PRESIDENCY OF THE REPUBLIC OF TÜRKİYE
DIGITAL TRANSFORMATION OFFICE

Introduction to Artificial Intelligence and Fundamental Concepts

Dr. Melike PALSÜ KURT

What Will We Learn?

- What is Artificial Intelligence?
  - Can Machines Think, and How Can They Think?
  - Turing Test
  - Chinese Room
  - The AI Lifecycle
- What Can Artificial Intelligence Do?
- What Can Artificial Intelligence Not Do?
- Fundamental Concepts
- Types of Data
- Data Preprocessing
- Feature Engineering
- Good Practice Examples

What is Artificial Intelligence?

What is Artificial Intelligence?

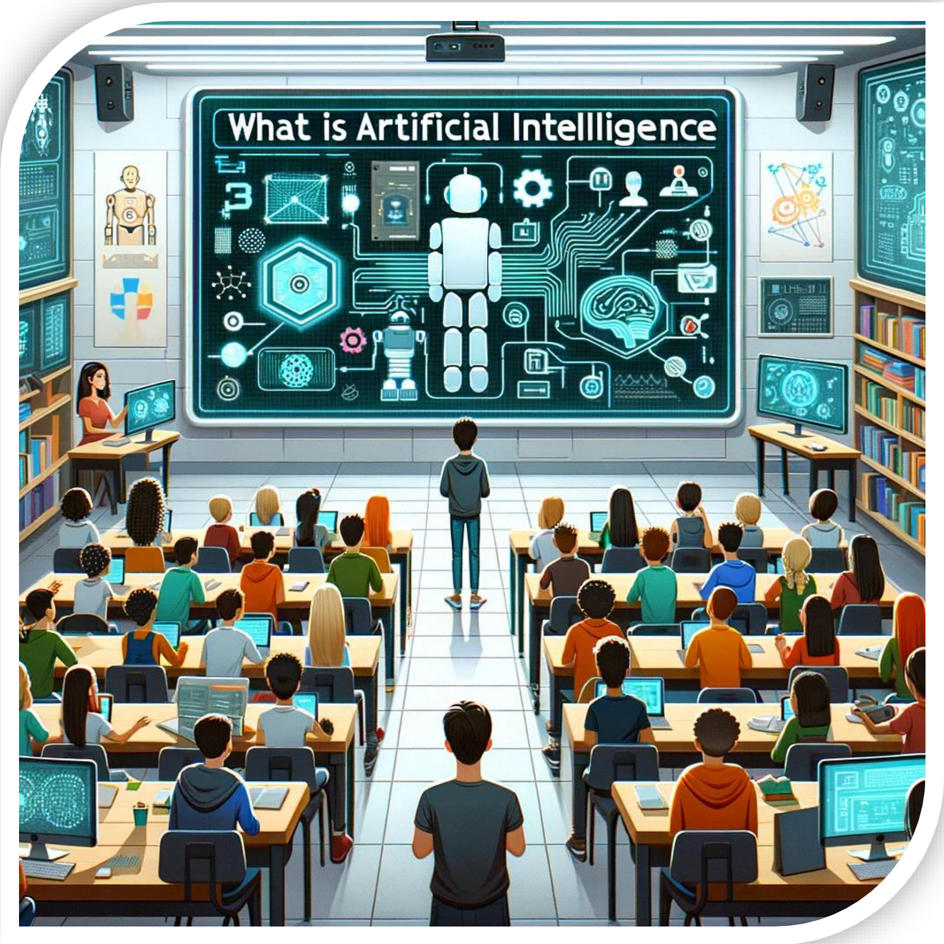

In its most general form, artificial intelligence (AI) is defined as the ability of a computer or a robot controlled by a computer to perform various tasks in a manner similar to intelligent beings.

National Artificial Intelligence Strategy of Türkiye 2021 - 2025

What is Artificial Intelligence?

Narrow AI is effective only in the specific areas it has been trained for, unlike human intelligence; it deals with a specific system and a specific problem.

For example: fraud detection, facial recognition, or social recommendations, etc.

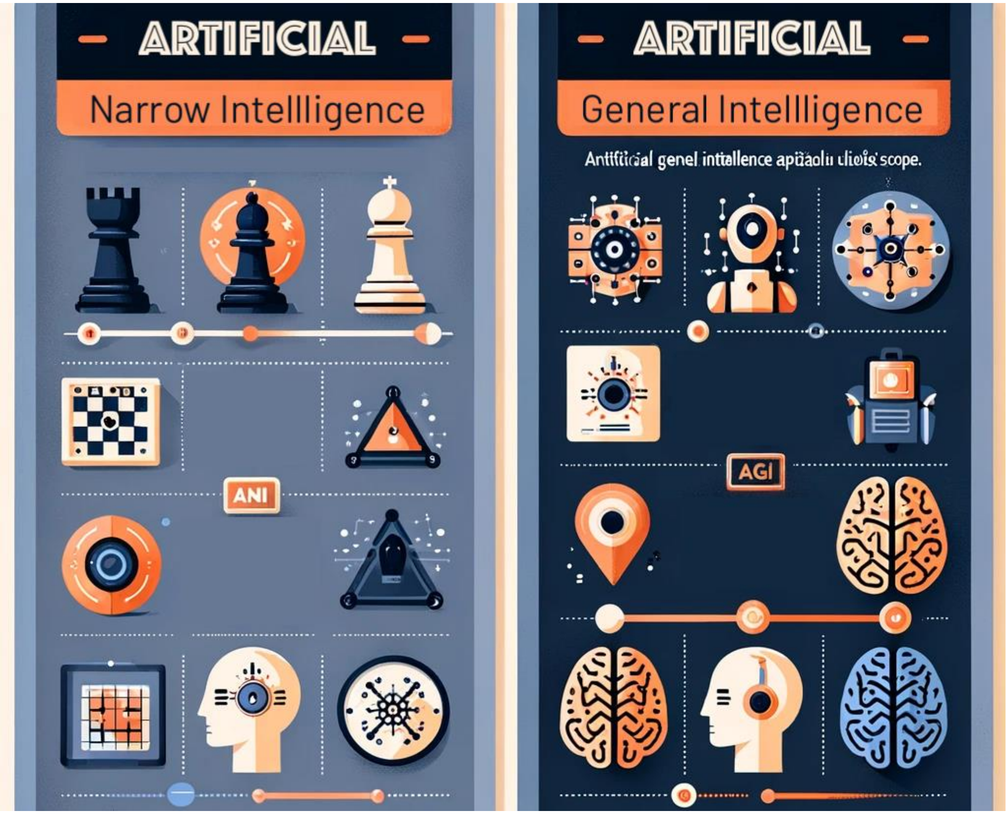

Although it does not yet exist, General AI is expected to function like human intelligence. The main challenge in creating General AI is modeling the world in a consistent and useful way with its entire body of knowledge.

The Dartmouth Summer Research Project 1956

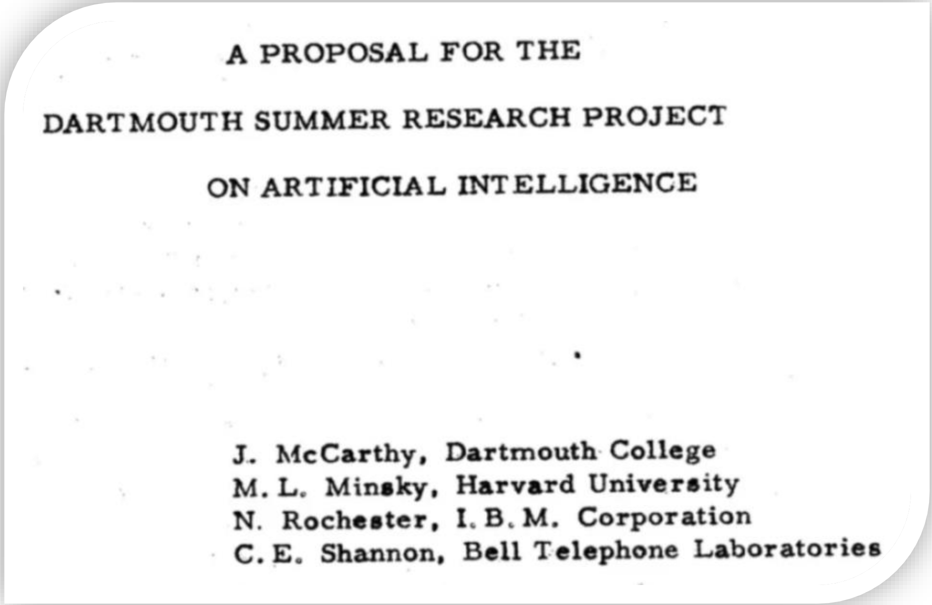
A PROPOSAL FOR THE

DARTMOUTH SUMMER RESEARCH PROJECT

ON ARTIFICIAL INTELLIGENCE

J. McCarthy, Dartmouth College
M. L. Minsky, Harvard University
N. Rochester, I. B. M. Corporation
C. E. Shannon, Bell Telephone Laboratories

Initially, the program was planned with the participation of 11 mathematicians, but over the two-month period, many more scientists joined. Among these scientists were:

•**John McCarthy**, who was awarded multiple times for his work on AI and is the developer of Lisp, one of the oldest and most powerful programming languages,

•**Marvin Lee Minsky**, co-founder of the AI Laboratory established at the Massachusetts Institute of Technology, and

•**John Nash**, the legendary American mathematician known for his work in game theory, differential geometry, and partial differential equations.

The Intelligence Revolution: The Birth of Artificial Intelligence

The first meeting of the conference was organized by mathematics professor John McCarthy.

McCarthy proposed that any feature of learning or intelligence could, in principle, be precisely described in order for a machine to simulate it.

Researchers gathered around this proposal began to explore ways to make machines more intelligent and started investigating how to better understand human intelligence by forming a general framework.

