Politechnika Wrocławska

Artificial Intelligence

Information Technology

Wojciech Myszka

Department of Mechanics, Materials Engineering and Biomedical Engineering

September/October 2025

unite! University Network for Innovation, Technology and Engineering

HR EXCELLENCE IN RESEARCH

Introduction

1. There is no real need to talk about it — it is on everyone's lips
2. Widely (???) used, usually without deeper reflection
3. Inseparably linked with computers (which was even reflected in the early names of the computer):
   - *mózg elektronowy* (electronic brain)
   - and in English *electronic brain* or sometimes *giant brain*

A Little History I

1. Golem *a being formed from clay in the shape of a man, but devoid of a rational soul, neshama, and therefore also the ability to speak.*

   It appears in the Old Testament, inseparably linked with Jewish folklore and the Talmud.
2. 1738 Jaques Vaucanson presents three automatons: the *Drummer*, the *Flute Player*, and the *Digesting Duck*.
3. 1769 **Wolfgang von Kempelen** presented the *Mechanical Turk* — a machine that was supposed to play chess (later automatons like *Ajeeb* or *Mephisto* worked on a similar principle).
4. 1818 **Mary Shelley** publishes the novel *Frankenstein; or, The Modern Prometheus*.

A Little History II

5. 1843 **Ada Lovelace** translates the article "Sketch of the Analytical Engine" invented by **Charles Babbage**. In notes to the translation, she develops what is considered the **first algorithm** that could enable Babbage's machine to calculate a complete sequence of Bernoulli numbers.
6. 1912 **El Ajedrecista** an automaton capable of playing three-piece endgames with a human (it worked algorithmically)
7. 1913 **Bertrand Russell** and **Alfred North Whitehead** publish *Principia Mathematica*, a groundbreaking work on formal logic and reasoning.
8. 1923 The term "robot" appears in **Karel Čapek**'s play R.U.R.

A Little History III

9. (1940s) **Isaac Asimov** publishes short stories later released as the collection **I, Robot**. The *Laws of Robotics* are formulated there.
10. 1943 the first article by Arturo Rosenblueth, Norbert Wiener, and Julian Bigelow *Behavior, Purpose and Teleology*, which initiates the field of **cybernetics**.
11. 1948–1949 **William Grey Walter** builds an autonomous, tortoise-shaped robot.
12. (1940–1950)

    Alan Turing works at Bletchley Park on automating the process of decoding ciphers used by the Germans.

A Little History IV

In 1950, Alan Turing delivers the lecture "Intelligent Machinery, A Heretical Theory" on the potential problems of automating rational processes. Alan Turing publishes the paper "Computing Machinery and Intelligence" and introduces the Turing Test.

13. (1950s) **John McCarthy** develops the **LISP** language for the IBM 704 computer. In 1960, the IBM 700–7000 series is produced, key systems in computer science and early artificial intelligence programming.
14. 1956

    Claude Shannon develops Theseus, a cybernetic mouse that solves a maze. Between 1948 and 1958, Claude Shannon presents a paper on information theory.

A Little History V

In 1956, the Dartmouth Conference, organized by John McCarthy, Marvin Minsky, and Claude Shannon, marked the birth of artificial intelligence as a separate scientific discipline.

15. 1957 Frank Rosenblatt, based on the idea of a neural network (developed in 1943), designs a network (Perceptron) that would be able to learn and simulates its operation on a computer. A physical device is created later...
16. 1959 **John McCarthy** and **Marvin Minsky** found the Artificial Intelligence Project within the Research Laboratory of Electronics and the Computation Center at MIT
17. 1960 **James L. Adam** builds the *Stanford Cart*, a prototype of an autonomous car.

