

ETHICS AND CURRENT ISSUES IN TECHNOLOGY



*Marshal Miller, Jeffrey Armstrong, Alexandra Briggs
and Abigail Michelini*

Ethics and Current Issues in TechnologyNorthampton
Community College

Ethics and Current Issues in Technology

This text is disseminated via the Open Education Resource (OER) LibreTexts Project (<https://LibreTexts.org>) and like the thousands of other texts available within this powerful platform, it is freely available for reading, printing, and "consuming."

The LibreTexts mission is to bring together students, faculty, and scholars in a collaborative effort to provide an accessible, and comprehensive platform that empowers our community to develop, curate, adapt, and adopt openly licensed resources and technologies; through these efforts we can reduce the financial burden born from traditional educational resource costs, ensuring education is more accessible for students and communities worldwide.

Most, but not all, pages in the library have licenses that may allow individuals to make changes, save, and print this book. Carefully consult the applicable license(s) before pursuing such effects. Instructors can adopt existing LibreTexts texts or Remix them to quickly build course-specific resources to meet the needs of their students. Unlike traditional textbooks, LibreTexts' web based origins allow powerful integration of advanced features and new technologies to support learning.



LibreTexts is the adaptable, user-friendly non-profit open education resource platform that educators trust for creating, customizing, and sharing accessible, interactive textbooks, adaptive homework, and ancillary materials. We collaborate with individuals and organizations to champion open education initiatives, support institutional publishing programs, drive curriculum development projects, and more.

The LibreTexts libraries are Powered by [NICE CXone Expert](#) and was supported by the Department of Education Open Textbook Pilot Project, the California Education Learning Lab, the UC Davis Office of the Provost, the UC Davis Library, the California State University Affordable Learning Solutions Program, and Merlot. This material is based upon work supported by the National Science Foundation under Grant No. 1246120, 1525057, and 1413739.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation nor the US Department of Education.

Have questions or comments? For information about adoptions or adaptations contact info@LibreTexts.org or visit our main website at <https://LibreTexts.org>.

The LibreTexts name, logo, and book covers are trademarked by LibreTexts, Inc. (a 501(c)(3) not-for-profit organization) and protected against unauthorized use.

This text was compiled on 08/24/2026

TABLE OF CONTENTS

Licensing

1: Introduction

- 1.1: Chapter Introduction
- 1.2: Machines, AI, and Humans
- 1.3: Types of AI
- 1.4: Applications of AI
- 1.5: Emerging Technology Trends

2: The History of AI

- 2.1: The Early Foundations (Pre-1950s)
- 2.2: The Birth of AI (1950s – 1960s)
- 2.3: The AI Winters and Specialization (1970s – 1980s)
- 2.4: Machine Learning and the Data Revolution (1990s – 2000s)
- 2.5: The Deep Learning Era (2010s)
- 2.6: The Age of Generative AI (2020s – Present)

3: The Current State of AI

- 3.1: AI in Everyday Life

3.2: AI Transforming Healthcare

- 3.2.1: Benefits and Concerns

- 3.3: AI Transforming Education

3.3.1: Pros and Cons

- 3.3.1.1: Benefits to Students
- 3.3.1.2: Benefits for Teachers

- 3.4: AI in Marketing Strategies

3.4.1: Good and Bad

- 3.4.1.1: AI for Entrepreneurs

- 3.5: Privacy and Safety Challenges with AI

- 3.6: Summary

3.7: Assignments

- 3.7.1: Case Study 1 - AI for Personal Use
- 3.7.2: Case Study 2 - AI in Education
- 3.7.3: Case Study 3 - AI for Entrepreneurs
- 3.7.4: Knowledge Check

4: AI and Ethics

5: Environmental Impacts of AI

6: AI in the Workspace

7: Legal Aspects of AI

7.1: Introduction

7.2: Copyright and Image Generation

7.3: What's the big deal?

7.4: Writing, Film, and Music

7.5: Case Study- Zarya of the Dawn

7.6: Teaching AI Ethics

7.7: Where does all that Data come from?

7.8: Protecting Personal Privacy

7.9: Case Study- AI Defamation

7.10: Teaching AI Ethics (1)

8: AI and Society

9: The Convergence of Thinking

10: AI and Politics

11: AI and Humanity

12: Summary

Detailed Licensing (1) (1)

Detailed Licensing

Glossary

Index

Licensing

A detailed breakdown of this resource's licensing can be found in **Back Matter/Detailed Licensing**.

This page titled [Licensing](#) is shared under a [CC BY-NC-SA 4.0](#) license and was authored, remixed, and/or curated by LibreTexts.

1: INTRODUCTION

1.1: Chapter Introduction

1.2: Machines, AI, and Humans

1.3: Types of AI

1.4: Applications of AI

1.5: Emerging Technology Trends

This page titled [1: Introduction](#) is shared under a [CC BY-NC-SA 4.0](#) license and was authored, remixed, and/or curated by Dr. Marshal Miller - Northampton Community College.