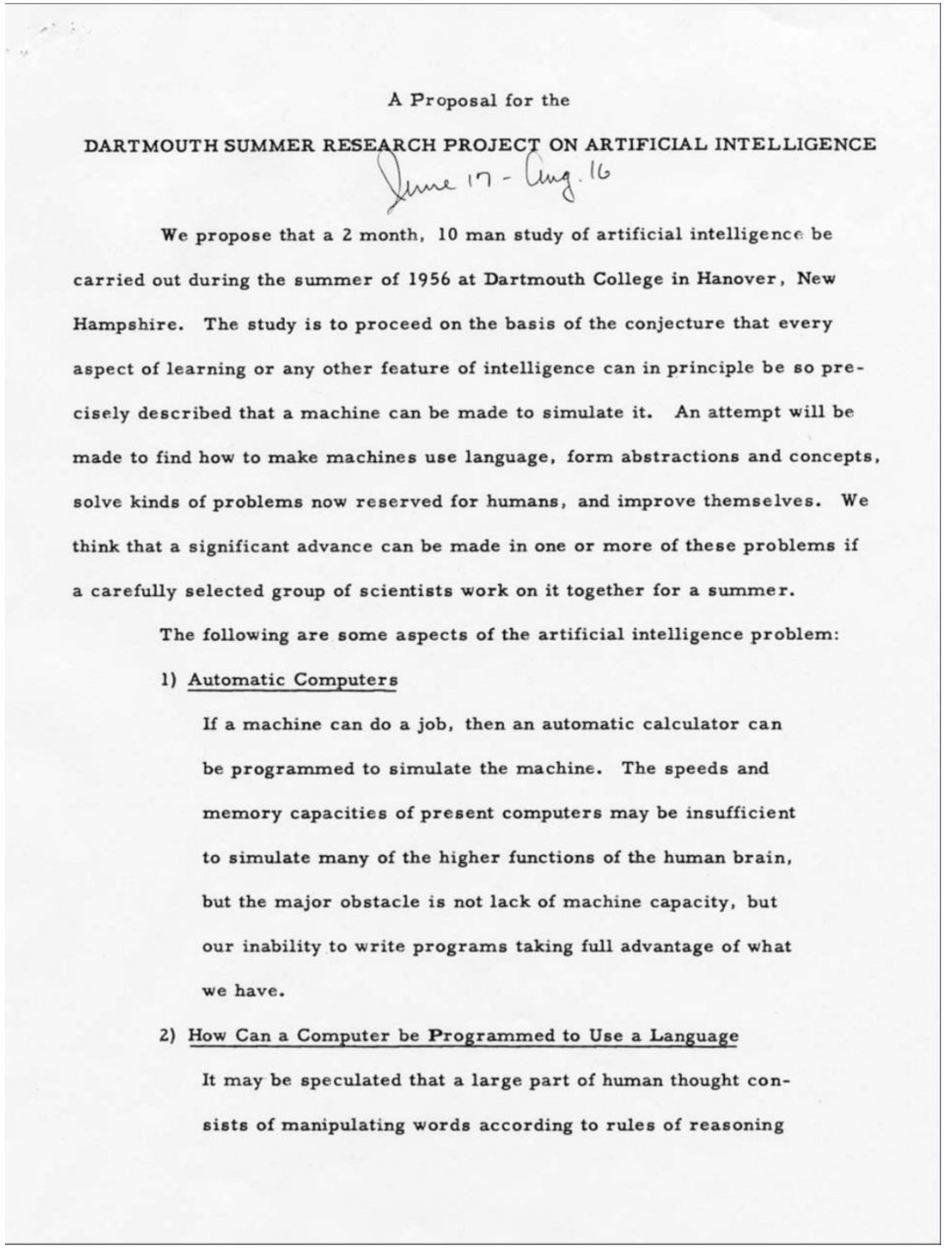

A Proposal for the

DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

June 17 - Aug. 16

We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

The following are some aspects of the artificial intelligence problem:

1) Automatic Computers

If a machine can do a job, then an automatic calculator can be programmed to simulate the machine. The speeds and memory capacities of present computers may be insufficient to simulate many of the higher functions of the human brain, but the major obstacle is not lack of machine capacity, but our inability to write programs taking full advantage of what we have.

2) How Can a Computer be Programmed to Use a Language

It may be speculated that a large part of human thought consists of manipulating words according to rules of reasoning

The Intelligence Revolution: The Birth of Artificial Intelligence

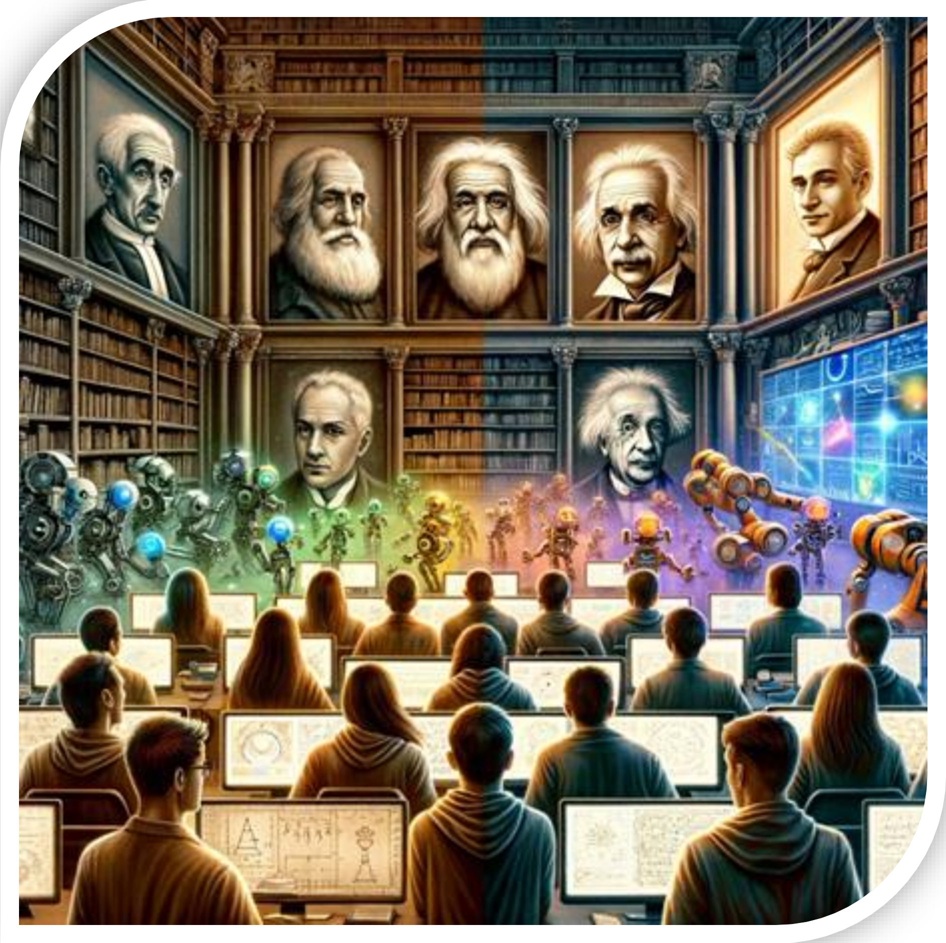

“A student in physics might reasonably feel that all the good ideas have already been taken by Galileo, Newton, Einstein, and the rest.

AI, on the other hand, still has openings for several full-time Einsteins and Edisons.”

Russell, S. J., & Norvig, P. (2010). Artificial intelligence a modern approach. London.

The Intelligence Revolution

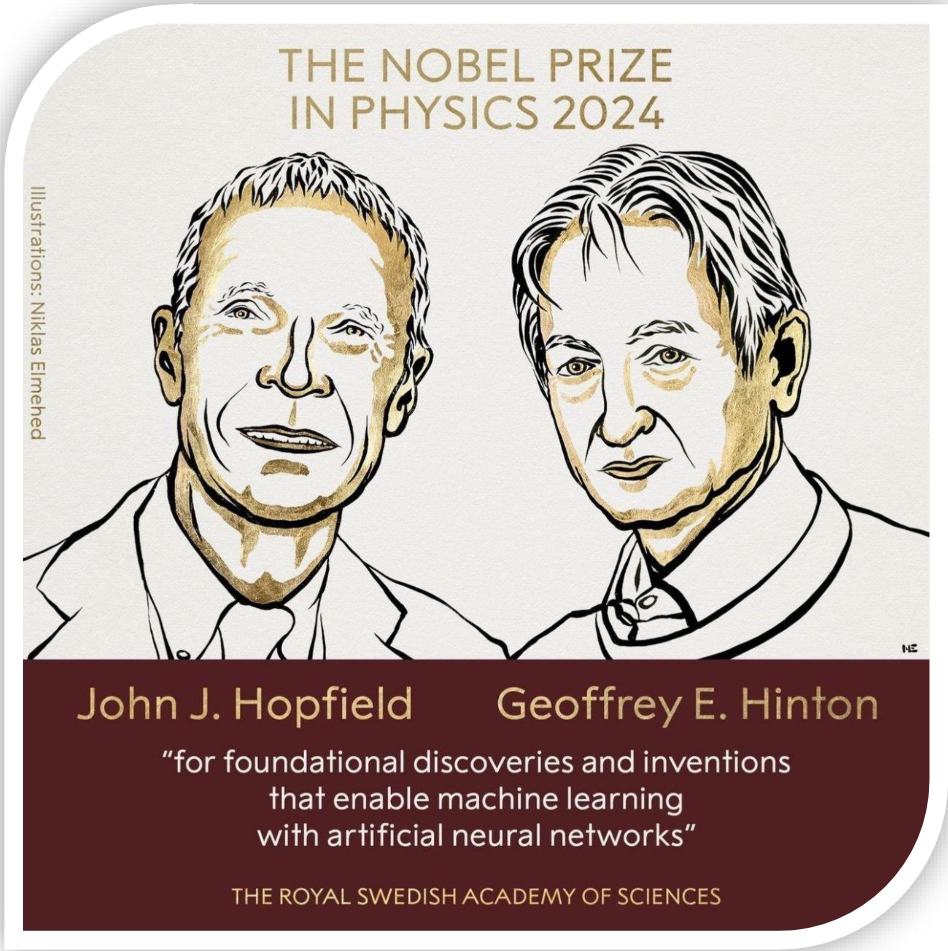

John Hopfield and **Geoffrey Hinton**, who developed artificial neural networks (ANN) and foundational methods for machine learning with the help of physics, were awarded the **2024 Nobel Prize in Physics** for their "fundamental discoveries and inventions enabling machine learning through artificial neural networks.»

Hopfield invented a method for constructing a relational memory system capable of storing and reconstructing patterns such as images, information, and other types of data.
Hinton, on the other hand, invented a technique that can identify features in data without human intervention, enabling tasks like recognizing specific elements in images.

