusions. It means that we don't tell the system the exact procedure, but it tries to find the procedure itself based on the patterns identified in the data. Using the example of dog recognition, we can state that the artificial intelligence system will be provided with several images of dogs, on which it will recognize their common features and, based on them, will subsequently be able to identify a dog that the system has never seen before.

The goal of machine learning is to enable computers to learn independently without the need for human assistance and to enable them to apply learned skills in practice. An

important component of this process is the quantity and quality of available data, as the success of an artificial intelligence system depends on the quality of the training data set. Therefore, ensuring sufficient and relevant data is key to success in machine learning. In recent years, the increase in computer performance and the availability of so-called big data (from English big data). It is the concept of aggregating huge amounts of data from the growing number of devices connected to the Internet and data that is generated through social networks. Training data can take the form of texts, images, videos, or audio recordings.

An artificial intelligence system is only as good as the data it learns from. If the training data is improperly prepared, there is little of it, or it does not reflect the diversity of the entire set, artificial intelligence can be biased. All deficiencies in the data will be taken over by the artificial intelligence system. So, if we were to train an artificial intelligence system to recognize dogs and cats, giving the system only photos of cats lying on the couch and only photos of dogs running on the grass, the system would be able to recognize a cat chasing a mouse on the grass as a dog. In the case of bias, we are therefore talking about an undesirable phenomenon. Therefore, it is important to have the right data collection techniques on which the artificial intelligence system will be trained.

The correct functionality of the trained model can then be tested during the **test phase on the test data set** (it represents data that the model has never seen before) and deployed in the production environment.

Types of artificial intelligence training

According to the way in which the learning of the artificial intelligence system takes place, we can divide the machine learning algorithms as follows:

- supervised machine learning,
- unsupervised machine learning,
- semi-supervised machine learning
- reinforcement learning.

Learning with a teacher

In this type of learning, the model is trained on a dataset consisting of **labelled data**, where each input is also associated with the correct output. The model is trained on this training data so that it can give the correct result for the given input. It is also called

supervised learning, because the model's learning process takes place on a training set of labelled data that had to be annotated - a description had to be assigned to them, for example by a human - who is therefore considered a teacher. The model then repeatedly predicts the outputs based on the inputs until the model reaches an acceptable level of performance – until it produces the correct outputs.

After training the model comes the testing phase. In the testing phase, the model is subsequently tested on a test set of data that no longer contains annotation. This is how the model's success is verified in the real world - for example, as a percentage of the model's match with our (not provided) annotation.

An example of this type of learning is a model that will recognize dogs and cats. First, it is necessary to supply the model with a sufficiently large set of images of dogs and cats, along with marking (with annotations) which images have a dog and which have a cat. Then the model is trained. After training, we show the model pictures of dogs and cats that it has never seen before (without annotations) and check whether it has correctly identified the animal. If not, we improve the dataset and train the model again.

Among the main problems that are solved with the help of this type of learning are:

- Classification – Used to assign input data to one of several categories. An example of a classification problem is assigning an email to the "spam" or "non-spam" class or assigning a diagnosis to a given patient.
- Regression – Focuses on predicting numerical values based on input data. For example, estimating real estate prices, predicting product demand, predicting student performance based on their academic results.

Learning without a teacher

Creating annotated data is a time-consuming and financially demanding process. If the model does not receive information about what the correct desired outcomes are, this is unsupervised learning. Compared to learning with a teacher, in this type of learning we only have a data set whose **data is not classified** (does not have assigned annotations), i.e. the correct answers are not known (the teacher is missing). Based on this data, the algorithm tries to recognize patterns and structures in the data. According to the recognized similarities, it then divides the inputs into groups with similar properties.
Among the problems that can be solved in this way of learning are:

- Clustering – combining data into groups according to their common properties, e.g. customer segmentation according to their purchasing behaviour.
- Association – finding association rules that describe a group of data and the relationships between them. They are used, for example, in the analysis of

customer purchasing behaviour to identify associations between different products - e.g. customers who buy a laptop also buy a mouse.

Partially supervised learning

This type of learning combines elements of tutored and untutored learning. The model is trained on a data set that contains a **lot of unlabelled data and only some labelled data**. Due to the time and financial complexity of data labelling, many real-world cases in the field of machine learning fall into the category of partially supervised learning. Untagged data is easier to collect and store. In this case of learning, it is possible to use unsupervised learning techniques to discover structures in the input data, as well as supervised learning techniques to estimate correct answers by repeatedly feeding the training set and running the training.

