Chapter 1: Introduction to Artificial Intelligence (AI) (1 week)
1. Definitions and fields of AI application.
2. Historical evolution of AI.
3. Introduction to the major areas:
• Machine Learning
• Deep Learning

Chapter 1: Introduction to Artificial Intelligence (AI) (Duration: 1 Week)

1.1 Introduction

Artificial Intelligence (AI) has become one of the most influential and rapidly developing fields in modern science and technology. From healthcare and transportation to education and industry, AI systems are increasingly integrated into everyday life, transforming the way humans interact with machines and make decisions. Understanding the fundamental concepts of Artificial Intelligence is therefore essential for students in computer science, engineering, and related disciplines. This chapter provides an introductory overview of Artificial Intelligence. It begins by defining AI and presenting its main fields of application, highlighting its importance in real-world problem solving. The chapter then traces the historical evolution of AI, from its early theoretical foundations to the modern era of data-driven and intelligent systems. Finally, the chapter introduces two major areas of AI Machine Learning and Deep Learning, which form the backbone of many contemporary AI applications.

1.1.1 Definition of Artificial Intelligence

Artificial Intelligence (AI) is a branch of computer science that focuses on the creation of systems capable of performing tasks that normally require human intelligence, such as learning, reasoning, problem solving, perception, and decision-making. The goal of AI is to design machines that can simulate, augment, or replicate aspects of human cognitive abilities. In the scientific literature, Artificial Intelligence (AI) has been defined in many different ways, reflecting the diversity of perspectives and objectives within the field. Some widely cited definitions include:

- **John McCarthy (1956):** Artificial Intelligence is the science and engineering of making intelligent machines.
- **Charniak and McDermott (1985):** The study of mental faculties through the use of computational models.
- **Poole et al (1998):** The design of intelligent agents.

- **Rich and Knight (1990):** The discipline that studies the possibility of making computers perform tasks at which humans are currently better than machines.
- **Bellman (1978):** The automation of activities associated with human reasoning, such as decision making, problem solving, and learning.
- **Winston (1992):** The study of the mechanisms that enable an agent to perceive, reason and act.
- **Russell and Norvig (1995):** Artificial Intelligence is the study of agents that perceive their environment and take actions to maximize their chances of success.
- **Nilsson (1998):** The study of entities that exhibit intelligent behavior.

☞ **Key idea for students:**

AI aims to **simulate or augment human intelligence** using machines.

1.1.2 Subdomains of Artificial Intelligence

As one might expect, creating intelligent agents is not a simple task. For this reason, Artificial Intelligence has been divided into several subdomains, each of which addresses a specific aspect of the overall problem. The main subdomains of AI include the following:

- **Knowledge Representation and Automated Reasoning:** As the name suggests, this area of AI focuses on how knowledge can be represented in a form suitable for computation. Such knowledge may be incomplete, uncertain, or even inconsistent. This subdomain also studies reasoning mechanisms that allow systems to draw conclusions and make decisions based on available knowledge.
- **General Problem Solving:** The objective of this subdomain is to develop general-purpose algorithms capable of solving a wide range of concrete problems, rather than being limited to a single specific application.
- **Natural Language Processing (NLP):** This subdomain aims to enable machines to understand, interpret, generate, and translate human language, whether written or spoken.
- **Computer Vision:** The goal of computer vision is to allow computers to understand and interpret visual information from images and videos, such as recognizing faces, objects, or characters.
- **Robotics:** Robotics focuses on the design and development of physical agents capable of perceiving their environment and acting within the real world. These systems combine perception, decision making, and control to perform tasks autonomously.
- **Machine Learning:** In this subdomain of AI, the objective is to design systems that can automatically improve their performance through experience. Machine learning enables programs to adapt and modify their behavior based on data.

Strong connections exist between these subdomains. For example, knowledge representation languages are often used as the foundation for expert systems, and many pattern recognition algorithms rely heavily on machine learning techniques. Moreover, Artificial Intelligence is closely related to other disciplines such as philosophy, psychology, neuroscience, cognitive science, linguistics, and economics.

1.2 Historical Evolution of Artificial Intelligence

The historical evolution of Artificial Intelligence (AI) began in the 1950s, progressing from early conceptual ideas to modern transformative systems. Key milestones include Alan Turing's foundational work, the 1956 Dartmouth Conference (which coined the term "artificial intelligence"), expert systems, the resurgence of machine learning in the 1990s, the deep learning revolution in the 2010s, and the rise of Generative AI in the 2020s.