The Intelligence Revolution

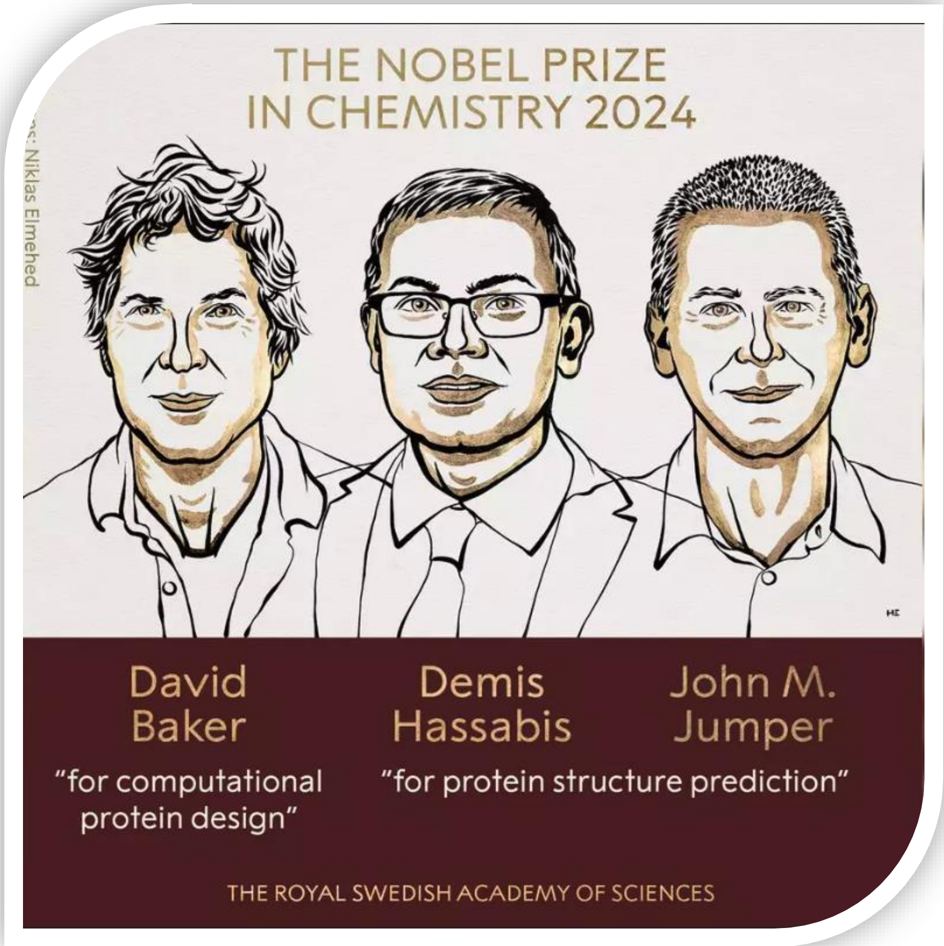

The **2024 Nobel Prize in Chemistry** was awarded to **David Baker** for his work in **computational protein design**, and to **Demis Hassabis** and **John Jumper** for their contributions to **protein structure prediction**. **Wedell**, **Hassabis**, and **Jumper** stated that their AI model **AlphaFold2**, which can predict the structure of nearly all 200 million proteins, has been made available for public use.

> **Adam Smith**, president of the world's oldest scientific academy, the Royal Society, stated that the award given to the developers of AlphaFold shows that *"the transformative role of artificial intelligence in science is clearly recognized."*

Can Machines Think, and If So, How?

“It becomes evident that no devilish intelligence is needed to understand these marvelous machines of our time. Mere common sense suffices.”

Cahit ARF

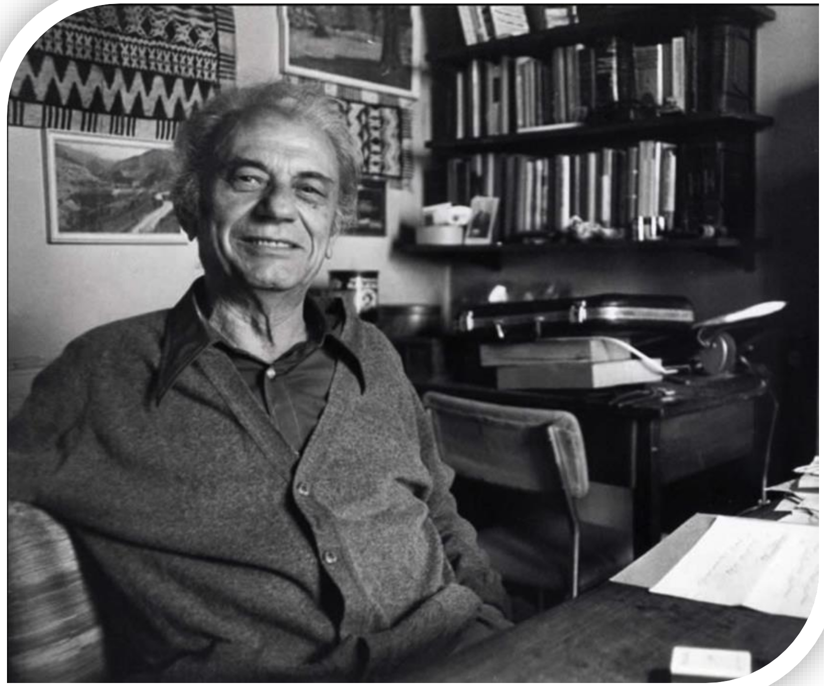

Cahit Arf’s article *“Can Machines Think, and If So, How?”* addresses fundamental questions in the fields of artificial intelligence and machine learning.

In this article, Arf discusses whether machines can possess human-like thinking abilities and how such a process might be realized. Drawing on his deep expertise in mathematics and computer science, he sheds light on the future potential of machines' cognitive capacities.

When Is a System Considered «Intelligent»?

Thinking Like a Human	Thinking Rationally
"The exciting new effort to make comput- ers think . . . machines with minds, in the full and literal sense." (Haugeland, 1985) "[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solv- ing, learning ..." (Bellman, 1978)	"The study of mental faculties through the use of computational models." (Charniak ve McDermott, 1985) "The study of the computations that make it possible to perceive, reason, and act ." (Winston, 1992)
Acting Like a Human	**Acting Rationally**
"The art of creating machines that per- form functions that require intelligence when performed by people." (Kurzweil, 1990) "The study of how to make computers do things at which, at the moment, people are better." (Rich ve Knight, 1991)	"Computational Intelligence is the study of the design of intelligent agents." (Poole et al., 1998) "AI ...is concerned with intelligent be- havior in artifacts." (Nilsson, 1998)

Acting Like a Human: The Turing Test

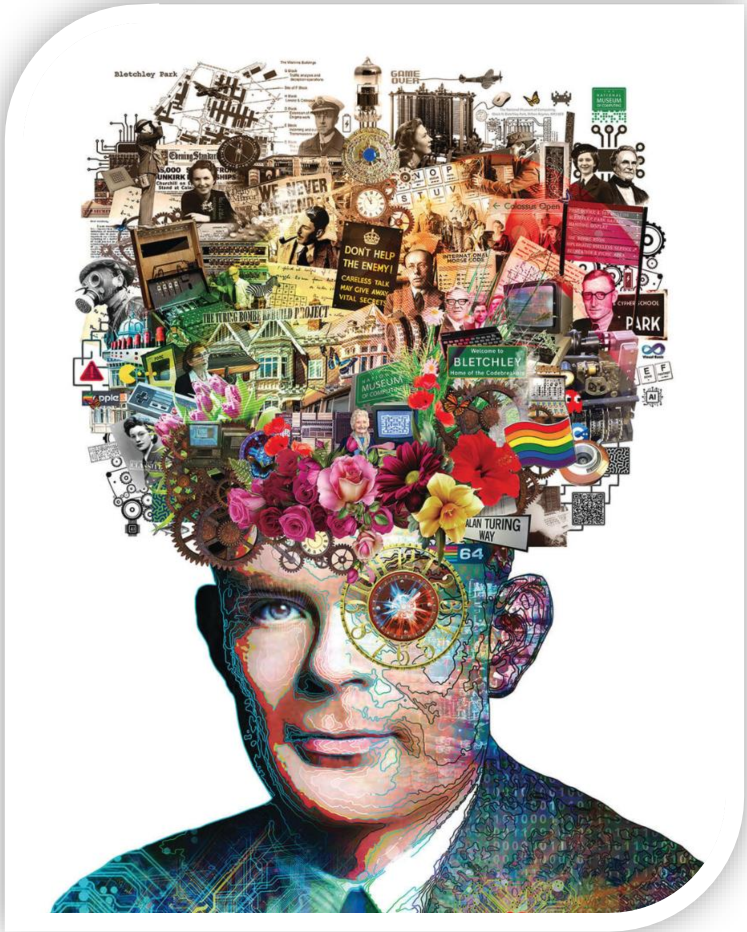

The **Turing Test** is a test proposed in 1950 by British mathematician and computer scientist **Alan Turing**, designed to evaluate whether a machine can exhibit intelligent behavior equivalent to, or indistinguishable from, that of a human.

In this test:
•A **judge** (a human) communicates in writing with both a human and a machine (typically a computer program),
•The judge tries to determine which is the human and which is the machine,
•If the judge cannot reliably distinguish the machine's responses from the human's, the machine is said to have passed the Turing Test.

Chinese Room

Chinese Room, proposed by **John Searle** in 1980, is a thought experiment. This experiment addresses the question of whether artificial intelligence can truly "understand" or gain consciousness. **Searle** argues that the ability of a computer program to process human language does not mean that the program actually understands the meaning or experiences mental states.

In the experiment, a person who does not know Chinese sits inside a room and responds to Chinese questions from outside using a complex rulebook — producing Chinese answers by simply following the instructions.

People outside the room may believe that the person inside understands Chinese. However, the person in the room does not actually understand Chinese; they are merely generating appropriate responses by following predefined rules.