Reward learning

The names reinforcement learning, or reward and punishment learning are also used for this type of learning. In reward learning, neither labelled nor unlabelled training data sets are used. Learning takes place in such a way that an agent is created and deployed in the environment. The rules for how the agent can behave and the so-called reward function. The agent subsequently **learns in interaction with the environment from its own experiences** (trial/error method) and feedback from the environment. Although he will not receive information about the correct answer, he will receive an evaluation of his action - a reward for achieving the desired goals and a punishment for failure. Based on this feedback (information about the correctness or incorrectness of its output/action), the model evaluates the usefulness of the action it has taken and adjusts its behavioural strategy. The agent repeatedly performs actions in the environment, considering the feedback already received. The goal is to optimize the behavioural strategy to maximize the sum of the rewards.

An example of the use of such a learning method is training robots to navigate in the environment, or training systems to play games. For example, we will create an agent for playing chess, define the allowed moves and the rule for winning. We reward him if he knocks out the opponent's piece or if he wins. We will punish him if his piece was kicked or if he lost.

Deep learning

Currently, a very successful approach to machine learning is deep learning. Deep learning tries to get even closer to true artificial intelligence inspired by the human brain and connecting neurons in multiple interconnected layers, creating a deep neural network. By creating deep neural networks, deep learning tries to emulate the functions

of the inner layers of the human brain and create insights from multiple layers of input data processing. Deep neural networks compare new information with what they already know, just like the human brain. Each time deep neural networks are given new data, their capabilities improve.

Deep learning is used, for example, in face and speech recognition, machine translation, or text generation. Its advantage is high accuracy in various tasks, adaptation to new data and the ability to learn from a large amount of data. On the other hand, it requires high computational complexity. It can also be a problem with the interpretability of its results.

1.4 Artificial neural networks

A frequently used option for implementing machine learning is artificial neural networks. An artificial neural network is a specific type of algorithm. It is designed to mimic the structure of the human brain. The neural network in humans consists of a few neural cells - neurons. Each neuron can be connected to many other neurons through connections - synapses. Neurons interconnected in this way form a neural network. Neurons receive and send information to other neurons. Through this process, some connections can be strengthened/strengthened or even weakened. The stronger the connection between the selected neurons, the easier information passes through the given connection. Therefore, the more a person learns or trains something, the better/easier it becomes, basically automatically without conscious thinking about the activity. If a person learns something new, new neural connections are created and selected existing ones are also strengthened.

Artificial neural networks have also been inspired by these biological neural networks and their activity. The basic unit of artificial neural networks is an artificial neuron.

Perceptron

A perceptron is a basic model of an artificial neuron representing a simplified simulation of a biological neuron. It was invented in 1957 by the American psychologist Frank Rosenblatt as a simple type of neural network that had the task of modelling decision-making processes in the brain. A perceptron consists of only one neuron and represents the general computing element of all artificial neural networks. The functioning of a perceptron is thus like the functioning of a biological neuron.

Abstractly, we can describe the functioning of an artificial neuron as follows (Fig. 1.1). Like a biological neuron, a perceptron can have multiple inputs through which it receives impulses (**inputs**) from neighbouring neurons. Each input is associated with a

certain value and weight that determines how important that input is to the output of the perceptron - some information is more important than others for solving the task. The input values are then multiplied by their **weights** and added to a single sum - we get the sum of the weighted inputs. The weighted sum is processed by the activation function, which compares it with the neuron's excitation threshold (**threshold**). Based on the result of the comparison, the perceptron provides a certain **output** through its output terminals. We can represent the output mathematically as follows: $y = f(\sum x.\, w)$.