A Little History VI

18. 1961 Unimate, the first industrial robot, is unveiled at General Motors.
19. 1966 Shakey, the first general-purpose robot, is demonstrated at the Stanford University Artificial Intelligence Laboratory. While conducting research on Shakey, Helen Chan Wolf makes breakthrough discoveries in cartography and image recognition, pioneering the field of facial recognition.
20. 1966 the ELIZA program, a chatbot simulating a psychoanalyst written by **Joseph Weizenbaum**, is created
21. 1968 **Stanisław Lem** in the SF short story *The Trial* describes a situation involving the use of a humanoid robot equipped with artificial intelligence as a spaceship pilot

A Little History VII

22. 1969 **Stanley Kubrick's** film *2001: A Space Odyssey* features HAL 9000, a malevolent artificial intelligence. Marvin Minsky was a consultant on this film.
23. 1971 **Nolan Bushnell** and **Ted Dabney** create *Computer Space*, the first arcade game utilizing artificial intelligence.
24. 1971 **Stanisław Lem** in the SF short story *Ananke* foresees the phenomenon known today as *bias* related to the training of artificial intelligence (which causes a disaster of a rocket landing on Mars)

A Little History VIII

25. 1981

    13% of UK households own a home computer. By 2017, this percentage would soar to 88%.

    IBM releases its first personal computer, the IBM PC

26. 1982 **Ridley Scott** adapts Philip K. Dick's novel "Do Androids Dream of Electric Sheep?" (1968) into the film "Blade Runner" starring Harrison Ford. The book reflects the increasing influence of machines and information on the concept of identity and humanity.

A Little History IX

27. 1989 English scientist **Tim Berners-Lee** invented the World Wide Web in 1989. The launch of the World Wide Web significantly increased the amount of data available for machine learning research.
28. 1990 **Rodney Brooks** publishes the paper “Elephants Don’t Play Chess” and becomes the director of the MIT AI LAB. Brooks revitalizes research into neutral networks.

    First successful speech recognition tests at the Stanford Research Institute.
29. 1999 **The Wachowski Brothers**/Wachowski Sisters directed the film *The Matrix*, depicting a struggle for power between machines and humanity.

Politechnika Wrocławska

A Little History X

30. 1999 Sony launches *Aibo*, the first commercially available robot dog.
31. 2000 **Cynthia Breazeal** develops *Kismet*, a program and robot that can recognize human emotions through facial analysis.
32. 2010 **Hiroshi Ishiguro** and the Intelligent Robots Lab introduce the android Geminoid.
33. 2015
    - The Future of Life Institute publishes the open letter "Research Priorities for Robust and Beneficial Artificial Intelligence" signed by 8,000 people, including Steve Wozniak, Stephen Hawking, and Elon Musk.

A Little History XI

- The Future of Life Institute publishes the document “Autonomous Weapons: An Open Letter from AI and Robotics Researchers,” endorsed by nearly 4,000 AI and robotics researchers.

34. 2015 ERIKA, an android news anchor, is presented by **Hiroshi Ishiguro** and the Intelligent Robots Lab at Osaka University.
35. 2016 Google DeepMind’s AlphaGo program defeats Go world champion Lee Sedol in the game of Go.
36. 2018
    - The European Union establishes ELLIS, the European Laboratory for Learning and Intelligent Systems.
    - Alibaba’s language processing artificial intelligence outperformed a human in a reading comprehension test at Stanford University.

What Constitutes Artificial Intelligence?

Politechnika Wrocławska

1. Machine Learning (ML)
    1.1 Supervised Learning
    1.2 Unsupervised Learning
    1.3 Reinforcement Learning
2. Generative AI
3. Natural Language Processing (NLP)
4. Expert Systems

What is Supervised Machine Learning? I

Supervised Learning is a machine learning method where the algorithm is trained on **labeled data**. This means that every **input data point** (e.g., image, text snippet, customer record) is assigned a **correct answer** (a label or output value) that acts as a “teacher” supervising the learning process.

How Does It Work?

This process involves the model learning to **map a function** from input data to output data.