AI: A Melting Pot of Interdisciplinary Theories

Probability	Bayes Rule(Bayes, 1763)
Mathematics	Least Squares Regression(Gauss, 1795)
Logic	First-order Logic (Frege, 1893)
Statistics	Maximum Likelihood (Fisher, 1922)
Neuroscience	Artificial Neural Networks (McCulloch/Pitts, 1943)
Economics	Minimax Games (von Neumann, 1944)
Optimization	Stochastic Gradient Descent (Robbind/Monro, 1951)
Algorithms	Uniform Cost Search (Dijkstra, 1956)
Control Theory	Value Iteration (Bellman, 1957)

The Artificial Intelligence Lifecycle

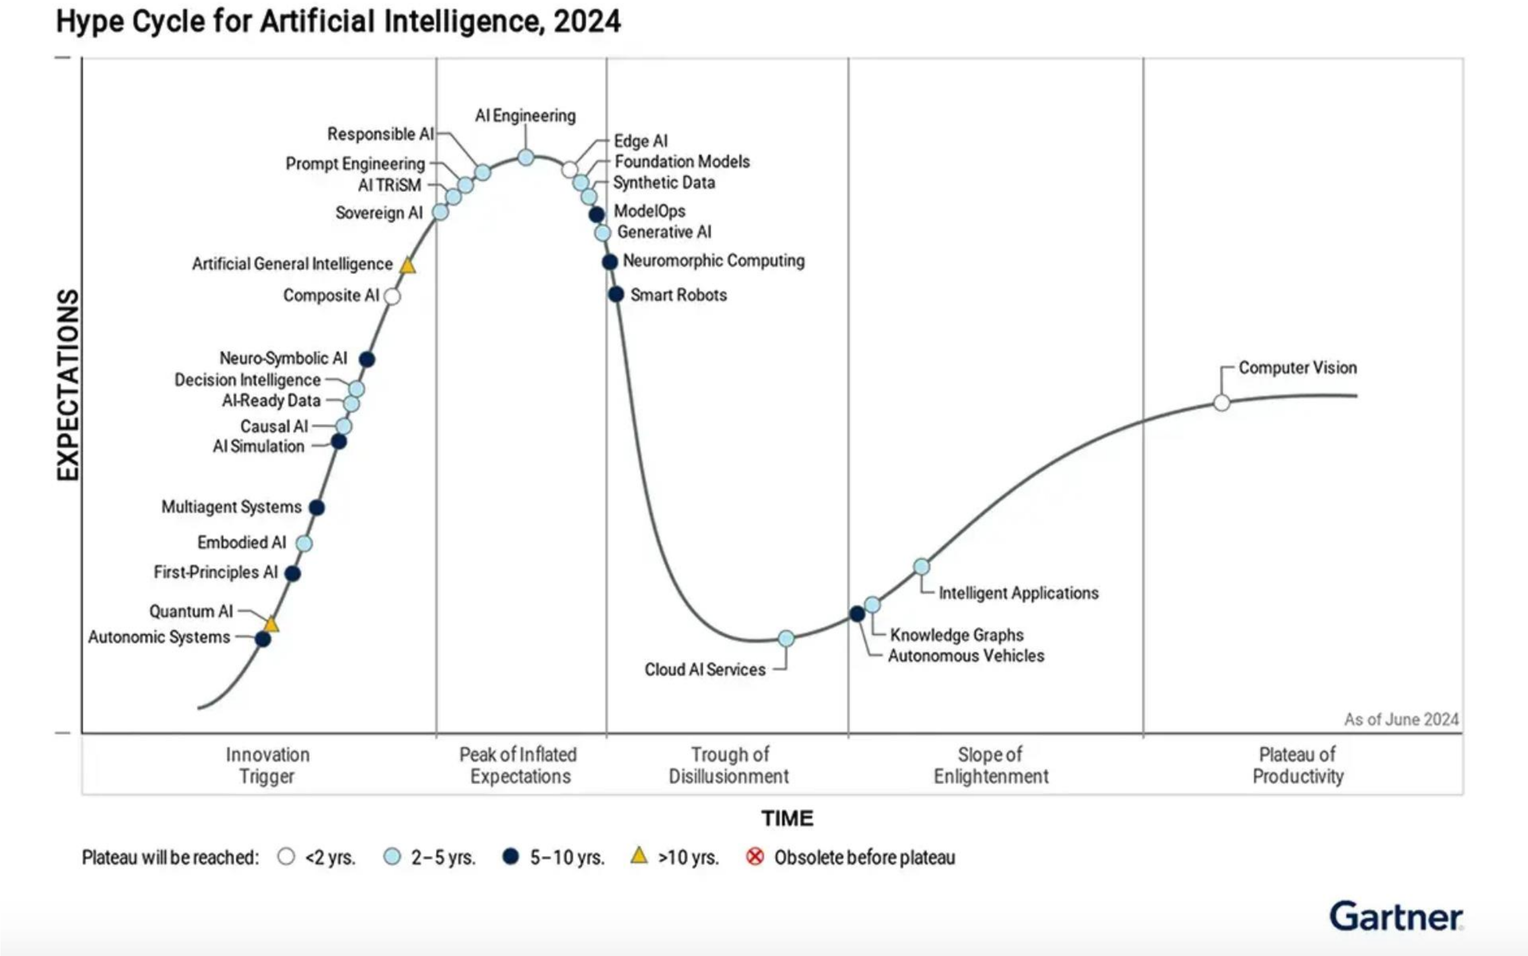

Gartner, 2024.

The Artificial Intelligence Lifecycle

Too Much Confidence

“Machines will be capable, within twenty years, of doing any work a man can do. ”
Herbert Simon

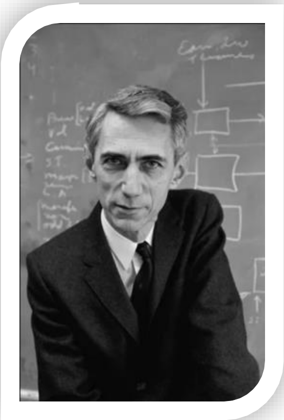

“I visualize a time when we will be to robots what dogs are to humans, and I’m rooting for the machines.”
Claude Shannon

“Within ten years, the problems of artificial intelligence will be substantially solved.”
Marvin Minsky

The Artificial Intelligence Lifecycle

The First Disappointment – The First AI Winter

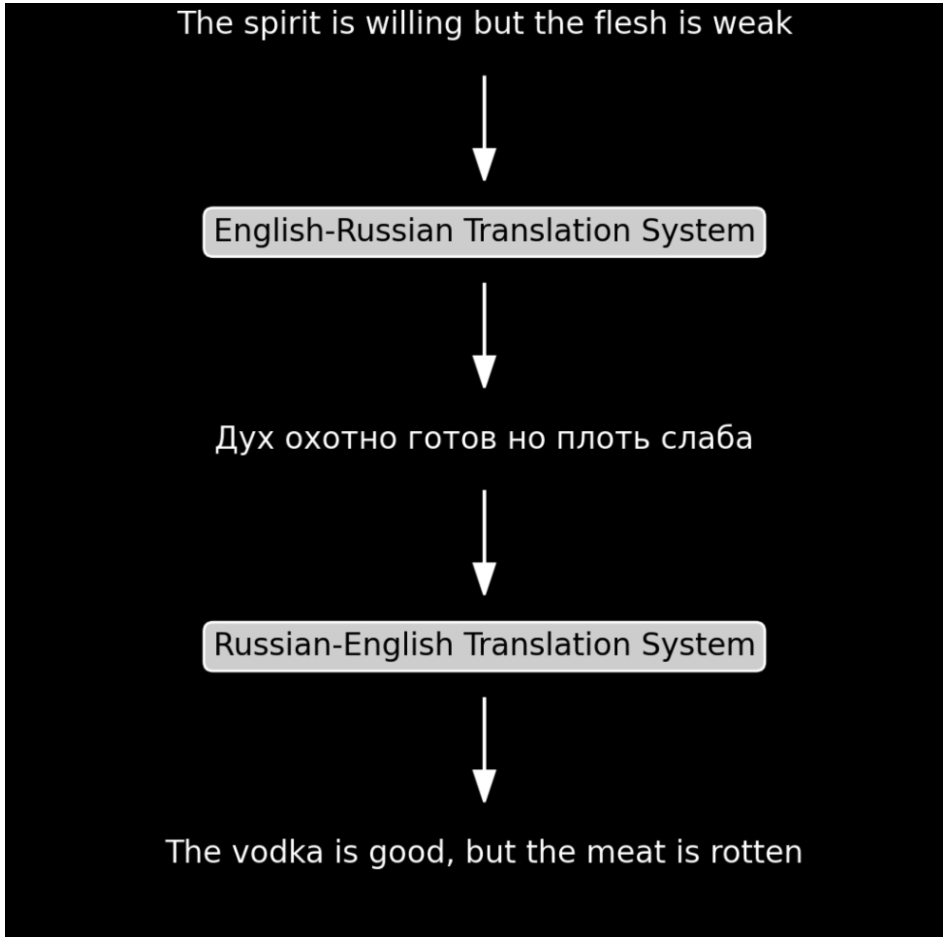

The **Automatic Language Processing Advisory Committee (ALPAC)** published a report in 1966 concluding that "machine translation of general scientific texts is not achievable, and likely will not be in the near future."

After spending 20 million dollars, the **National Academy of Sciences, Engineering, and Medicine** terminated government funding for AI research based on this report.

The Artificial Intelligence Lifecycle

The Second AI Winter

There was no single event that triggered the Second AI Winter; rather, it emerged from the accumulation of several factors:

- **Unfulfilled Promises:**
  In the 1980s, exaggerated claims and unfulfilled promises in the field of AI raised expectations. During this period, researchers and companies asserted that artificial intelligence would soon be able to replicate human intelligence and solve complex problems. However, most of these promises were not realized.
- **Technological Limitations:**
  In the early 1980s, hardware and software technologies were not advanced enough to provide the computing power and data storage capacity required by AI research.

The Artificial Intelligence Lifecycle

The Second AI Winter

There was no single event that triggered the Second AI Winter; rather, it resulted from the accumulation of several factors:

- **Reduction in Funding:**
  The inability of AI projects to deliver expected results and their commercial failures led to significant budget cuts, particularly in government and private sector funding. This caused a slowdown in research and development.
- **Disappointment and Skepticism:**
  The failure to meet rising expectations in AI led to disappointment and a skeptical attitude in both academic and industrial circles. As a result, interest from researchers and investors declined.

What can AI do?