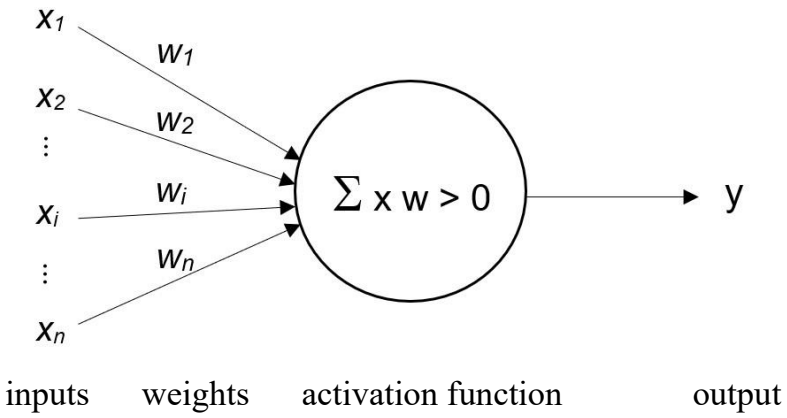

Fig. 1.1 Representation of a perceptron

If the resulting value of the perceptron is above the excitation threshold, the perceptron returns one value in the output (e.g. 1 – represents classification into one class), otherwise it returns another value (e.g. 0 – represents classification into the second class).

$$\text{output} = \begin{cases} 1 & ak\ \sum x\,.\,w > 0 \\ 0 & ak\ \sum x\,.\,w \leq 0 \end{cases}$$

In everyday life, it can be easily imagined as a situation where we are going to make a certain decision. We will collect all the pros and cons of the given decision. Each pros and cons can still have some importance. After considering all the pros and cons regarding their importance, we will either implement the given decision or not.

The perceptron learns by finding the weights that best solve the problem. For example, if we know the correct outputs for specific inputs (learning with a teacher), we run a calculation and verify that the perceptron has given the correct result. If the result was not correct, the weights will be adjusted according to their merit in the final output.

Multilayer artificial neural networks

Multilayer artificial neural networks are a type of artificial neural networks that consist of multiple layers of interconnected artificial neurons. These networks are much more complex than simple perceptron's and can solve much more complex problems in various areas, such as image processing (face recognition, object classification), natural language processing (machine translation, speech recognition), or making predictions (trend prediction, anomaly detection).

Multilayer neural networks consist of the following layers of artificial neurons (Fig 1.2):

- one **input layer** - its task is to receive input data and transfer it to the hidden layer,
- one or more **hidden layers** - they are placed between the input and output layers; they process the input data,
- one **output layer** – provides a result based on the processing of input data.

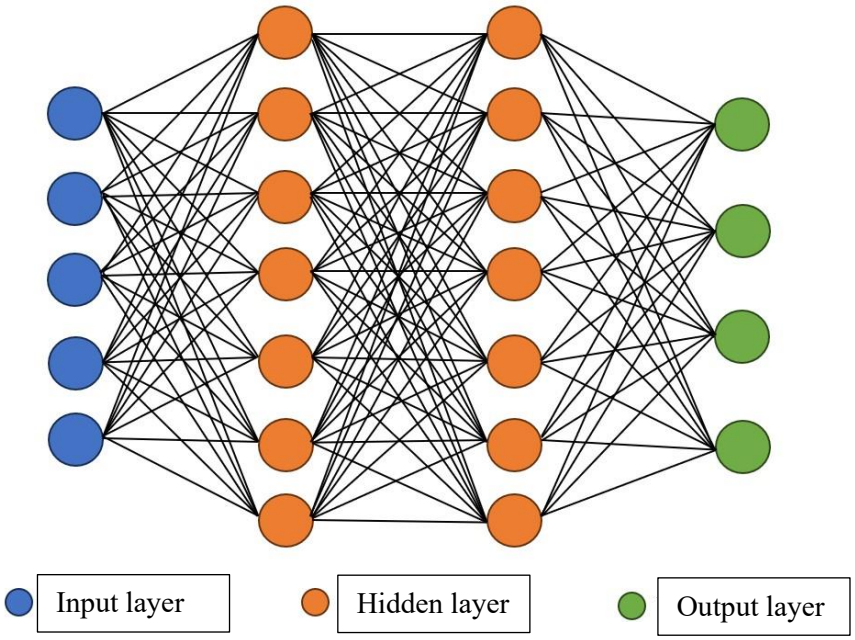

Fig 1.2 Illustration of a multi-layer neural network

Each artificial neuron has several weighted inputs along with outputs and threshold settings. How many layers are used, how many artificial neurons each layer contains, and how the layers are interconnected, we call the architecture of the network. If a neuron in a hidden layer is connected to all neurons in the previous and subsequent layers, we speak of a fully connected network (Fig. 1.2). Hidden layers make it possible to model more complex patterns and relationships in data.