Before we explore **Generative AI**, let's go back to **before 1950**, when the concept of Artificial Intelligence had not yet emerged:

- **1939–1945:** During **World War II**, the **Enigma machine** was used by **Nazi Germany** to encrypt **military** communications.
- **1940–1942: Alan Turing**, at **Bletchley Park**, developed mathematical and computational methods to break Enigma codes. The creation of the **Bombe machine** automated the decryption process.
- **1942 onward:** Turing's work demonstrated that machines could perform **human-like reasoning tasks**, inspiring his later theoretical exploration of **Artificial Intelligence** in the **1940s**.
- **1943 – The First Mathematical Model of a Neural Network:** In **1943**, **Warren McCulloch and Walter Pitts** introduced the first mathematical model of a neural network. Their model formalized how networks of simplified neurons could perform basic logical functions, laying the foundation for future research in artificial neural networks and machine learning.

1) Birth of AI (1950–1956)

- **Artificial Intelligence** was first proposed by **Alan Turing** in the late **1940s**. In **1950**, he asked, *"**Can machines think**?"* and introduced the **Turing Test** (or ***Imitation Game***), in which a person must determine whether they are interacting with a human or a machine. The test aimed to evaluate machine intelligence and whether a machine could exhibit human-like behavior, sparking widespread interest and research in AI.
- In **1956**, **John McCarthy** officially coined the term *"**artificial intelligence**"* at the **Dartmouth** Workshop, establishing **AI** as a formal field of study.

- Significant progress followed in the **1950s** and **1960s** as machines became faster, cheaper, and capable of storing data. In **1959**, **Arthur Samuel** introduced **Machine Learning** with a self-learning checkers program. Around the same time, **Frank Rosenblatt** developed the first operational **perceptron**, an early neural network for training computers from data.

2) Early Enthusiasm and AI Winters (1960s–1980s)

- **First Generative AI Model ELIZA** developed in the **1960s** by **Joseph Weizenbaum**, **ELIZA** was one of the first generative **AI chatbots**. It used simple keyword patterns to mimic a psychotherapist, demonstrating early natural language processing. While it could communicate convincingly, **ELIZA** had no real intelligence and served more as a simulation than a true AI.
- Overhyped expectations and limited computing power led to periods of reduced funding and interest, known as the ***AI winters***, during the **1970s** and late **1980s**.

3) Expert Systems and Machine Learning (1980s–1990s)

- The **1980s** saw the rise of expert systems designed to emulate human expertise.
- In the **1990s**, the focus shifted toward machine learning.
- A major milestone in the **1990s** was **IBM's *Deep Blue* defeating world chess champion Garry Kasparov in 1997**. This achievement demonstrated the remarkable potential of **data-driven approaches** and computational power in solving complex, strategic problems. Unlike traditional AI systems based on hand-coded rules, Deep Blue relied on advanced algorithms and brute-force search to evaluate millions of possible moves, highlighting how combining machine learning, search strategies, and high-performance computing could surpass human expertise in specific domains.

4) Deep Learning Revolution (2000s–2010s)

- The availability of large datasets (***Big Data***) and increased computing power fueled the growth of deep learning and neural networks.
- Breakthroughs included advanced image recognition (***ImageNet, 2012***) and natural language processing.
- Systems such as IBM's ***Watson*** demonstrated the ability to understand and process complex language tasks.

5) Generative AI and the Modern Era (2020–Present)

- The **2020s** have seen the emergence of large-scale, pre-trained transformer models, giving rise to Generative AI capable of producing text, images, and code.
- Current trends emphasize ***agentic AI***, providing systems with autonomy to achieve complex goals and make independent decisions.