1. **Labeled Data:** The algorithm receives a large dataset where each sample is a pair: (input data, expected result).
   - *Example:* (photo of a cat, label "cat"), (features of an apartment, sale price).
2. **Training:** The model analyzes these pairs and iteratively adjusts its internal parameters (weights) to minimize the difference (error) between its **predictions** and the **correct labels** (expected results). This is continuous correction, similar to a student checking their answers against a key.
3. **Generalization:** The goal is to teach the model **general patterns** and dependencies, not just memorize the training data. After training, the model is tested on **new, unseen data** to check if it can correctly predict the outcome for them.

Main Types of Tasks

Supervised learning is used to solve two main types of problems:

Task Type	Description	Example
Classifica	Predicting a **discrete (categorical) label** or class.	Determining whether an email is **SPAM** or **NOT SPAM** (binary classification), or recognizing whether a photo shows a **dog**, **cat**, or **rabbit** (multi-class classification).
Regressi	Predicting a **continuous numerical value**.	Predicting the **price of a house** based on its size and location, or forecasting tomorrow's **temperature**

What is Unsupervised Machine Learning?

Unsupervised Learning is a method where the algorithm learns from **unlabeled data**, i.e., data without predetermined correct answers. Its goal is to independently **discover hidden patterns, structures, and relationships** in the dataset to organize and understand it.

This is like a situation where you give an algorithm a pile of photos (e.g., animals, fruits, buildings) and ask it to sort them, but you don't tell it what they depict. The algorithm must find similarities and differences itself to create natural groups.

How Does Unsupervised Learning Work?

In contrast to supervised learning, where the model corrects its errors based on known labels, there is no “teacher” in unsupervised learning. The model explores the data itself:

1. **No Labels:** The algorithm receives raw input data (e.g., customer demographic data, measurement results). There is no column telling which group this customer belongs to or what these measurements mean.
2. **Pattern Discovery:** The algorithm uses mathematical and statistical techniques to identify similarities and differences between data points. It looks for natural clusters or ways to simplify the data.
3. **Creating Structure:** The model does not predict values, but rather **transforms the data** or **organizes it** into logical structures.

Main Tasks I

Unsupervised learning is used to solve exploratory problems:

1. Clustering (Grouping)

   Involves grouping similar data points into **clusters** (groups).

   - **Goal:** Finding natural divisions in the data that were previously unknown.
   - **Example: Customer segmentation**. The algorithm analyzes customer purchasing behaviors and automatically groups them into segments (e.g., “thrifty bargain hunters,” “premium customers”), allowing the company to better target advertisements. The most popular algorithms are **K-Means** and **Hierarchical Clustering**.

Main Tasks II

2. Dimensionality Reduction

   Involves reducing the number of features (variables) in a dataset while retaining as much essential information as possible.

   - **Goal:** Simplifying complex data, reducing noise, and speeding up computations.
   - **Example:** Image analysis. A digital image can have thousands of pixels (dimensions). Dimensionality reduction allows identifying only the key pixel features important for object recognition, while eliminating unnecessary noise. The most commonly used algorithm is **PCA** (Principal Component Analysis).

Main Tasks III

3. Association Rules

   Involve finding dependencies and rules between variables in large transactional datasets.

   - **Goal:** Discovering which items are frequently purchased together.
   - **Example: Market Basket Analysis.** Discovering a rule like: "If a customer buys bread and butter, there is an 80% probability they will also buy cheese." This allows for optimizing the arrangement of goods in a store or creating personalized online recommendations.

What is Reinforcement Learning? I

Reinforcement Learning (RL) is a machine learning paradigm where an **agent** (computer program) learns to make optimal **decisions** through **interaction** with an **environment** to **maximize the cumulative reward** over the long term.

This process is modeled on behavioral psychology and trial-and-error learning, similar to training an animal or learning to play a video game.

The Basic Operating Principle (The Loop)

RL involves a continuous interaction between the agent and the environment:

1. **State Observation:** The agent observes the current state of the environment (e.g., position in a game, sensor data from a robot).