Likewise, in this case of multilayer neural networks, learning takes place through training on a training data set, where results are verified for specific inputs and the

weights of neural connections are adjusted based on them. It is often difficult to know in advance what weight to assign to which characteristic. One way to figure this out is to start with a certain estimate of the weights and then let the algorithm adjust that estimate as it gains experience over time. Properties that appear to be good predictors of successful problem solving are given increased weights. Properties that, on the other hand, are not related to the successful solution of the problem, will have their weights reduced or will not be included in the evaluation at all. When training a neural network, it is an optimization process where the goal is to adjust the weights of the connections so that the neural network provides the desired input for the given output.

If a multilayer neural network contains many hidden layers, we are talking about **deep neural networks and deep machine learning**.

1.5 Types of artificial intelligence

Despite the great success and progress in development, the artificial intelligence systems developed so far are very specifically focused. It is possible to state that artificial intelligence systems can solve tasks as well or even better than a human, but each artificial intelligence system can only solve a specific task. A system that could beat a human at chess could not recognize dogs and cats in pictures.

From the point of view of the complexity of its ability to solve tasks, artificial intelligence is divided into:

- **artificial narrow intelligence**, ANI – it is also called as **weak artificial intelligence**,
- **artificial general intelligence**, AGI – it is also called as **strong artificial intelligence**,
- **artificial superintelligence**, ASI.

Artificial narrow intelligence

It is the most basic type of artificial intelligence. Every artificial intelligence system we encounter today can be classified into this type of artificial intelligence. This type of artificial intelligence is aimed at **solving only a specific task** (spam detection, face recognition, etc.). It is also called weak artificial intelligence, not because it is not powerful enough, but because it is still far from the real human one, which can solve any tasks from various fields.

Artificial general intelligence

So far, this concept is only being worked on at a hypothetical level. General artificial intelligence should be like human intelligence and should be able to **solve any intellectual task that a human can handle**. It should have such complex knowledge and cognitive abilities that its performance is indistinguishable from a human, even if its speed and data processing ability is much greater. The goal is to create a generally intelligent system/machine. However, there is no working example of general artificial intelligence yet.

Artificial superintelligence

A concept superior to general artificial intelligence. It describes artificial intelligence that can **surpass human cognitive abilities in a wide range of tasks**. Artificial superintelligence is thus far more intelligent than any human being in all areas of intellect and problem-solving ability.

1.6 Areas of artificial intelligence research

The great importance of artificial intelligence lies in the automation of people's mental work. From practice, we can observe that tasks that are difficult for humans are easy for machines - e.g. a computer can solve mathematical operations much faster than a human. On the contrary, many tasks that humans solve without consciously thinking about them represent a high level of difficulty for machines - e.g. grasping an object (chess piece), visual recognition of objects (whether it is a dog or a cat) and so on. This is also due to the fact that, at first glance, simple tasks, such as grasping an object with a person's hand, consist of a number of sub-tasks that are difficult for machines - it is necessary to look around in space and find the object, plan a route to it, approach to a reasonable distance , to move the hand by contracting different muscles in a specific time sequence and finally to grasp the object with just enough force so that it does not fall out of our hand or that we do not crush it.

Initial research in the field of artificial intelligence focused precisely on cognitive abilities and solving various intellectual tasks (such as playing chess). It was also because these tasks represented difficulty and the essence of intelligence for people. However, it turned out that computers can solve tasks of this type of very well thanks to their ability to perform many calculations within a second. Currently, areas such as natural language processing, computer vision and robotics are a challenge for artificial intelligence and researchers.

Natural language processing

Natural language processing (NLP) deals with the development of systems that can **understand and generate human language**. Therefore, the terms human language generation and speech recognition are also used in this context. It is the ability of a machine to understand human language in written or spoken form, identify words and sentences in it and convert them into a form that can be understood by a computer.

Until now, it was necessary to communicate with computers through programming languages that are precise, unambiguous, and structured. Natural human language is a challenge for computers due to its inaccuracy, use of synonyms, homonyms, different word forms, slang, dialects, etc. A modern approach to this issue is represented by the so-called **large language models** (LLMs), with which artificial intelligence tries to obtain the semantic and syntactic properties of words in sentences by repeatedly reading many human-written texts.