Adoption Rates of AI Tasks in Industry

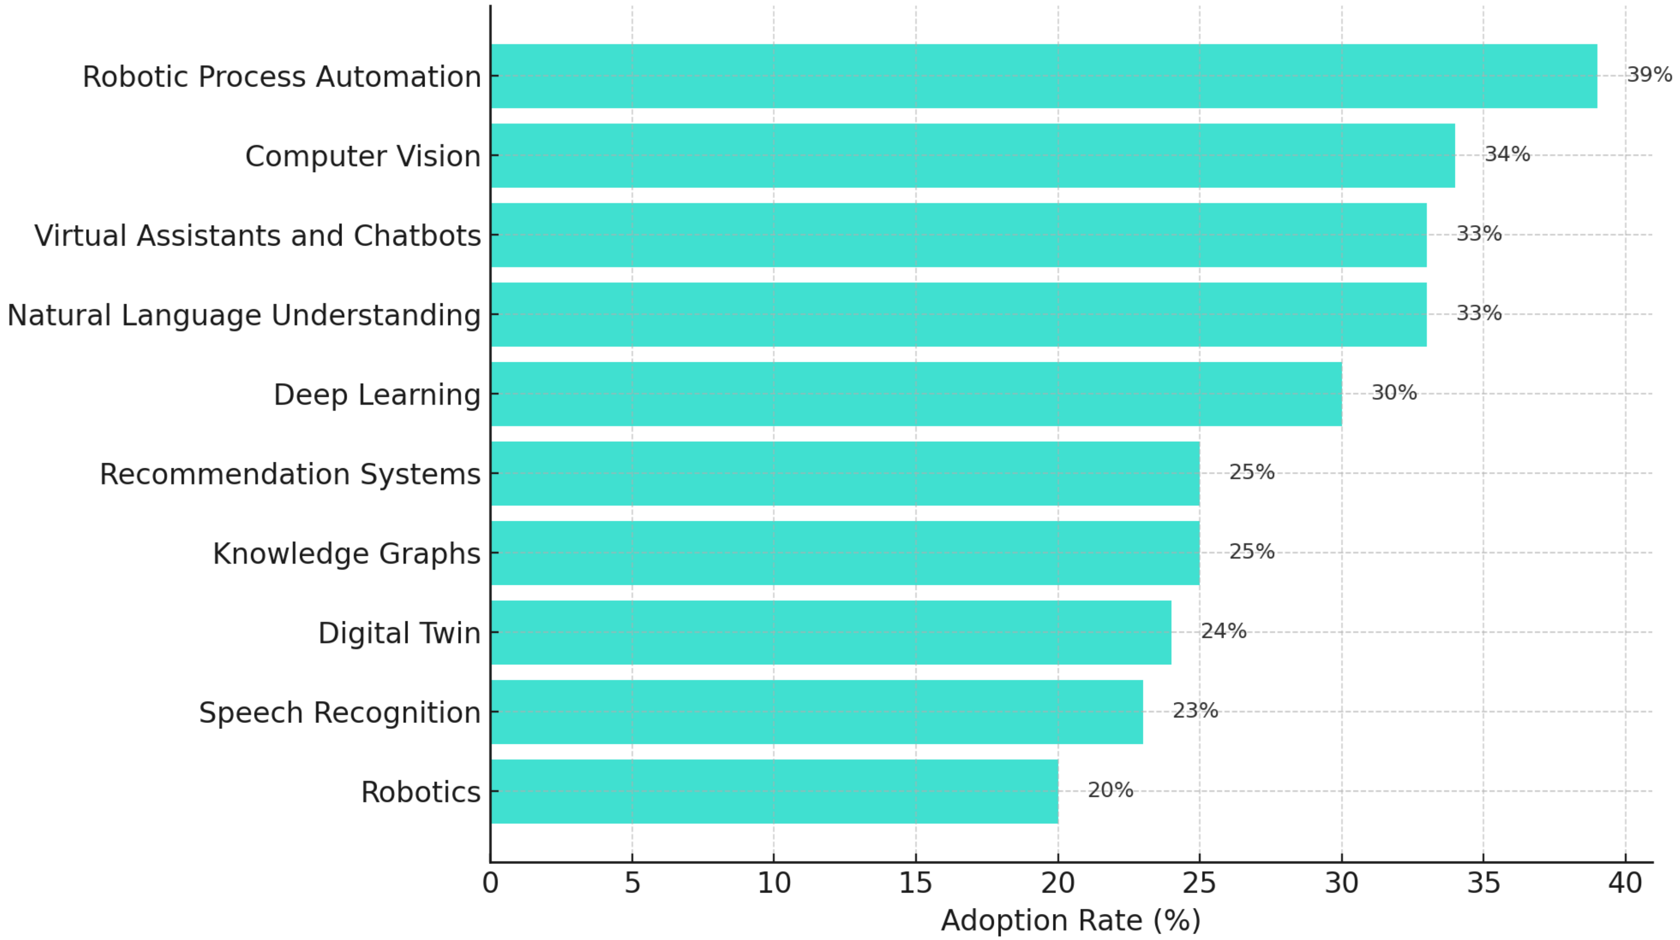

What Can’t Artificial Intelligence Do?

The “argument from disability” makes the claim that “a machine can never do X.”

As examples of X, Turing lists the following:

- Be kind,
- resourceful,
- beautiful,
- friendly,
- have initiative,
- have a sense of humor,
- tell right from wrong,
- make mistakes,
- fall in love,
- enjoy strawberries and cream,
- make someone fall in love with it,
- learn from experience,
- use words properly,
- be the subject of its own thought,
- have as much diversity of behavior as man,
- do something really new.

Artificial Intelligence Is Not a Superhero!

When Should AI be Used?

The Right Formula for Deciding on AI Use

1. Use Cases
What are the effective, measurable, and quickly solvable use cases?

2. Skills
Do you have the necessary skills to apply AI techniques?

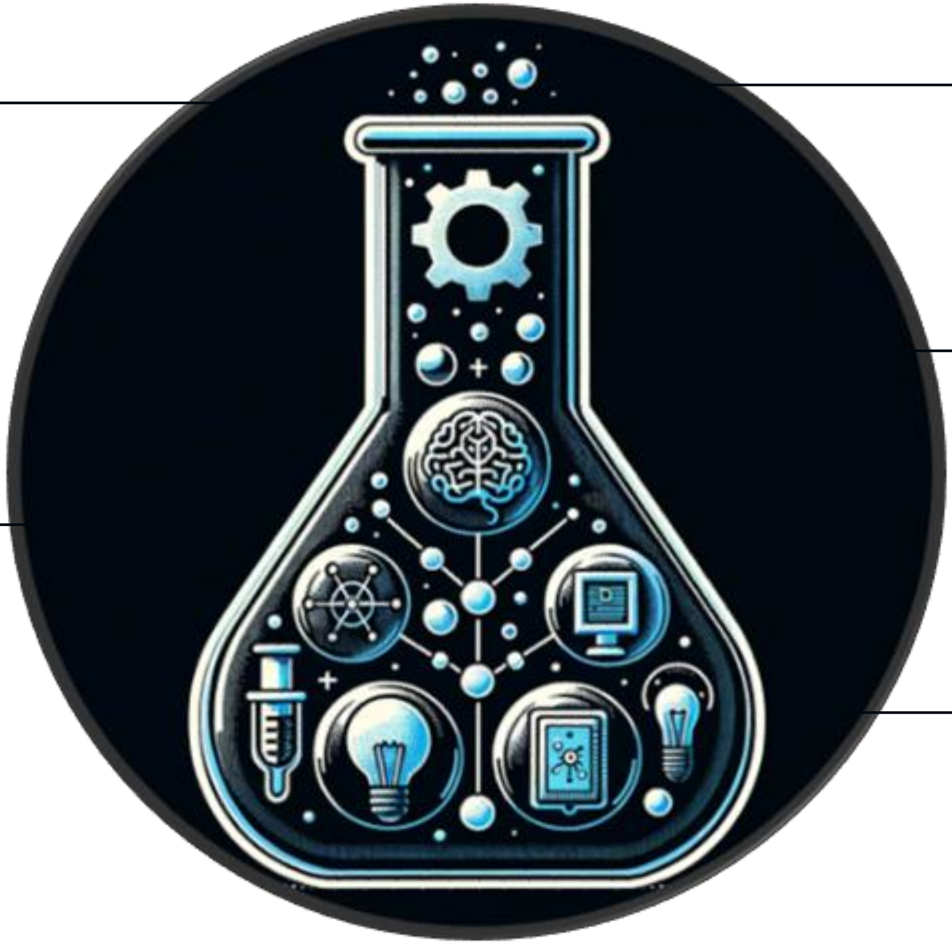

3. Data
Do you have relevant data for the selected use cases?

4. Technology
Considering the existing skills and available data, what techniques can be applied to solve the problem?

5. Organization
Based on the outcomes of concept studies, what is the organization's area of domain expertise?

Core Concepts

Core Concepts

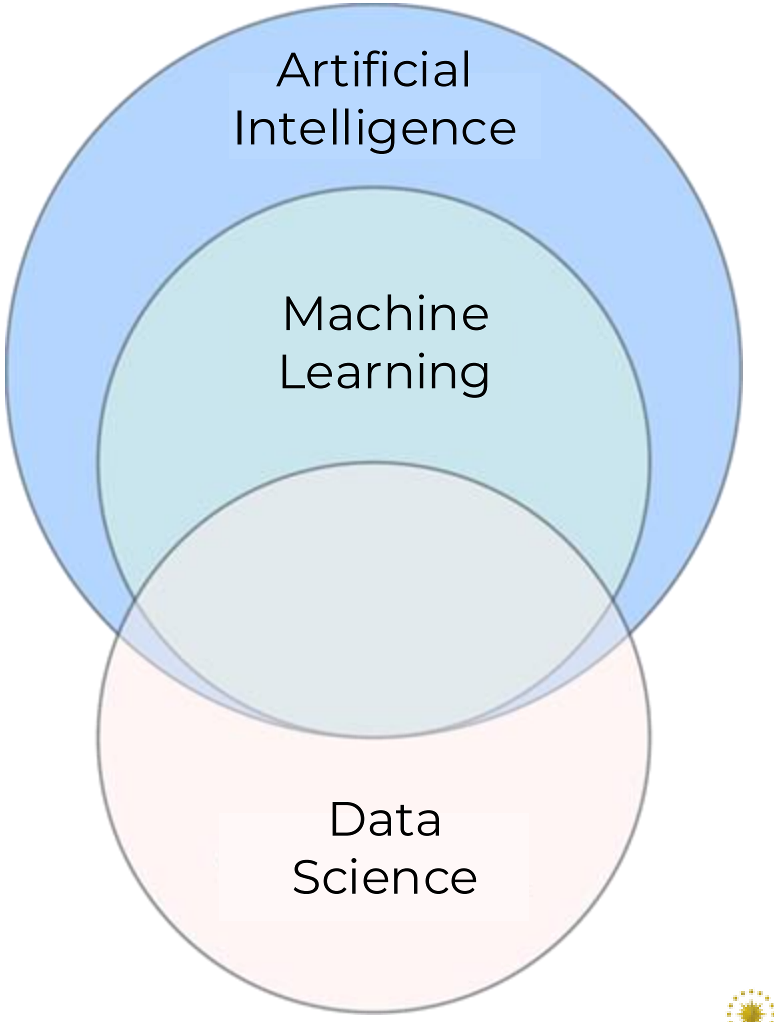

Artificial Intelligence: In its broadest sense, it refers to the ability of a computer or a computer-controlled robot to perform tasks typically requiring human intelligence.