1.1: Chapter Introduction

Learning Objectives

Upon successful completion of this chapter, you will be able to:

- Define Artificial Intelligence;
- Compare machines, humans, and AI;
- Distinguish between the types of AI and learning methods;
- Differentiate between the types of extended reality;
- List some of the emerging technology trends and their impact on business;

Information systems have evolved at a rapid pace ever since their introduction in the 1950s. Today devices that you can hold in one hand are more powerful than the computers used to land a man on the moon in 1969. The last 10 years has seen the proliferation of intelligent devices, devices that can process information and make suggestions or provide responses. How does Netflix seem to know what types of programs or movies we like? How does Amazon post product displays to our account that sparks our interest? How does YouTube seem to provide video feeds that align with what we have previously watched and protect younger viewers from harmful content? This is achieved through artificial intelligence using algorithms and machine learning.

Artificial intelligence, the ability of a computer or machine to think and learn, and mimic human behaviour is changing the way companies do business. AI uses software algorithms to simulate human intelligence processes such as reasoning and speaking within computers and other IoT devices. AI simulates human intelligence processes through robotics, intelligent agents, expert systems, algorithms, and natural-language processing.

This chapter will provide an introduction to artificial intelligence and provide some examples of how businesses are leveraging this technology. It will also explore other types of emerging technology. **Emerging technology** includes new technology and technology that is continuously evolving.

1.2: Machines, AI, and Humans

Current developments in intelligent machines focus on mimicking human behavior to perform complex tasks. Applications of AI are proliferating and being integrated into our everyday lives. Explore the examples in the activity below. Which of these examples of everyday technology incorporate AI?

Interactive Content

This interactive element is not available in this format. Visit the link or scan the QR code to access it online.

[View this Interactive Content online](#)



Machines vs AI

Machines and AI systems may appear to be the same, but their systems and mode of operation are vastly different.

- Machines that incorporate AI are typically capable of analyzing data, and they appear to “make decisions”, similar to humans.
- Regular machines simply execute a series of commands and are unable to improve based on the data they receive.

Example: Machine vs AI

A great example of the difference between machines and AI systems is the platform [Grammarly](#) versus your typical spell check. While a spelling checker is able to correct your spelling, grammar, and punctuation, Grammarly can do all that, as well as check for issues in sentence structure, misused words, and more through the power of AI.

A key distinction between a computer following the code vs an AI system executing this is that with AI, the algorithm would become smarter as more data is entered and the number of errors would decrease over time. AI models are created through **algorithms**, which are a set of rules or processes to solve a specific problem or task.^[1]

Example: Algorithm

Imagine you were given some numbers and were tasked with developing an algorithm that sorts the numbers from lowest to highest. What would be the set of steps that you would need to follow to accomplish this task?

One solution would probably be to first define the range of numbers and set some criteria for the computer to follow. Next, write a specific set of instructions on how you would like to order the numbers. For example, you may want the computer to sort the numbers from lowest to highest. If there are several datasets, sometimes a **loop**, which is a part of code that allows a command to run again and again, would be used to ensure the computer completes the command over and over until all datasets are sorted. The algorithm would then execute the commands one by one, and any errors would be attributed to a mistake in the code.

Interactive Content

This interactive element is not available in this format. Visit the link or scan the QR code to access it online.

[View this Interactive Content online](#)



AI vs Humans

The distinction between AI and simple machines has been made, but how is AI different from humans? AI and Humans are different in a number of ways. Apart from humans being biological, live organisms, a key difference is the type and level of intelligence we possess.^[2] We are able to make decisions based on our intelligence, emotions, self-awareness, and creativity, whereas most machines simply perform tasks based on code written by programmers. The main goal behind the advancement of machines and computers is to make these systems more efficient and “smarter” to support humans, which is now being accomplished through AI.

However, tasks that involve emotional intelligence or intuition cannot be automated. A machine will make decisions based on facts and statistical data but cannot, for example, recognize emotions, and thus would not be able to make rational decisions. AI lacks the “Human Factor”. The table below compares natural (human) and artificial intelligence for a range of different abilities on a scale of achievement (low to high).

Ability	Natural	Artificial
Creativity and imagination	high	low
Learn from experience	high	med
Use sensors	high	med
Adapt to new situations	high	med
Perform complex calculations	high	high
Transfer information	high	high
Perform calculations accurately and efficiently	high	high

^[3]

1. de Ponteves, H. (2019). AI Crash Course: A fun and hands-on introduction to machine learning, reinforcement learning, deep learning, and artificial intelligence with Python. Birmingham: Packt Publishing Ltd.
2. Vadapalli, P. (2020, September 15). *AI vs Human Intelligence: Difference Between AI & Human Intelligence*. upGrad blog. <https://www.upgrad.com/blog/ai-vs-human-intelligence/>.
3. Stair, R. & Reynolds, G. (2017). Principles of Information Systems. Cengage.

This page titled [1.2: Machines, AI, and Humans](#) is shared under a [CC BY-SA 4.0](#) license and was authored, remixed, and/or curated by [Patrick McClanahan](#).

1.3: Types of AI

AI systems are different from humans and machines, and the degree of variation helps in classifying the different types of AI. AI can be classified based on ability, and functionality. There are three categories of AI based on ability: narrow AI, general AI and super AI.