What is Reinforcement Learning? II

2. **Action Taking:** The agent decides what action to perform based on its current **policy** (strategy).
3. **New State and Reward:** The environment reacts to the action, transitioning to a new state and returning a **reward signal** (positive) or **penalty** (negative).
4. **Learning:** The agent uses the received reward to **improve its policy**, learning which actions in a given state lead to higher future rewards.

The agent's goal is not just to gain an immediate reward, but to **maximize the sum of rewards over the long term** (cumulative reward).

Key Elements I

A reinforcement learning system consists of the following elements:

- **Agent** (Learner/Decision-Maker): This is the algorithm or model that learns and makes decisions.
- **Environment:** The world in which the agent operates (e.g., a computer game, stock market simulation, physical world for a robot). The environment defines possible states and reaction rules.
- **State:** The current situation in the environment that the agent observes.
- **Action:** The choice the agent makes in a given state.
- **Reward:** A number that is immediate feedback from the environment, indicating how good the last action was.

Key Elements II

- **Policy:** The agent's strategy; a function that maps states to actions (determines what action the agent should take in a given state).
- **Value Function:** An estimate of the predicted **cumulative reward** that the agent can obtain starting from a given state and following the policy. It helps the agent evaluate the long-term profitability of states, not just the immediate reward.

Exploration vs. Exploitation

A key challenge in RL is finding a balance between:

- **Exploitation:** Choosing the action the agent believes to be the best (most rewarding) based on current knowledge.
- **Exploration:** Choosing a new, unknown action to gather additional information about the environment, which may lead to the discovery of a better strategy.

The agent must explore enough to find the optimal path, but also utilize the knowledge gained (exploit) to maximize the reward.

Examples of Applications

Reinforcement learning is ideal for problems that require **sequential decision-making**:

- **Computer Games:** AlphaGo and AlphaZero agents playing Go and chess.
- **Robotics:** Teaching robots grasping, walking, or navigation in complex environments.
- **Autonomous Vehicles:** Making decisions about accelerating, braking, and turning in dynamic traffic.
- **Resource Management:** Optimizing energy consumption in data centers (e.g., DeepMind used RL to optimize cooling in Google).
- **Finance:** Creating trading algorithms on the stock market.

Fundamentals of Generative AI Operation I

Generative Artificial Intelligence is a type of AI capable of creating new content such as text, images, music, and even code. Unlike traditional AI systems designed for specific tasks (e.g., image recognition or language translation), Generative AI can "understand" and mimic the patterns of the data it was trained on to then generate new, unique, and realistic outputs.

The fundamentals of its operation:

1. **Training Data:** Generative AI requires vast amounts of training data. For instance, if it is to create images, it needs thousands or even millions of photos. If it is to write texts, it needs an extensive corpus of texts. This data teaches the model "what things look like" or "what things sound like" that it is supposed to generate.

Fundamentals of Generative AI Operation II

Politechnika Wrocławska

2. **Model Architecture:** There are many different architectures for generative AI models, but the two most popular are:
   - **Generative Adversarial Networks (GANs):** Consist of two cooperating (and competing) neural networks:
     - **Generator:** Its task is to create new data (e.g., images) from input noise. It tries to create data as realistic as possible to fool the discriminator.
     - **Discriminator:** Its task is to distinguish between real training data and data generated by the generator. The better the generator fools the discriminator, the better the generated data becomes.

     This competition process continues until the generator can create data that the discriminator cannot distinguish from real data.

Fundamentals of Generative AI Operation III

- **Transformers:** Although initially designed for natural language processing (NLP) tasks, they have proven extremely effective in generating text, and have also been adapted for other modalities, including images (e.g., in DALL-E, GPT-3/4 models). A key element of Transformers is the *attention mechanism*, which allows the model to weigh the importance of different parts of the input data when generating the output. Language models based on Transformers learn to predict the next word in a sequence, allowing them to create coherent and contextual texts.