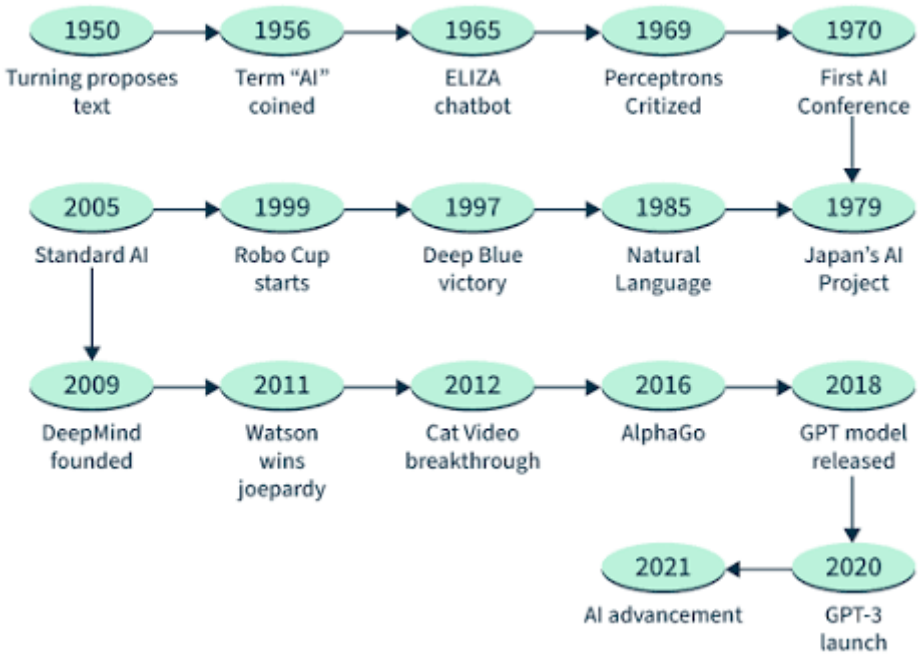

Figure 1.1: Evolution of AI.

Generative AI has rapidly evolved since the mid-2010s and is now considered one of the most significant advancements in technology. Its growth continues at an extraordinary pace, making it essential for organizations and societies to stay informed about the latest developments. With this momentum, Generative AI promises even greater innovations in the near future.

1.3 Major Areas of AI: Machine Learning and Deep Learning

Artificial Intelligence (AI) encompasses a wide range of techniques and approaches. Two of its most important and widely applied areas are **Machine Learning (ML)** and **Deep Learning (DL)**. Both focus on enabling computers to learn from data, but they differ in complexity, methodology, and applications.

1.3.1 Machine Learning (ML)

Machine Learning is a subset of AI that allows systems to **learn patterns from data** and make predictions or decisions without being explicitly programmed for every task. ML algorithms can be **supervised** (learning from labeled data), **unsupervised** (finding patterns in unlabeled data), or **reinforcement-based** (learning through trial and error).

Examples of Machine Learning:

- Spam email detection
- Predicting stock market trends
- Recommender systems (Netflix, Amazon)

1.3.2 Deep Learning (DL)

Deep Learning is a specialized branch of Machine Learning that uses **artificial neural networks** with multiple layers (hence "deep") to model **complex, high-dimensional patterns**. DL is particularly effective for tasks such as image and speech recognition, natural language processing, and autonomous driving.

Examples of Deep Learning:

- Image recognition (e.g., identifying objects in photos)
- Voice assistants (e.g., Siri, Alexa)
- Generative AI (e.g., ChatGPT, DALL·E)

1.3.3 Relationship between ML and DL

Deep Learning is a subset of Machine Learning. While ML often relies on manually engineered features, DL can automatically extract features from raw data, making it highly effective for complex tasks that traditional ML struggles with.

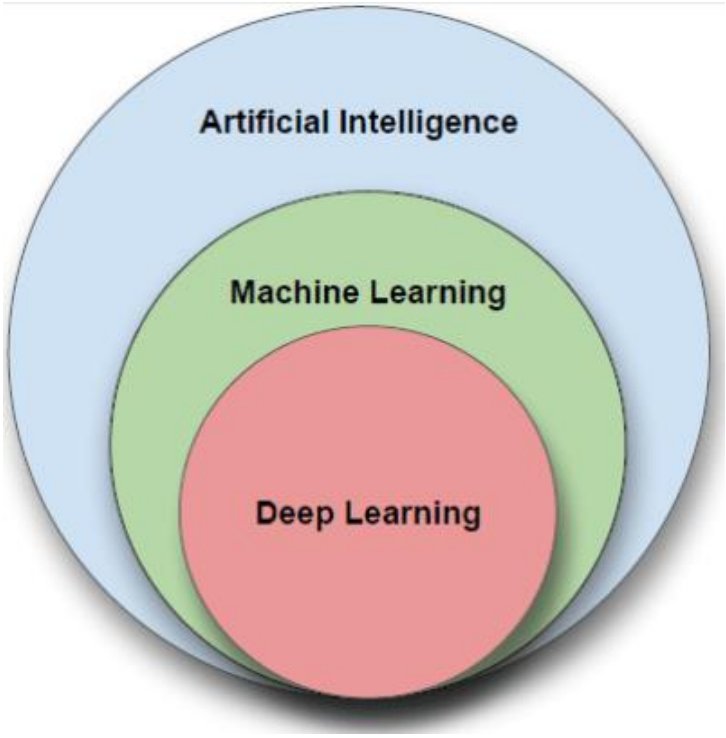

Figure 1.2: Relation between Machine Learning Deep Learning.