Types of AI: Based on Ability

[Categories of Artificial Intelligence Based on Ability are narrow AI, general AI and super AI](#)

Figure 1.3.1: Categories of Artificial Intelligence Based on Ability

Ability of AI	Description
Artificial Narrow Intelligence (A.N.I)	A form of weak AI refers to a computer's ability to perform a single task well. This is the type of AI we see today. Natural language processing used in smart assistants like Siri and Google are examples of narrow AI.
Artificial General Intelligence (A.G.I)	A form of strong AI that mimics human intelligence. It means that machines have the ability to apply what they have learned across different tasks, take in new information, and apply reason.
Artificial Super Intelligence (A.S.I)	A form of supernatural AI that exceeds human capabilities.

Despite having developed complex machines and applications such as Siri, Alexa, or other virtual assistants, all AI systems that currently exist are classified as A.N.I. According to a general consensus among researchers, A.G.I and A.S.I are still decades away.^[1]

Example: Chatbots

Chat-bots and/or virtual assistants (i.e. Siri, and Alexa) are computer programs that use AI and natural language processing (NLP) to understand customer questions and automate responses to them, simulating human conversation.^[2] These bots are cloud-based and able to gather information, such as user preferences or browsing history, to provide solutions that are specific to each individual. Chat-bots improve the individual customer experience as the Cloud is able to provide personalized and relevant information.

In the hospitality industry, Best Western in collaboration with IBM Watson Advertising managed to use AI and NLP to create a more personalized customer experience for those looking to book Best Western hotels for their holidays. Specifically, "Conversations", an AI-powered advertising model from IBM, was used to provide one-on-one connections and recommend options based on the user's intent.^[3] For example, Conversations was able to provide users with ads containing information on their destination, and other tips and tricks!

Types of AI: Based on Functionality

AI can also be categorized based on functionality: reactive machines, limited memory, theory of mind, and self-awareness.

[Categories of Artificial Intelligence Based on functionality](#)

Categories of Artificial Intelligence Based on functionality include reactive machines, limited memory, theory of mind and self awareness

Function of AI	Description of Function
Reactive Machines	AI that conducts tasks for specific objectives. It is the first step of any AI system where no learn Garry Kasparov.

Function of AI	Description of Function
Limited Memory	Is the ability of AI to use past experience and information to make predictions.
Theory of Mind	AI with the awareness that others have their own beliefs, objectives, and intentions. This type of
Self-Aware	AI with the sense of self and consciousness. This type of AI is not yet developed.

Types of AI: Based on Application

There are a number of applications of AI ranging from

Application	Description of Application
Automation	The use of machines, control systems and information technologies to optimize productivity in the production of goods and delivery of services. ^[1]
Robots	The use of machines to perform tasks traditionally done by humans.
Intelligent Agents	The use of software to perform specific tasks.
Expert Systems	Aim to emulate the human ability to make decisions in specific contexts.
Machine Learning	Uses data and algorithms to imitate the way humans learn.
Deep Learning	Uses a process that replicates the human brain in data processing and also creates patterns for decision making.
Natural Language Processing	Allows computers to understand and communicate in human language.

Many of these applications of AI will be explored in further detail in the following sections.

1. automation. (n.d.). Definitions.net. <https://www.definitions.net/definition/automation>.
2. Dilmegani, C. (2021, Aug 11). Will AI REACH singularity By 2060? 995 experts' opinions on AGI," AIMultiple, 19-Jun-2021. <https://research.aimultiple.com/artificial-general-intelligence-singularity-timing/> ↵
3. Ibm.com. (n.d.) What is a chatbot? <https://www.ibm.com/topics/chatbots> ↵
4. Ibm.com. (2021). IBM Watson Advertising Conversations - Overview. <https://www.ibm.com/products/watson-advertising-conversations>

1.4: Applications of AI

Autonomous Technology

One of the most widely used applications of AI is autonomous technologies. By combining software, sensors, and location technologies, devices that can operate themselves to perform specific functions are being developed. Some examples include: medical nanotechnology robots (nanobots), self-driving cars, or unmanned aerial vehicles (UAVs).

A **nanobot** is a robot whose components are on the scale of about a nanometer, which is one-billionth of a meter. While still an emerging field, it is showing promise for applications in the medical field. For example, a set of nanobots could be introduced into the human body to combat cancer or a specific disease.

A **UAV**, often referred to as a “drone,” is a small airplane or helicopter that can fly without a pilot. Instead of a pilot, they are either run autonomously by computers in the vehicle or operated by a person using a remote control. While most drones today are used for military or civil applications, there is a growing market for personal drones.

Robots



Figure 1.4.1: [Photo](#) (Phasmatisnox via Wikipedia CC-BY)

Robots are automated machines that can execute specific tasks with very little or no human intervention and are able to accomplish tasks with both speed and precision ^[1]. The development and deployment of robots is most common in manufacturing in the replacement of humans in repetitive tasks. Robots are also used in medicine, education, restaurants and hotels, and entertainment. Some of the most popular robots are [ASIMO](#) by Honda and Boston Dynamics robots like [ATLAS](#). Robots can increase productivity and accuracy, but are costly. For more advantages and disadvantages, see the table below.