Fundamentals of Generative AI Operation IV

3. **Learning Process:** During training, the generative model learns hidden patterns and dependencies in the data. It does not "understand" in the human sense, but creates a statistical representation of the data. For example, a model generating cat images will learn what visual features (fur, ears, eyes) are typical for cats and how they are arranged relative to each other.
4. **Generating New Content:** Once trained, the model can be used to generate new, previously non-existent content. In the case of GANs, the generator receives random "noise" as input and transforms it into new data. In the case of language models, the user provides a "prompt," and the model continues the text by predicting subsequent words.

Key Features of Generative AI:

Fundamentals of Generative AI Operation V

- **Creativity:** Capable of creating unique and original content.
- **Realism:** Aims to create content that is difficult to distinguish from real content.
- **Multimodality:** The ability to generate different types of content (text, images, sound).
- **Controlled Generation:** Many modern models allow the user to influence the generation process by providing textual descriptions, styles, or other parameters.

Generative AI has broad applications, from assisting artists and designers, through automatic generation of marketing content, to creating realistic simulations and supporting innovation in many fields.

What is Natural Language Processing?

Natural Language Processing (NLP) is an interdisciplinary field at the intersection of **Artificial Intelligence (AI)**, **computer science**, and **linguistics**, which enables computers to **understand, interpret, and generate human natural language** (speech and writing).

The goal of NLP is to bridge the communication gap between humans and machines, giving computers the ability to analyze language logically and contextually, just as humans do.

How Does NLP Work?

NLP uses advanced algorithms, machine learning (including deep neural networks), and linguistic rules to break down language into machine-understandable elements:

1. **Syntactic Analysis:** Focuses on sentence structure, grammar, and rules for constructing phrases. This helps the computer understand how words are connected to each other.
   - *Example:* Recognizing which word is the subject and which is the predicate.
2. **Semantic Analysis (Meaning):** Focuses on the meaning of words, phrases, and the entire text, often in context.
   - *Example:* Distinguishing whether the word "zamek" (castle/lock) in a sentence means a building or a mechanism in a door.
3. **Pragmatic Analysis:** Understanding how external context affects the meaning and intention of the message (e.g., irony, purpose of the utterance).

Key NLP Tasks and Applications I

NLP is the driving force behind many technologies we use every day:

NLP Task	Description	Example Application
Machine Translation	Automatic translation of text or speech between languages.	Google Translate.
Sentiment Analysis	Determining the emotional tone of text (positive, negative, neutral).	Analyzing customer opinions on social media.
Named Entity Recognition (NER)	Identifying and classifying proper nouns in text (people, places, organizations, dates).	Automatic indexing and database creation from documents.

Key NLP Tasks and Applications II

Speech Recognition	Converting spoken language into written text.	Voice assistants (Siri, Alexa) and recording transcription.
Natural Language Generation (NLG)	Creating coherent, readable text from structured data.	Automatic writing of sports summaries or financial reports.
Q&A Systems/Chatbots	Answering user questions or conducting conversations in natural language.	Virtual customer service assistants.

Key NLP Tasks and Applications III

Text Summarization	Automatically creating shortened versions of long documents while retaining key information.	Summaries of scientific or press articles.

What are Expert Systems?

Expert systems are computer programs belonging to the field of **Artificial Intelligence (AI)** that have the ability to **mimic the decision-making process and knowledge of a human expert** in a very narrow and specialized domain.

Instead of relying on traditional, procedural code, these systems use codified knowledge to solve complex problems, propose solutions, diagnose, or advise.

Main Components of an Expert System I

The architecture of an expert system consists of several key, interacting elements:

1. Knowledge Base

   This is the heart of the system, containing all the specialized knowledge in the given domain. This knowledge is usually represented in the form of:

   - **Facts:** Basic information and data.
   - **Rules (If-Then):** Decision rules (e.g., *IF* the patient has a fever and a cough, *THEN* there is a probability of flu).
   - **Heuristics:** Informal rules and "rules of thumb" used by experts.