Machine Learning: A subfield of AI that learns from examples.

Data Science: The collection of techniques used to extract value from large volumes of data.

Basic Terminology

Basic Terminology

Name	Height (cm)	Weight (kg)
Player 1	182	77
Player 2	175	69
Player 3	188	83
Player 4	190	85
Player 5	178	73
Player 6	180	75
Player 7	185	78
Player 8	179	74
Player 9	183	76
Player 10	177	70
Player 11	174	72

Feature Vector

Feature

Position
Goalkeeper
Right Back
Center Back
Midfielder
Left Winger
Striker
Midfielder
Right Winger
Center Back
Left Winger
Forward

Label

Data Point

Player → Position

$x_i \rightarrow y_i$

$x_i = (height, weight), i = 1, 2, \ldots, 11.$

$y_i = 1, 2, \ldots, 9.$

Dataset

Basic Terminology

Input

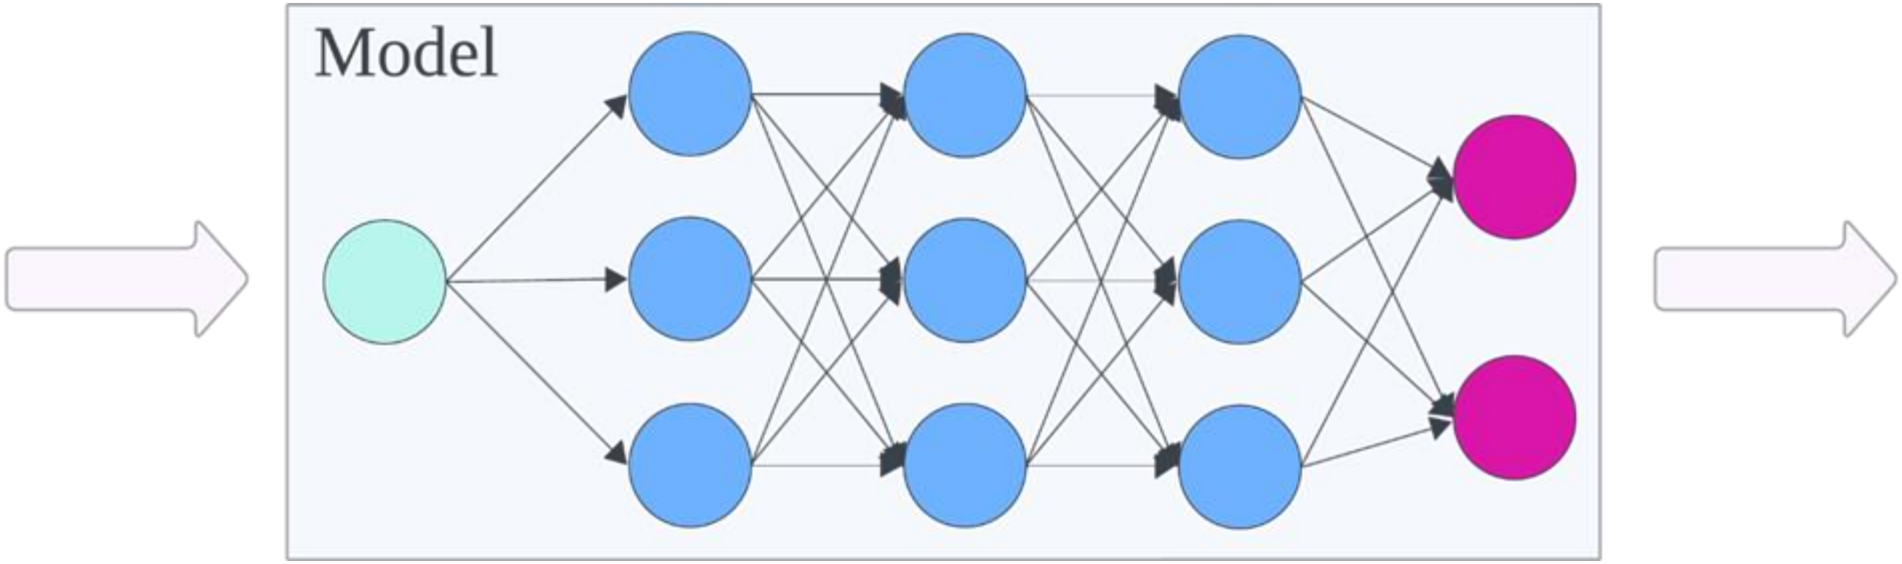

Output

What is Data?

What is Data?

Daily Data Production: 2.5 – 2.9 Billion GB

This volume continues to grow steadily due to factors such as:

• increased internet usage,

• social media interactions,

• growing data usage by businesses,

• the widespread adoption of smart devices, and

• the expansion of innovative technologies like the Internet of Things (IoT).

What is Data?

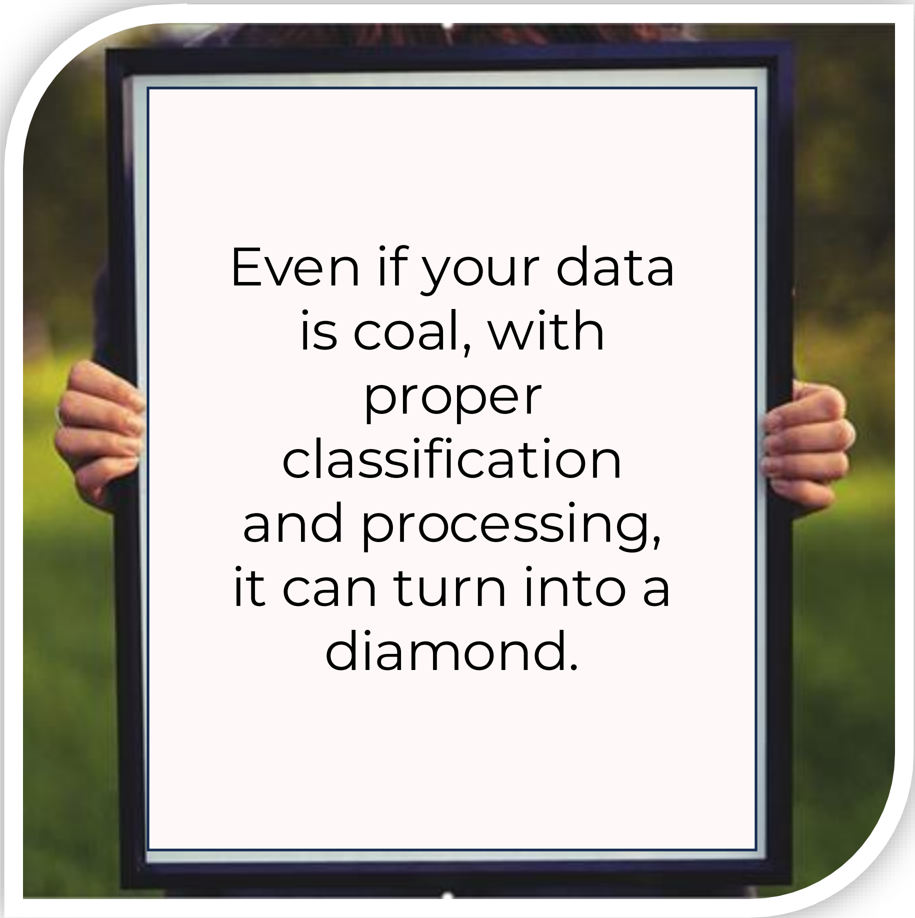
Even if your data is coal, with proper classification and processing, it can turn into a diamond.

DIKW Pyramid

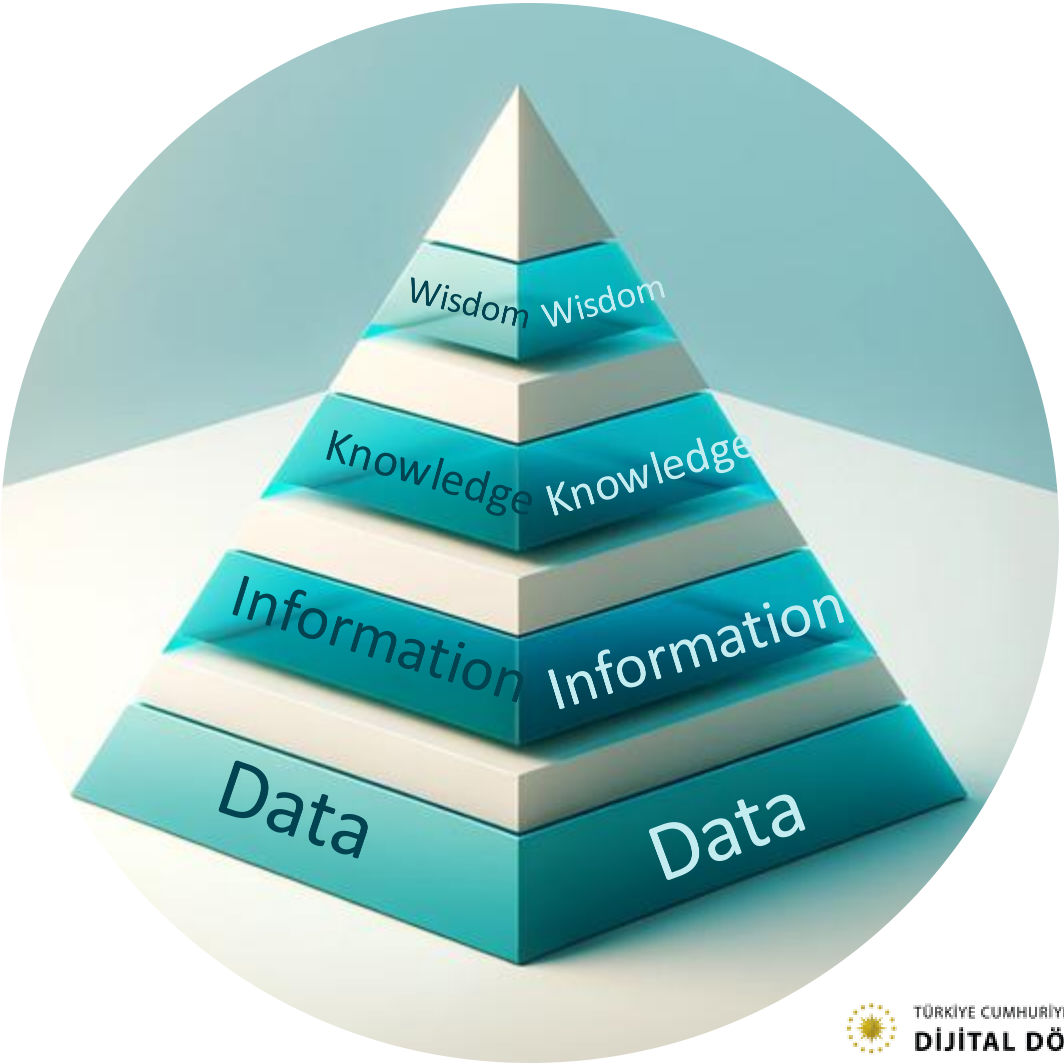

Data→ Information → Knowledge → Wisdom

DIKW Pyramid

Knowledge...