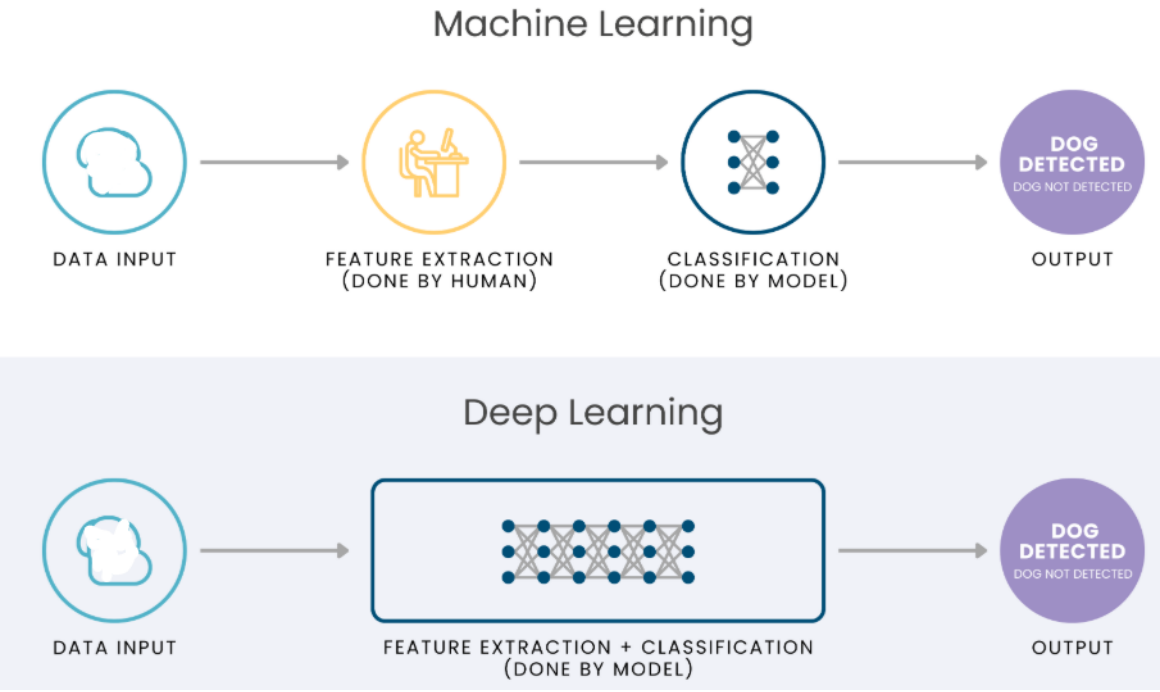

Figure 1.3: Machine Learning (ML) vs. Deep Learning (DL).

In traditional machine learning, a human person manually identifies features and labels data *before* feeding it to the ML model. From there, the ML models rely on straight-forward algorithms such as decision trees to learn to identify patterns. Deep learning models are different in that they're designed to mimic the way the human brain processes information using artificial neural networks known as "deep neural networks" that doesn't require a human for manual feature extraction.

1.4 Real-World Applications

1. Machine Learning (ML) applications

- Email Spam Filtering (e.g. Automatically detecting and filtering spam emails using algorithms trained on labeled data)
- Customer Service Chatbots (e.g. Providing automated responses to customer inquiries based on learned patterns)
- Speech Recognition (e.g. Transcribing spoken language into text in apps like Google Voice)
- Credit Scoring (e.g. Assessing the creditworthiness of loan applicants through predictive modeling)

2. Deep Learning (DL) applications

- Image Recognition (e.g. Identifying objects, faces, and scenes in photos and video)
- Natural Language Processing (NLP) (e.g. Google Translate)

- Autonomous Driving (e.g. Tesla, Cruise)
- Voice Assistants (e.g. Alexa, Siri)

3. Artificial Intelligence (AI) applications

Encompasses all ML and DL applications above and anything else that performs tasks previously requiring human intelligence, including:

- Rule-Based Chatbots: Simple AI chatbots that operate based on predefined rules and decision trees, rather than learning from data.
- Automated Decision-Making Systems: AI systems used in legal or administrative settings to apply rules and regulations without learning from data (e.g., tax preparation software).
- Game AI: Non-learning AI used in video games to control non-player characters (NPCs) using scripted behaviors and predefined rules.

I.5 Conclusion

In this chapter, we introduced the fundamental concepts of Artificial Intelligence and examined its role in modern technology. We explored the definition of AI and its wide range of applications across various domains, demonstrating its growing impact on society. The historical evolution of AI was also discussed, illustrating how advances in computing power, data availability, and algorithms have shaped the development of intelligent systems.

Furthermore, the chapter presented an overview of the major areas of Artificial Intelligence, with particular emphasis on Machine Learning and Deep Learning. These approaches enable machines to learn from data and solve complex problems that were previously considered exclusive to human intelligence. This foundational understanding prepares students for deeper exploration of AI techniques and algorithms in the following chapters.

Reference

S. J. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*. Englewood Cliffs, NJ, USA: Prentice Hall, 1995.