Main Components of an Expert System II

2. Inference Engine

   This is the “brain” of the system that processes and interprets the knowledge from the Knowledge Base. Its task is to:

   - Analyze user input data.
   - Apply rules and facts to draw logical conclusions.
   - Use inference methods (e.g., **forward chaining** – from facts to a conclusion; **backward chaining** – from a hypothesis to facts that support it).

3. User Interface

   Allows the user to input data (e.g., symptoms, test results, parameters) and receive a diagnosis or advice in an accessible format.

Main Components of an Expert System III

4. Explanation Facility

   An optional but crucial component that distinguishes an expert system from a regular program. It explains to the user **why** the system made a particular decision or suggested a specific conclusion, increasing trust in its recommendations.

Applications of Expert Systems

Expert systems are highly specialized, which allows them to achieve high accuracy in a narrow field. They are used, among others, in:

- **Medicine:** Diagnosis of diseases (e.g., the famous **MYCIN** system), recommending treatment plans.
- **Finance:** Credit rating, fraud detection, investment advice.
- **Engineering:** Diagnosing machine and equipment failures, planning and designing complex systems.
- **Management:** Supporting strategic decisions, risk management.

What are the Fundamentals of LLM Operation? I

The basis of operation for **Large Language Models (LLMs)** is the mechanism of **predicting the next word** (or, more precisely, the **token**) in a sequence, based on the analysis of massive textual datasets and an advanced neural network architecture called the **Transformer**.

Here are the key pillars of their functioning:

1. Transformer Architecture I

At the heart of every modern LLM is the Transformer architecture, introduced in 2017. It enabled a breakthrough in language processing thanks to:

- **Tokenization:** Input text (e.g., a sentence) is divided into smaller units called **tokens** (these can be whole words, parts of words, or single characters). These tokens are then converted into **numerical vectors (word embeddings)** because computers operate on numbers, not words. These vectors encode meaning and context.

1. Transformer Architecture II

- **Attention mechanism/Self-attention:** This is the key element of the Transformer. It allows the model to **weigh the importance** of different tokens in the input text relative to the token it is currently processing. This enables the model to understand long-range dependencies in a sentence and maintain context, which is essential for correct natural language understanding (e.g., in the sentence “The dog ate the bone because it was hungry,” the attention mechanism links “it was hungry” to “The dog,” not “the bone”).

2. Pre-training on Vast Datasets I

LLM models are pre-trained on **gigantic textual corpora**, which include billions or even trillions of tokens, originating from various sources such as:

- Books
- Articles and websites (e.g., Wikipedia)
- Source code
- Dialogues and transcripts

The learning process is usually **unsupervised learning** and primarily involves:

- **Predicting the Next Token:** The model learns which token is most likely to appear after a given sequence of tokens.

2. Pre-training on Vast Datasets II

- **Filling in Missing Tokens (Masking):** In some architectures, the model is trained to guess which token was intentionally hidden.

As a result of this process, the model establishes a huge number of **parameters** (weights in the neural network) that represent statistical patterns and dependencies in language.

3. Text Generation as Sequence Prediction I

When a user enters a query (the **prompt**), the LLM generates a response in the following steps:

1. **Prompt Analysis:** The prompt is tokenized and processed by the Transformer to establish the context.
2. **Generating the First Token:** The model calculates the **probability** of every possible next token occurring, based on the entire prompt context. The token with the highest (or randomly weighted) probability is selected.
3. **Iteration:** The newly generated token is added to the sequence, and the model uses the entire new sequence (prompt + first token) as context to predict the **next token**.

3. Text Generation as Sequence Prediction II

4. **Continuation:** The process repeats, token by token, until a complete and coherent response is generated (e.g., reaching a token limit or generating an end-of-sentence token).

The LLM does not "understand" text in the human sense. It operates as an incredibly advanced **prediction machine**, generating the most logical and probable sequences of words that align with the statistical patterns it learned during training.