Information: Processed, organized, or structured data that gains meaning through contextual or relational connections.

Data: Raw facts and figures without context. On their own, they may be qualitative or quantitative, but they do not carry meaning until interpreted.

DIKW Pyramid

Wisdom:
The highest level of understanding. It involves the application of knowledge to enable sound and ethical decisions and predictions about the future. Wisdom is knowing what the next step should be, what needs to be done in the long term, and why it matters.

Knowledge:
A synthesis of information from various sources that provides a more comprehensive understanding of a topic or situation. Acquired knowledge is information combined with experience, context, interpretation, and reflection.

DIKW Pyramid

Data	• 3024 x 4032 photo.
Information	• A cat.
Knowledge	• My neighbor’s cat, Sütlaç.

Data-Driven Approach

The term **"data-driven"** refers to a decision-making approach based on concrete data and analysis rather than intuition or guesswork.

Today, businesses, governments, and other organizations adopt the data-driven approach to guide their strategic decisions, policies, and improvements.

This approach goes beyond intuitive predictions and aims to make more conscious and objective decisions by leveraging the power of data and analytics.

Data Types

Structured vs. Unstructured Data

- Can be viewed in rows, columns, and relational databases
- Numbers, dates, and strings
- Comprises approximately 20% of enterprise data (Gartner)
- Requires less storage space
- Easier to manage and protect with traditional solutions

- Cannot be viewed in rows, columns, or relational databases
- Images, audio, video, word processing files, emails, spreadsheets
- Comprises approximately 80% of enterprise data (Gartner)
- Requires more storage space
- Harder to manage and protect with legacy solutions

Structured vs. Unstructured Data

EXAMPLE	STRUCTURED PART	UNSTRUCTURED PART
Phone Calls	Who called, who answered, and the time of the call	Audio recording of the conversation
Insurance Forms	Date of submission, who filled it out, who wrote it, and the type of request	Content of open-ended questions
Podcasts	Date published, host name, and podcast category	Audio recording and transcript of the podcast
Server Logs	Timestamp, source microservice, and log level (info, error, debug, etc.)	Content of the log

Real vs. Synthetic Data

REAL DATA

- Data obtained from natural environments and originating from real-world events or situations.

SYNTHETİC DATA

- Data derived from real data or entirely created through artificial means, possessing similar characteristics to real-world data but not directly based on actual events or behaviors.

Biased Data

- It refers to the inadequacy of available data to represent the problem space, population, or the phenomenon being studied.
- In most large data spaces, it is observed that the data is biased.
- If the data used for training a model is biased, the model outputs are likely to be biased as well.

Biased Data

A model learns from the data it is given —
biased input results in biased output.

High-Quality Data

AI applications succeed when they are powered by the right amount of quality data.

High-Quality Data

Clean Data: Data that has been cleaned of inconsistencies and errors.
Relevant Data: Data that is directly related to the problem and contributes to its solution.
Diversity: Data that includes different scenarios and variations to improve model generalization.
Balanced Data: Data that fairly represents each class in classification problems and does not contain bias.
Up-to-Date Data: Data that reflects the current state of the problem and has not lost its relevance.
Accessibility: Data that can be easily accessed and processed when needed.
Legal Compliance: Data that is collected and used in accordance with legal regulations.

High-Quality Data

So, how do we obtain high-quality data?

Data Preprocessing

Data Preprocessing

Consists of 6 stages:

- Data Collection
- Data Operations
- Modeling and Development
- Deployment
- Results Collection
- Validation

Transitions between steps do not have to be sequential.

An Overview of Data Preprocessing

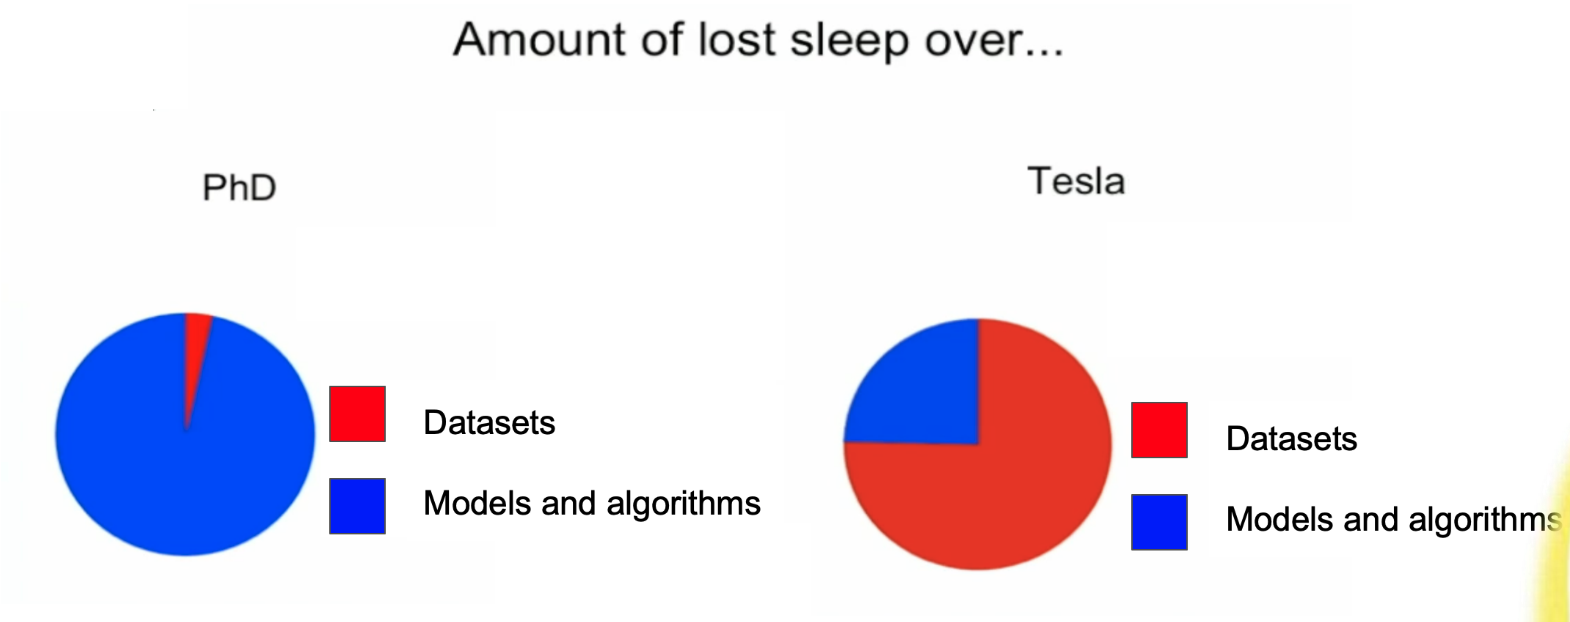

Slide from Andrej Karpathy, Tesla Director of AI (2021)

Feature Engineering

What is Feature Engineering?

Feature engineering is the process of preparing datasets to be used in machine learning models.

In this process, more meaningful and modeling-relevant features are extracted, selected, transformed, and created from raw data.

> *“ Coming up with features is difficult, time-consuming, requires expert knowledge. ‘Applied machine learning’ is basically feature engineering.”*
>
> *Andrew NG*

What is Feature Engineering?