Robot Advantages	Robot Disadvantages
• Reduction of risk. Tasks that are dangerous to humans can be done by robots (i.e. defusing bombs, travelling to space). • Error reduction and increased level of precision. • No downtime. Robots can work everyday all day. • Good at performing repetitive tasks that may otherwise injure humans • Faster decisions can be made by robots programed with artificial intelligence than humans as they are able to process complex data at a much higher speed	• Reduction in specific employment opportunities in jobs being replaced by robots • High capital investment required by organizations to utilize and maintain • Lack of creativity and emotion in the performance of tasks or decision-making processes • Reduction in human interaction • Require constant power to run and maintain

With artificial intelligence, robots will be able to independently estimate the events around them and to make decisions on actions which they're required to make for their given objective. Robots can already both record human movement skills and replicate them as machine learning improves drive efficiency and mobility. [\[2\]](#)

Intelligent Agents



Figure 1.4.1: IOS 7 Siri (Kavakavins111 on Wikipedia CC-BY-SA)

Intelligent agents process the inputs it receives, and makes decisions/ takes action based on that information. Sensors allow intelligence agents to perceive their surroundings, while actuators allow them to act on that perception. A software agent, a human agent, and a robotic agent are all examples of agents, each with its own set of sensors and actuators to sense its environment and then perform actions through its actuators.

AI agents that can learn from their previous experiences or have learning capabilities are known as learning agents. Beginning by acting on basic knowledge, learning agents subsequently learn to act and adapt on their own automatically. Conversational intelligence tools are tools that have the ability to auto-record meetings, transcribe, and apply AI to speech. They can be helpful for individuals working in sales, and could possibly replace some functions performed by customer relationship management systems.

Siri, Alexa and Google Assistant are examples of intelligent agents that can be accessed from smartphones. These tools can allow users to open applications, send messages, make calls, play voicemails and check the weather.

Natural Language Processing (NLP)

Natural Language Processing (NLP) allows computers to extract meaning from human language. Natural Language Processing's goal by design is to read, decipher, and comprehend human language. In Natural Language Processing, humans converse with the machine which records the conversation and converts the audio to text. After that, the system analyses the data and converts it to audio. The machine then plays the audio file back in response to the human.

Algorithms use natural language processing to detect natural language rules, resulting in the conversion of unstructured language input into a format that computers can recognize. Natural Language Processing (NLP) is used in many translation applications such as Google Translate, various interactive voice response (IVR) applications found in call centres, and software tools that check for grammatical accuracy of texts, like Microsoft Word and Grammarly. Additionally, it is also used in personal assistant tools like Siri, Cortana and Alexa.

Expert Systems

Expert systems (ES) are designed to emulate the human ability to make decisions in specific contexts. An expert system, by means of a reasoning, can perform at a level comparable to or better than a human expert does within a specified area. The goal of ES is therefore to solve complex problems by following a line of reasoning that is derived from human knowledge. This reasoning is normally represented by if-then-else statements, instead of conventional procedural code. Expert systems have been reliably used in the business world to gain competitive advantages and forecast market conditions. In a time where every decision made in the business world is critical for success; the assistance provided from an expert system can be an essential and highly reliable resource.

Expert Systems in Fraud Detection

Expert systems can be used in fraud detection. A company can establish rules for detecting fraud and then allow the system to identify fraudulent behaviour based on the set of rules created. For example, a rule can be created to look for transactions that are originating from a particular location (one that may be known for criminal activity) and are over a certain amount. Since expert systems are based on rules that are programmed, it may be easy for cybercriminals to circumvent the system. However, expert systems can still be useful as part of a larger fraud defense strategy that may also incorporate more intelligent systems ^[3]

-
1. TechTarget. (2021, May 14). What are robots and how do they work? Retrieved January 12, 2022, from <https://www.techtarget.com/searchent...finition/robot> ↵
 2. Evgeniy Bryndin, Robots with Artificial Intelligence and Spectroscopic Sight in Hi-Tech Labor Market, International Journal of Systems Science and Applied Mathematics. Vol. 4, No. 3, 2019, pp. 31-37. doi: 10.11648/j.ijssam.20190403.11 ↵
 3. Lu, C. (2017, February 19). How AI is helping detect fraud and fight criminals. Retrieved December 3, 2021, from <https://venturebeat.com/2017/02/18/h...ght-criminals/>.

1.5: Emerging Technology Trends

In addition to AI; a number of other emerging technologies are forecast to have an impact on how businesses operate and communicate. The following are some trends that are driving these new and innovative technologies.

Wearable

The last section explored VR technologies, which include devices that can be worn to simulate virtual environments. Technology that can be worn or 'wearables' has been around for a long time, with applications such as hearing aids and, later, Bluetooth earpieces. Product lines have expanded to include the Smartwatch, body cameras, sports watch, and various fitness monitors and other wellness devices. Energy harvesting and haptic devices are examples of future developments in wearable technology. Energy harvesting allows body heat to be converted to solar power, and haptic devices allow one to control virtual objects. Wearable haptic devices will be integrated into clothing to help with directions, or to assist individuals navigate the virtual world ^[1]

Connected

As discussed in chapter 3, the Internet of Things (IoT) refers to devices that have been embedded into a variety of objects including appliances, lamps, vehicles, lightbulbs, toys, thermostats, jet engines, etc., and then connecting them via Wi-Fi, Bluetooth, or LTE to the internet. Think of IoT as devices that you wouldn't normally consider being connected to the internet, and the connection is independent of human intervention. This interconnectedness or uploading of data is virtually automatic. So a PC is not an IoT, but a fitness band could be. Interconnected devices are becoming ubiquitous, meaning they are everywhere. Today there are IoTs for monitoring traffic, air quality, soil moisture, bridge conditions, consumer electronics, autonomous vehicles, and the list seemingly never stops.