Natural language processing is already currently used in the automatic generation of subtitles on YouTube or in the development of voice assistants (Siri, Alexa, Google Now).

Computer vision

Computer vision deals with the development of systems that can **analyse and interpret digital images and videos to extract information about the world**. Thanks to this possibility, artificial intelligence systems can recognize objects in the real world. Given enough data, they can recognize pretty much anything.

Thanks to computer vision, we can already log in to smartphones via the face recognition function, or we can even identify objects that we don't even know (e.g. different types of plants, mushrooms) through the smartphone camera. It is currently also used in cars:

- parking assistance using sensors and cameras,
- adaptive cruise control can slow down or even stop the vehicle when approaching another braking car or an obstacle,
- traffic sign recognition function,
- lane keeping function - can turn the steering wheel slightly to keep the car from leaving the lane.

Robotics

Robotics deals with the development of robots that can interact with the physical world and perform complex tasks. The robot can perceive its surroundings and perform various tasks in this environment autonomously or with minimal human intervention. Robots are usually composed of a physical body, a control system ("brain"), sensors, and a

propulsion system. The advantages of using robots include that they can work non-stop, do not get tired, do not need to rest, reduce the risk of errors, are not influenced by emotions, can work faster and with greater precision, and can be deployed even in dangerous or harmful environments.

Robotics has been used in industry for years. But without artificial intelligence, it is necessary to manually program and calibrate the robots. With the use of artificial intelligence, it is possible to achieve the autonomy of these robots. Robotics is also used in healthcare in surgery and patient care, in agriculture for harvesting crops at the optimal time, milking animals and monitoring, in logistics in warehouses when working with goods, in space research, in the home during cleaning work or mowing the lawn.

1.7 Generative artificial intelligence

Generative artificial intelligence is currently a very popular and rapidly developing branch of artificial intelligence. Unlike other AI models that focus on recognizing and analysing existing data, generative AI focuses on **synthesizing and creating new data or content**. Generative artificial intelligence learns based on already existing data, from which it subsequently creates new data like the original. Such generative models thus allow computers to create authentic and realistic images, sounds, videos, and other types of content that a person often has no doubt is not real.

Depending on the type of content being created, the following generative models exist:

- **Models for text generation** can create texts such as articles, poems, professional texts and the like. They use the so-called large language models. These models are trained on a large amount of text data and can generate new text on different topics, written in different styles or in different languages.
- **Models for sound and music generation** can create new sounds, music, or sound effects. These models are trained on audio recordings and can generate sound in different melodies or in different rhythms.
- **Models for image generation** are used to generate new images, photographs, illustrations or works of art. They are trained on large datasets of images and can generate images in different styles, etc.
- **Video generation models** can generate moving images (videos) along with sound. With the use of these models, it is possible to create various visual effects, as well as realistic videos.

Generative artificial intelligence finds application in many fields such as copywriting, artistic creation, design, graphics, game development, film and audio industry and many others.

Generative artificial intelligence can help people be more creative and efficient at work. On the other hand, it can be misused to create false information and propaganda. Media content that is manipulated and usually created by artificial intelligence is called **deepfake**. The danger of deepfake content lies in its convincingly realistic appearance, in which artificial intelligence can create it.
A person can often have a problem distinguishing the content created in this way from reality and believing the fake content. A deepfake can usually take the form of a fake image, audio recording or video. In the case of a picture/photo, its content can be deleted, supplemented, or otherwise modified, for example, the background is changed, the person's face is replaced by another face (so-called face swap), etc. In deepfake audio recordings, the voice of a specific person can speak sentences that the person did not say in real life. Deepfake videos can also have their content modified in such a way that the person in the video can say sentences that they didn't say, while the movement of the person's lips is adapted to this. That makes the video even more realistic. Alternatively, even in the case of a video, the entire face of a person can be replaced by another face. Deepfake is usually associated with the risks it brings. However, deepfake can also have positive uses, for example in the film industry during dubbing. Currently, the actor was dubbed by another person (different voice) into another language. With the use of artificial intelligence and deepfake, the actor can speak in his original voice in a different language in the film.

1.8 Possibilities of using artificial intelligence

Artificial intelligence can be used in many areas of human activity. The possibilities of its application are wide and varied. We have already given several examples of the use of artificial intelligence in the previous subsections. Some examples of the use of artificial intelligence in practice in various spheres are presented in the following lines.