Politechnika Wrocławska

4. Fine-tuning and RLHF I

After initial pre-training, the base model is often **fine-tuned** to better serve specific goals and increase its utility:

- **Instruction Tuning:** The model is taught to respond better to user instructions and queries.
- **Reinforcement Learning from Human Feedback (RLHF):** This stage is crucial for conversational models (like ChatGPT). It involves:
  1. Generating multiple responses to a single prompt.
  2. Evaluating these responses by human raters for quality, helpfulness, and safety.
  3. Training a **Reward Model** based on these human evaluations.
  4. Using the Reward Model to further optimize and improve the main LLM using reinforcement learning algorithms.

4. Fine-tuning and RLHF II

Thanks to RLHF, the model learns not only **what** to say but also **how** to communicate, so that its responses are more natural, relevant, and consistent with user expectations.

History of LLM

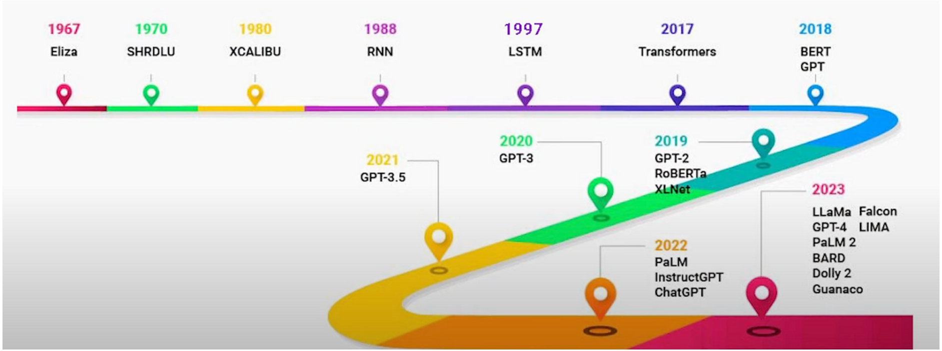

Mathematical Foundations of Artificial Intelligence I

1. **Linear Algebra** Data in AI are represented as mathematical structures – vectors, matrices, and tensors – which are the domain of linear algebra.
   - *Vectors and Matrices:* Input data (e.g., image features, words in text, measurement values) are encoded as vectors and matrices. Operations such as matrix multiplication are the basis of neural network operation (each layer is essentially a matrix operation).
   - *Transformations:* Linear algebra describes how data is transformed in a multidimensional space, which is the essence of learning processes, e.g., dimensionality reduction (PCA).
   - *Eigenvalues and Eigenvectors:* Used for principal component analysis and data compression.

Mathematical Foundations of Artificial Intelligence II

2. **Probability and Statistics** AI and ML rely on data, and real-world data is inherently uncertain and random. Probability and statistics allow modeling and managing this uncertainty.
   - *Probability Theory:* Used to model randomness and determine the probability of events.
   - Statistics: Used for data analysis, model evaluation and validation (e.g., hypothesis testing), and measuring algorithm performance (e.g., mean squared error).
3. **Differential and Integral Calculus**

   It is essential for optimization, which is the key process where ML algorithms learn by minimizing the error (cost function).

- *Derivatives and Gradients:* The most important concept is the gradient – a vector of partial derivatives that indicates the direction of the fastest increase in the cost function or minimization of the error.
- *Vector Calculus:* Extends differential and integral calculus to functions of multiple variables, which is typical for operations on multidimensional data.

4. **Optimization**

Optimization methods are the essence of the machine learning process.

Politechnika Wrocławska

???!