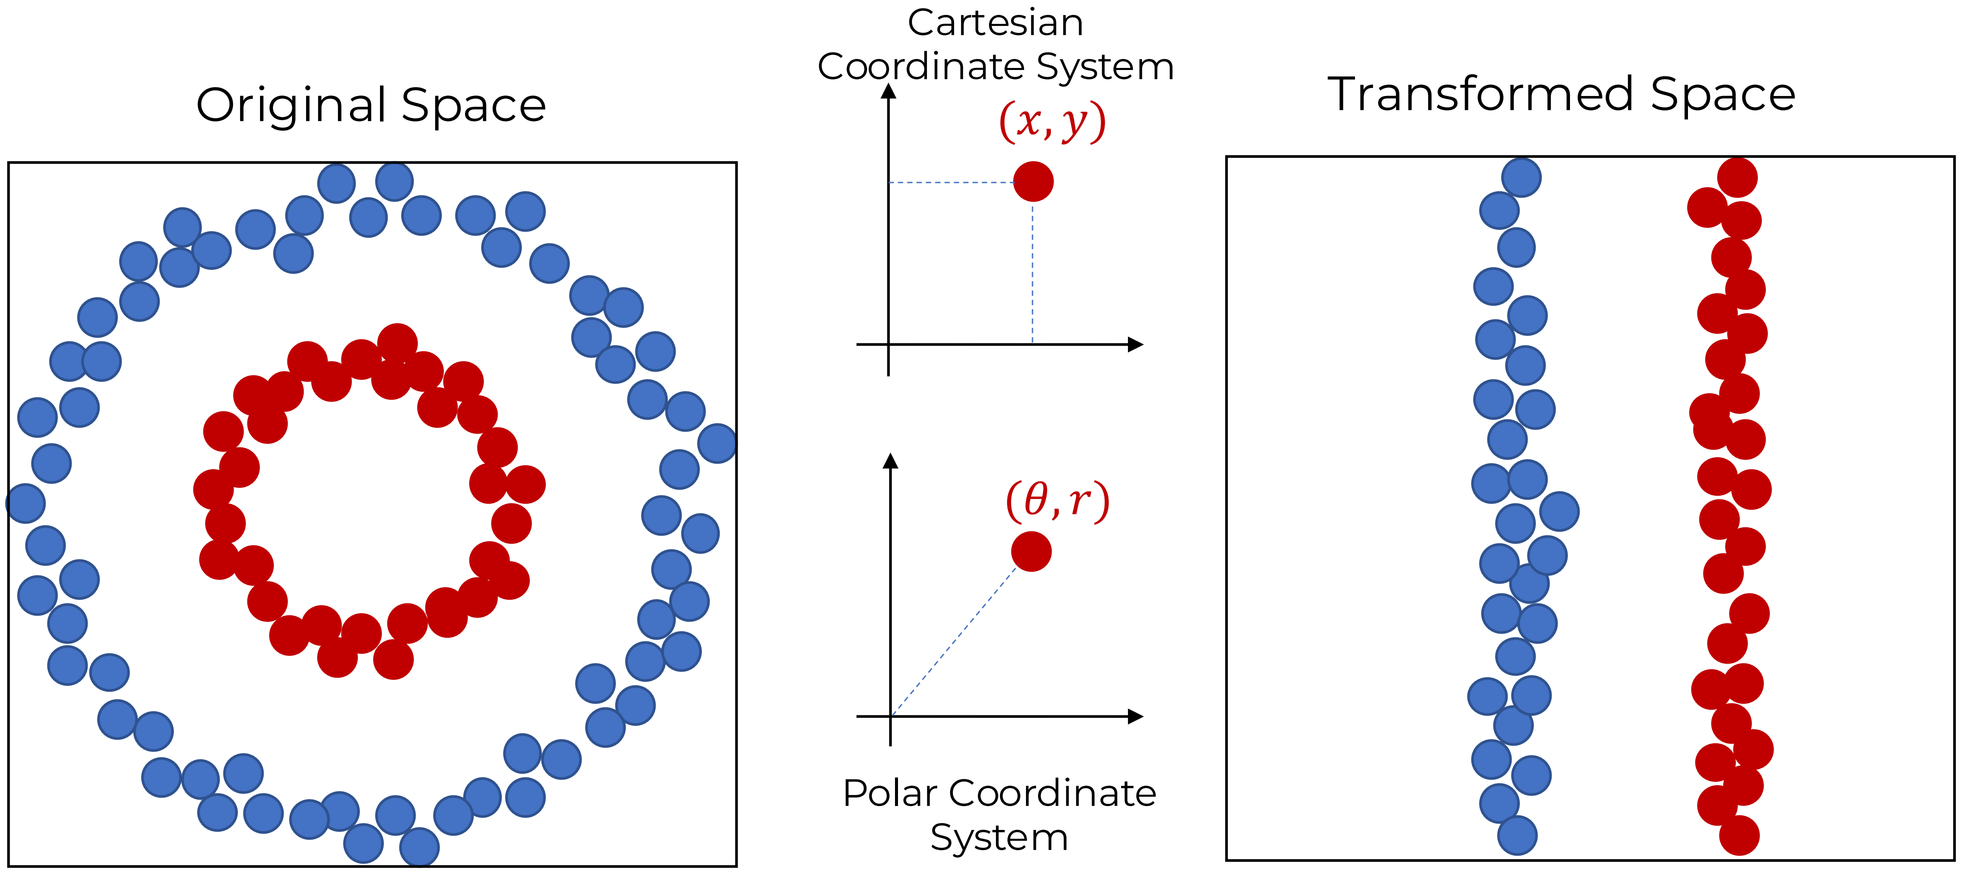

What is Feature Engineering?

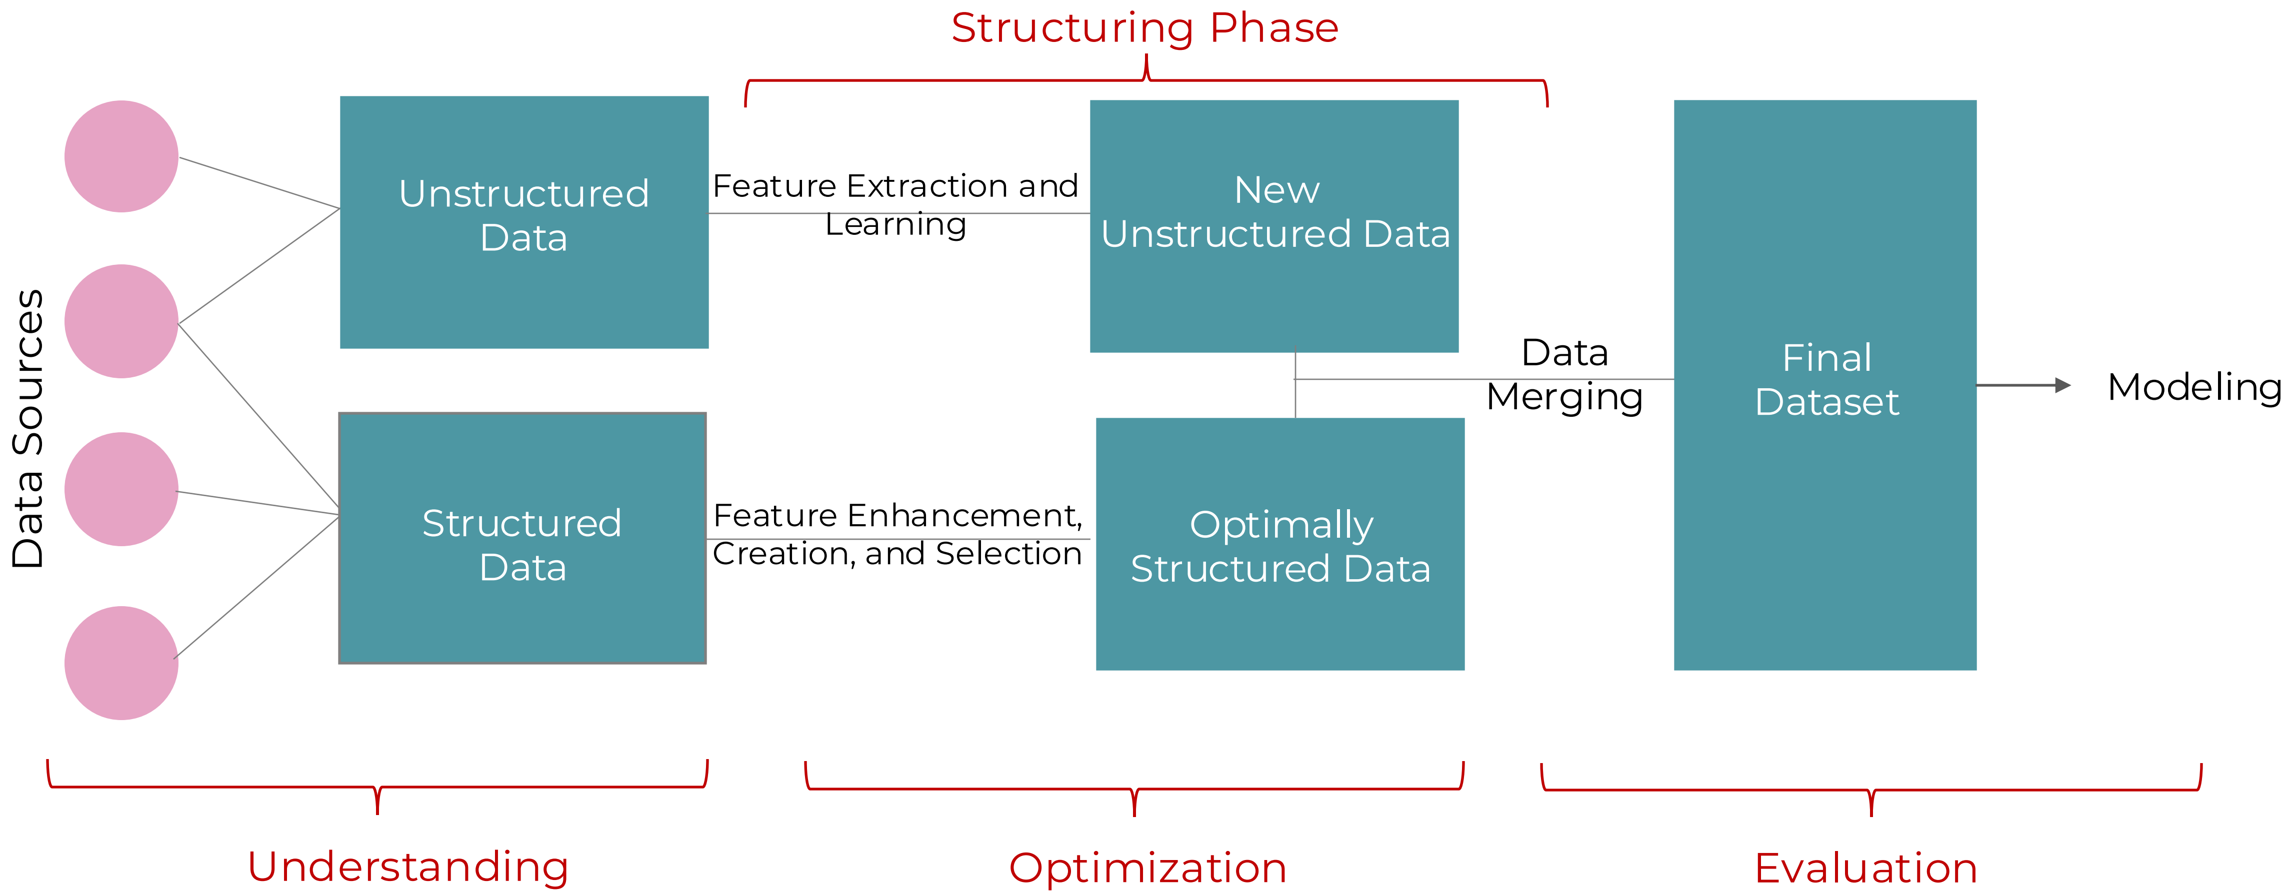

Good Practice Examples

Good Practice Examples

TÜRK BEYİN PROJESİ

Turkish Brain Project

SAYISAL GÖZ the digital eye

The Digital Eye - Mammography

Dijital Genç YAPAY ZEKÂ EKOSİSTEMİ

Digital Youth Artificial Intelligence Ecosystem

AI Transformation in the Public Sector (KAYZED)

Thank You