Principally three factors have come together to give us IoT: inexpensive processors, wireless connectivity, and a new standard for addresses on the internet known as IPv6. Processors have become both smaller and cheaper in recent years, leading to their being embedded in more devices. Consider technological advancements in your vehicles. Your car can now collect data about how fast you drive, where you go, radio stations you listen to, and your driving performance such as acceleration and braking. Insurance companies are offering discounts for the right to monitor your driving behavior. On the positive side, imagine the benefit of being informed instantly of anticipated traffic delays each time you adjust your route to work in the morning. Benefits from IoTs are virtually everywhere. Here is a quick list.

Optimization of Processes	IoT in manufacturing monitor a variety of conditions that impact production, including temperature, humidity, and barometric pressure – all factors that require adjustment in the application of manufacturing formulas.
Component Monitoring.	IoT is added to components in the manufacturing process, then monitored to see how each component is performing.
Home Security Systems.	IoT makes the challenge of monitoring activity inside and outside your home easier.
Smart Thermostats.	Remote control of home thermostats through the use of IoT allows the homeowner to be more efficient in the consumption of utilities.
Residential Lighting	IoT provides remote control of lighting.

While there are many benefits to these constantly connected devices privacy and security of personal data and information should also be considered.

Collaborative




Figure 1.5.1: [Waze Screen Shot](#) (D. Bourgeois [CC-BY-NC](#))

As more people use smartphones and wearables, it will be simpler than ever to share data with each other for mutual benefit. Some of this sharing can be done passively, such as reporting your location in order to update traffic statistics. Other data can be reported actively, such as adding your rating of a restaurant to a review site. The smartphone app [Waze](#) is a community-based tool that keeps track of the route you are traveling and how fast you are making your way to your destination. In return for providing your data, you can benefit from the data being sent from all of the other users of the app. Waze directs you around traffic and accidents based upon real-time reports from other users.

Super Computers

Quantum computing, a technology that applies quantum mechanics to build novel supercomputers, is a type of high-performance computer used to solve large-scale computational tasks. Quantum computers can operate at speeds that are exponentially faster than common computers and can make calculations based on the probability of an object's state before it is measured. Google demonstrated that its quantum computers can solve a problem that no classical computer could ever solve. For the third year in a row, IBM managed to double its quantum computing power. Additionally, several web service providers, including Amazon, announced plans for [cloud-based](#) quantum computing services. Quantum computers can make drug development, power storage, manufacturing, and agriculture better, faster, and more [sustainable](#). They may also unravel cybersecurity infrastructure around the world, making them a potential threat to national security. The field of quantum computing is still in its infancy, and the question of scale remains unsolved.

1. Perkoic, M. (2022, September 29). How Smart Wearables Are Shaping Our Future. Forbes.<https://www.forbes.com/sites/forbesbusinesscouncil/2022/09/29/how-smart-wearables-are-shaping-our-future/?sh=53a8fdde6b24>

This page titled [1.5: Emerging Technology Trends](#) is shared under a [CC BY-SA 4.0](#) license and was authored, remixed, and/or curated by [Patrick McClanahan](#).

2: THE HISTORY OF AI

2.1: The Early Foundations (Pre-1950s)

2.2: The Birth of AI (1950s – 1960s)

2.3: The AI Winters and Specialization (1970s – 1980s)

2.4: Machine Learning and the Data Revolution (1990s – 2000s)

2.5: The Deep Learning Era (2010s)

2.6: The Age of Generative AI (2020s – Present)

2.1: The Early Foundations (Pre-1950s)

The dream of creating artificial beings with human-like intelligence is as old as civilization itself. Long before the advent of contemporary computational systems, the conceptualization of artificial entities possessing cognitive capabilities permeated mythological narratives and literary works across cultures. Ancient Greek myths told of Talos, the bronze automaton that guarded Crete, while Jewish folklore spoke of the golem, a clay figure brought to life through mystical means. These stories reflected humanity's enduring fascination with the possibility of breathing life into inanimate matter.



Frankenstein: [Photo](#) (Tom McKinnon via Flickr CC BY-SA 2.0)

As the industrial age dawned and mechanical marvels became reality, writers like Mary Shelley explored the ethical implications of artificial creation in "Frankenstein" (1818), while Karel Čapek introduced the term "robot" in his 1920 play "R.U.R." These literary works didn't merely entertain—they posed profound questions about consciousness, identity, and what it truly means to be alive.

Throughout the twentieth century, the theoretical underpinnings emerged through rigorous mathematical formulations and logical frameworks. Mathematicians and philosophers began to seriously consider whether machines could think, setting the stage for what would become one of the most transformative fields in human history:

- **Ancient Mythological Traditions:** Narratives depicting automatons and artificially constructed beings proliferated throughout Greek and various other ancient mythological systems, establishing early conceptual precedents for artificial intelligence. From Hephaestus's golden servants in Homer's *Iliad* to the mechanical marvels described in ancient Chinese texts, these stories revealed humanity's persistent desire to create intelligent companions and servants.
- **Alan Turing (1950):** Published the seminal work "*Computing Machinery and Intelligence*," proposing the revolutionary "Imitation Game" (subsequently recognized as the Turing Test) as a definitive criterion for evaluating machine intelligence. Turing's elegant thought experiment asked a deceptively simple question: if a machine could engage in conversation so convincingly that a human interrogator couldn't distinguish it from another human, could we say that the machine was truly thinking?