AI Slop I

"**AI slop**" (loosely translated as "AI swill," "AI mush," or "digital sewage") refers to **low-quality, mass-generated content** (text, images, videos, audio) using generative artificial intelligence tools, which:

1. **Are created with minimal effort and without human reflection.** The priority is quantity and speed, not content or quality.
2. **Are repetitive, chaotic, worthless, or lacking deeper meaning.**
3. **Flood the internet**, especially social media, blogs, and search results, making it difficult to find valuable, original content. They often aim to generate views, clicks, or advertising revenue (so-called "content farming" or "SEO slop").
4. **May contain subtle errors, inaccuracies, simplifications, or distortions** (so-called AI "hallucinations") that are presented as truth.

AI Slop II

5. **Can sometimes be easy to recognize** (e.g., unnatural-looking hands in images, strange proportions, chaotic backgrounds, grammatical and stylistic errors in text), but their sheer volume and increasingly better quality make this harder.

In short: AI slop is digital "garbage" or "noise" that is a side and undesirable effect of the mass use of generative artificial intelligence. This phenomenon leads to a decrease in the overall quality of the internet and trust in online content.

Problems I

1. Energy Consumption
   - Training
   - Chat
2. "Free Will"
3. The Real World
   - Humans operate in physical reality
   - Artificial Intelligence operates in...?
4. Humans cooperate, work in groups, exchange views — AI "rolls the dice"...
5. Sources of "Knowledge"

Problems II

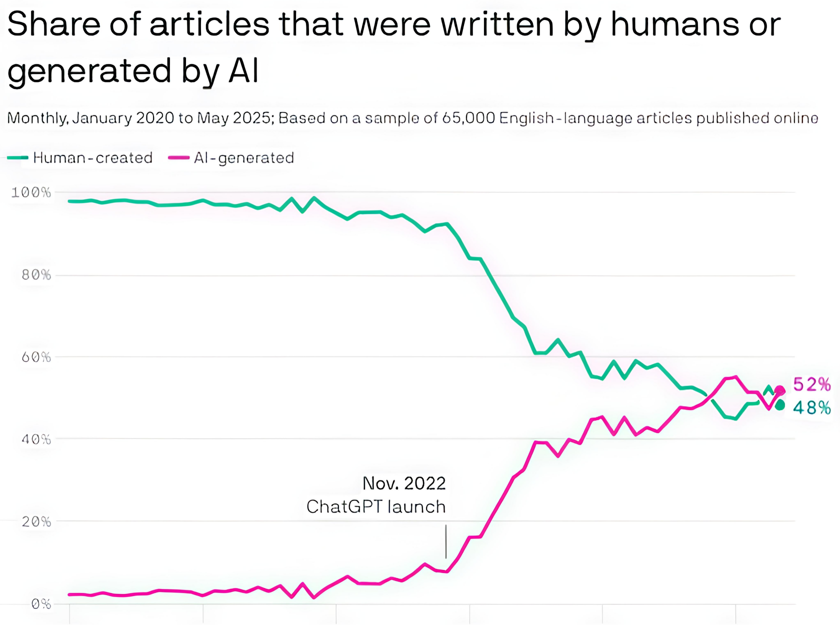
Share of articles that were written by humans or generated by AI
Monthly, January 2020 to May 2025; Based on a sample of 65,000 English-language articles published online
Human-created
AI-generated
100%
80%
60%
40%
20%
0%
Nov. 2022
ChatGPT launch
52%
48%

Politechnika Wrocławska

An AI-powered podcast startup publishes over 3,000 episodes a week, all hosted by AI-generated presenters.

???

Paint fancy clock showing half past six time

Here is a fancy clock showing half past six:

Paint glass full of red wine

Here you go! A glass full of red wine:

Colophon

Politechnika Wrocławska

This presentation was prepared using the markup language Markdown; the source text was converted to LaTeX using the program pandoc.

The program pdfTeX was used to generate the slides in PDF format. The beamer class and the Wrocław University of Science and Technology presentation template were utilized. The iwona font was used.

The substantive content was partly generated in the author's mind and partly by Artificial Intelligence (Google Gemini, Perplexity,...).

No AI was harmed during the preparation of the lecture (which cannot be said about the undersigned who put everything together).