Healthcare:

- analysing X-rays and other images and identifying diseases,
- analysing the patient's health data and predicting the risk of disease,
- surgical robots can perform complex operations,
- tailoring the treatment of patients based on their individual characteristics.

Industry:

- product quality management using camera imaging,
- prediction of machine failures based on data from their sensors and maintenance planning,
- monitoring and management of production lines to increase their efficiency and minimize errors.

Trade and marketing:

- recommending the purchase of a product based on the previous purchasing behaviour of customers,
- automation and personalization of marketing campaigns,
- demand forecasting,
- chatbots for customer support.

Transportation:

- designing the most optimal route considering the traffic or the condition of the vehicle,
- autonomous self-driving vehicles,
- optimization of transport in transport networks and prevention of traffic jams.

Agriculture:

- monitoring and analysing soil data to deliver the exact amount of nutrients and water to plants,
- monitoring fields via drones to suggest the correct application of fertilizers and pesticides only where needed.

Science and research:

- help scientists analyse large amounts of data and perform complex simulations.

Education:

- optimization of curricula based on individual needs and results.

We present only a selection of several possibilities of using artificial intelligence. With the continued development of artificial intelligence, it is possible to expect an increasing scope and innovativeness of the possibilities of its use.

1.9 Challenges and ethics of artificial intelligence

Thanks to artificial intelligence, modern technologies are changing from hitherto passive objects that perform something only based on human commands, to actors who can make decisions and act according to their decisions. Artificial intelligence can be used in a variety of ways, with a positive or negative impact. We can make an analogy with a hammer, which is a useful and beneficial tool, but it can also be used as a weapon. Artificial intelligence also brings with it several ethical questions and challenges that need to be solved. The task of the ethics of artificial intelligence is to think about the impact of the actions of artificial intelligence and **to deal with questions about its**

correct behaviour, which should be in accordance with the moral, ethical, and other principles of our society.

There are several challenges and questions that need to be considered in the context of artificial intelligence and ethics. Among them is:

- **Legal liability** – For the actions and potential errors that machines and artificial intelligence systems make through their actions. Who is responsible for the actions taken by the artificial intelligence?
- **Copyright** – Who owns the copyright of content created by generative artificial intelligence?
- **Data protection and privacy** – Artificial intelligence systems use a lot of data, which can also be of a personal nature, for example, a patient's medical records or a person's financial situation. It is important that this data is sufficiently protected, that the privacy of individuals is respected, and that this data is not misused.
- **Minimization of damage or injury to health** – For example, in the case of autonomous vehicles, when it will be exposed to a situation where an accident can no longer be avoided and it must decide whether to endanger the pedestrian in front by not braking, or to swerve to the left and endanger its crew and the crew of the oncoming vehicle with a head-on collision, or to turn right and endanger pedestrians on the sidewalk. How should he decide?
- **Transparency and explainability** – Artificial intelligence and machine learning algorithms are often too complex and can be thought of as a "black box" that we cannot see into. However, there may be situations when it is important to justify how the given decision was reached, for example in the field of aviation or medicine, when human lives are at stake.
- **Bias, prejudice and discrimination –** In the case of insufficient quality and objectivity of training data, for example, discrimination against a certain group of residents may occur.
- **Deepfake –** Dishonest use for spreading misinformation and defaming other persons.
- **Employment of people –** Advances in artificial intelligence and the associated automation of activities that were performed by humans may lead to the demise of some professions and the loss of jobs. Another point of view is that it is a matter of relieving people from activities that do not need to be explicitly performed by a person. This can increase the productivity and efficiency of problem solving. Like the invention of the steam engine, the tractor, or the computer.
- **Technological singularity** – General artificial intelligence and artificial superintelligence have the potential to become so intelligent that they would

surpass the intelligence of humans. At the same time, these systems would be able to optimize themselves, update, improve and increase their intelligence more and more. In such a case, we speak of a technological singularity. Such machines could be difficult to predict and control. It is therefore important to ensure that they are safe, that they do not pose a threat to humans, and that humans remain able to control and manage them.

There are various challenges associated with artificial intelligence related to ethics, morality, security, and other areas. There are no easy answers to these challenges and questions. Their solution is an important task for the whole society.