This page titled [2.1: The Early Foundations \(Pre-1950s\)](#) is shared under a [CC BY-NC-SA 4.0](#) license and was authored, remixed, and/or curated by Dr. Marshal Miller - Northampton Community College.

2.2: The Birth of AI (1950s – 1960s)

1956: The Dartmouth Conference

In the summer of 1956, a small group of visionary researchers gathered at Dartmouth College for what would become a pivotal moment in scientific history. John McCarthy, a young mathematician with a passion for mechanical reasoning, joined forces with Marvin Minsky, Nathaniel Rochester from IBM, and Claude Shannon, the father of information theory. Together, they orchestrated a groundbreaking symposium that would formally launch artificial intelligence as a distinct academic discipline.

The conference proposal boldly declared their intention to explore "how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves." This ambitious gathering, though lasting only two months, established the foundational principles and research directions that would guide AI development for decades to come. McCarthy coined the term "artificial intelligence" itself, giving a name to humanity's quest to create thinking machines.

1950s – 1960s: Early Optimism

The years following Dartmouth were marked by infectious enthusiasm and remarkable breakthroughs that seemed to validate the researchers' boldest predictions. Allen Newell and Herbert Simon developed the **Logic Theorist**, recognized as the inaugural operational AI program, which could prove mathematical theorems by mimicking human reasoning processes. When the program successfully proved 38 of the first 52 theorems in Whitehead and Russell's "Principia Mathematica," it demonstrated that machines could engage in what appeared to be genuine logical thinking.

Joseph Weizenbaum's ELIZA program captured public imagination by conducting seemingly intelligent conversations, using simple pattern matching to respond to human input in a way that felt remarkably natural. Though ELIZA's methods were relatively simple, users often attributed human-like understanding to the program, revealing both the power of well-designed interfaces and the human tendency to anthropomorphize responsive systems.

During this golden age of early AI, researchers maintained optimistic projections that general human-level artificial intelligence was not just possible, but imminent. Herbert Simon famously predicted in 1965 that "machines will be capable, within twenty years, of doing any work a man can do." This confidence, while ultimately premature, drove unprecedented innovation and attracted significant funding to the field.

2.3: The AI Winters and Specialization (1970s – 1980s)

The First "AI Winter"

The euphoria of the 1960s gradually gave way to sobering reality as the 1970s progressed. The grand promises of human-level AI within decades began to ring hollow as researchers encountered fundamental challenges they hadn't anticipated. Early programs, while impressive in narrow domains, proved brittle and unable to handle the complexity and ambiguity of real-world problems. The "combinatorial explosion" problem became apparent—as problems grew more complex, the computational resources required grew exponentially, quickly overwhelming even the most powerful computers of the era.

The Lighthill Report in the UK and similar assessments in the United States delivered harsh critiques of AI's progress, leading to dramatic cuts in research funding. Government agencies and private investors, disappointed by the gap between promises and deliverables, redirected their resources elsewhere. This period, dubbed the first "AI Winter," forced the field into a period of introspection and strategic refocusing toward more narrowly defined, specialized computational tasks.

Yet this apparent setback proved beneficial in unexpected ways. Researchers learned valuable lessons about the complexity of intelligence and began developing more realistic approaches to AI development. The focus shifted from creating general intelligence to solving specific, well-defined problems—a strategy that would eventually prove more fruitful.

The Rise of Expert Systems (1980s)

YouTube Video

This interactive element is not available in this format. Visit the link or scan the QR code to access it online.

[View this YouTube Video online](#)



From the ashes of the first AI winter emerged a new approach that would revitalize the field and demonstrate AI's practical value. Expert systems represented a fundamental shift in strategy—instead of trying to replicate general human intelligence, researchers focused on capturing the specialized knowledge of human experts in narrow domains.

MYCIN, developed at Stanford University, exemplified this approach by diagnosing blood infections and recommending treatments with accuracy that often exceeded that of human doctors. DENDRAL could identify molecular structures from mass spectrometry data, while XCON helped configure computer systems for Digital Equipment Corporation, saving the company millions of dollars annually.

Artificial intelligence experienced a commercial renaissance through these sophisticated programs engineered to address domain-specific problems such as medical diagnostic procedures, financial planning, and industrial troubleshooting. Companies like Symbolics and Lisp Machines Inc. emerged to provide specialized hardware and software for AI applications, creating a thriving ecosystem around expert systems technology.

Despite their commercial success and practical utility, these systems fundamentally lacked the adaptive learning capabilities characteristic of human cognition. They were essentially sophisticated databases of rules and facts, unable to learn from experience or adapt to new situations outside their programmed expertise. This limitation would eventually contribute to a second AI winter, but not before expert systems had proven AI's commercial viability and established important precedents for knowledge representation and reasoning.

This page titled [2.3: The AI Winters and Specialization \(1970s – 1980s\)](#) is shared under a [CC BY-NC-SA 4.0](#) license and was authored, remixed, and/or curated by Dr. Marshal Miller - Northampton Community College.

2.4: Machine Learning and the Data Revolution (1990s – 2000s)

1997: Deep Blue

On May 11, 1997, the world watched in fascination as IBM's Deep Blue computer made history by defeating world chess champion Garry Kasparov in a six-game match. This wasn't just a victory for artificial intelligence—it was a watershed moment that demonstrated the raw computational power and sophisticated algorithmic approaches that machines could bring to complex strategic thinking.

Deep Blue's triumph represented the culmination of decades of research in game-playing algorithms and parallel processing. The system could evaluate 200 million chess positions per second, using specialized chess chips and advanced search algorithms to peer deeper into the game tree than any human player could manage. Kasparov, widely regarded as the greatest chess player of all time, found himself outmaneuvered by a machine that combined brute-force calculation with sophisticated positional evaluation.

The victory sent shockwaves through both the AI community and popular culture, reigniting public interest in artificial intelligence and demonstrating that machines could excel in domains requiring strategic thinking and pattern recognition. More importantly, it established a significant milestone in computational processing power and algorithmic search methodologies that would influence AI development for years to come.

The Shift to Data

As the new millennium approached, AI researchers began a fundamental paradigm shift that would transform the entire field. The rule-based expert systems of the 1980s, while successful in narrow domains, had proven too rigid and brittle for many real-world applications. A new generation of researchers began exploring probabilistic and statistical approaches that could learn patterns from data rather than relying on hand-coded rules.

This transition from manually programmed rule-based systems (symbolic AI) toward probabilistic statistical approaches and machine learning paradigms was driven by both theoretical advances and practical necessity. Researchers developed algorithms that could automatically discover patterns in data, adapt to new information, and handle uncertainty—capabilities that rule-based systems struggled to achieve.

The exponential growth of internet infrastructure provided unprecedented access to massive datasets, which subsequently became the fundamental resource fueling these innovative algorithmic approaches. Web pages, digital images, user interactions, and countless other forms of digital information created a treasure trove of training data that machine learning algorithms could exploit. Companies like Google began demonstrating the power of data-driven approaches, using vast amounts of web data to improve search results and introduce new services.

This data revolution coincided with advances in computational power and storage capacity, making it feasible to process and analyze datasets that would have been unimaginable just decades earlier. The combination of abundant data, improved algorithms, and increased computing power set the stage for the remarkable AI advances that would follow in the 2010s.

This page titled [2.4: Machine Learning and the Data Revolution \(1990s – 2000s\)](#) is shared under a [CC BY-NC-SA 4.0](#) license and was authored, remixed, and/or curated by Dr. Marshal Miller - Northampton Community College.

2.5: The Deep Learning Era (2010s)

2012: AlexNet

The year 2012 marked a revolutionary turning point in artificial intelligence history when Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton unveiled AlexNet at the ImageNet Large Scale Visual Recognition Challenge. This deep convolutional neural network didn't just win the competition—it obliterated the previous best results, reducing the error rate from 26% to 15% in a single leap that stunned the computer vision community.

AlexNet's breakthrough represented more than just improved accuracy; it demonstrated the transformative power of deep neural network architectures when combined with massive datasets and GPU computing power. The network's ability to automatically learn hierarchical features from raw pixel data—detecting edges at lower levels, shapes at middle levels, and complex objects at higher levels—showed that machines could develop their own internal representations of visual concepts without human guidance.

This achievement in computer vision capabilities demonstrated performance levels that surpassed human capabilities in image classification tasks, sparking a renaissance in neural network research that had been dormant since the previous AI winter. Suddenly, every major technology company was investing heavily in deep learning research, and universities worldwide were launching new AI programs to meet the exploding demand for expertise in this field.

Virtual Assistants

The 2010s witnessed the widespread mainstream adoption of AI-powered virtual assistants that brought artificial intelligence directly into consumers' daily lives. Apple's Siri, launched in 2011, pioneered the integration of sophisticated natural language processing and advanced voice recognition technologies into a consumer device that millions could carry in their pockets.

Siri's introduction marked the beginning of a new era where people could have natural conversations with their devices, asking questions, setting reminders, and controlling smart home devices through simple voice commands. The technology that made this possible combined automatic speech recognition, natural language understanding, dialogue management, and text-to-speech synthesis in a seamless user experience that felt almost magical to many users.

Google Assistant and Amazon's Alexa soon followed, each bringing their own innovations and capabilities to the virtual assistant landscape. These systems demonstrated how AI could enhance human productivity and convenience while gradually accustoming the public to interacting with intelligent machines as part of their daily routines. The success of virtual assistants also drove improvements in cloud computing infrastructure, as these systems required massive computational resources to process and respond to millions of user queries in real-time.

Beyond their immediate utility, virtual assistants served as trojan horses for AI adoption, introducing sophisticated machine learning capabilities to users who might never have considered themselves early adopters of artificial intelligence technology. They paved the way for broader acceptance of AI-powered tools and services across all sectors of society.

2.6: The Age of Generative AI (2020s – Present)

Transformers and Large Language Models

The introduction of the revolutionary "Transformer" architecture in 2017 through the groundbreaking paper "Attention Is All You Need" fundamentally transformed the entire field of artificial intelligence. This architectural innovation, developed by researchers at Google, solved long-standing problems in sequence modeling by introducing the concept of "attention mechanisms" that allowed models to focus on relevant parts of input data regardless of their position in the sequence.

The Transformer's elegance lay in its simplicity and parallelizability—unlike previous recurrent neural networks that processed sequences step by step, Transformers could process entire sequences simultaneously, dramatically reducing training time while improving performance. This breakthrough subsequently facilitated the development of massive-scale models including GPT-3, BERT, and numerous other sophisticated language processing systems that could understand and generate human-like text with unprecedented fluency and coherence.

As these models grew larger—from millions to billions to trillions of parameters—they began exhibiting emergent capabilities that their creators hadn't explicitly programmed. Large language models demonstrated the ability to engage in complex reasoning, write creative fiction, solve mathematical problems, and even generate computer code, all through the simple task of predicting the next word in a sequence.

The release of ChatGPT in late 2022 brought generative AI to mainstream attention, sparking global conversations about the implications of artificial intelligence for education, work, creativity, and society as a whole. Suddenly, millions of people had access to AI systems that could engage in sophisticated conversations, answer complex questions, and assist with a vast array of tasks—from writing emails to debugging code to creating poetry.

This generative AI revolution has ushered in a new era where artificial intelligence is no longer confined to specialized applications but has become a general-purpose tool that can augment human capabilities across virtually every domain of knowledge work. As we stand at this inflection point, we're witnessing the emergence of AI systems that may represent the first steps toward the artificial general intelligence that researchers have pursued since the field's inception at Dartmouth nearly seven decades ago.

3: THE CURRENT STATE OF AI

In this module, you will explore how AI is used in real-world sectors like healthcare, education, and marketing. You'll critically analyze these applications for ethical risks, examine real-life safety and privacy challenges, and propose your own ethically responsible AI solution. Through case analysis, research, and design, you'll deepen your understanding of how AI can both help and harm and how we can do better.

Learning Objectives

- Analyze AI use in sectors like healthcare, education, and marketing, with attention to ethical risks.
- Identify privacy and safety challenges in real-world AI deployments.
- Choose a real-world challenge where a safe and ethical AI solution could apply.

3.1: AI in Everyday Life

3.2: AI Transforming Healthcare

3.2.1: Benefits and Concerns

3.3: AI Transforming Education

3.3.1: Pros and Cons

3.3.1.1: Benefits to Students

3.3.1.2: Benefits for Teachers

3.4: AI in Marketing Strategies

3.4.1: Good and Bad

3.4.1.1: AI for Entrepreneurs

3.5: Privacy and Safety Challenges with AI

3.6: Summary

3.7: Assignments

3.7.1: Case Study 1 - AI for Personal Use

3.7.2: Case Study 2 - AI in Education

3.7.3: Case Study 3 - AI for Entrepreneurs

3.7.4: Knowledge Check

3.1: AI in Everyday Life

YouTube Video

This interactive element is not available in this format. Visit the link or scan the QR code to access it online.

[View this YouTube Video online](#)



Artificial intelligence has seamlessly integrated into various aspects of daily life, enhancing convenience and efficiency. Virtual assistants like Siri and Alexa use AI to understand voice commands, helping users manage tasks, set reminders, and access information hands-free. Similarly, AI-driven recommendation systems on platforms like Netflix and Spotify personalize entertainment choices based on user preferences.

In addition to entertainment and productivity, AI plays a critical role in improving safety and accessibility. Navigation apps such as Google Maps leverage AI to provide real-time traffic updates and optimal routes. Moreover, AI-powered tools like facial recognition and smart home devices contribute to security and automation, making everyday routines more streamlined and secure.

This page titled [3.1: AI in Everyday Life](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by Dr. LaTonya Motley-Northland Pioneer College.

3.2: AI TRANSFORMING HEALTHCARE

Artificial Intelligence (AI) is revolutionizing healthcare by enhancing diagnostic accuracy and efficiency. Machine learning algorithms analyze vast amounts of medical data, enabling early detection of diseases such as cancer through imaging technologies. These advancements help healthcare professionals make informed decisions and improve patient outcomes.

AI-powered tools are also streamlining administrative tasks, such as patient scheduling and medical record management. By automating these processes, healthcare providers can focus more on patient care, reducing burnout and increasing operational efficiency. This shift allows for a more patient-centered approach to healthcare delivery.

Furthermore, AI is driving innovation in personalized medicine. By analyzing genetic information and patient history, AI systems can recommend tailored treatment plans that cater to individual needs. This level of customization not only enhances treatment effectiveness but also fosters a deeper understanding of complex medical conditions.

3.2.1: Benefits and Concerns

3.2.1: Benefits and Concerns

Artificial Intelligence is transforming the education sector by offering personalized learning experiences and automating administrative tasks. AI-powered tools can adapt to individual learning styles, helping students grasp complex concepts more effectively. Additionally, educators benefit from reduced workloads, allowing them to focus on teaching and mentoring.

However, the integration of AI in education also raises concerns. Issues such as data privacy, the potential for bias in AI algorithms, and the risk of over-reliance on technology must be addressed. Striking a balance between leveraging AI's benefits and mitigating its challenges is essential for fostering an equitable and effective learning environment.

3.3: AI TRANSFORMING EDUCATION

Artificial Intelligence is revolutionizing the education sector by personalizing learning experiences for students. AI-powered tools can analyze individual learning patterns and adapt content to suit the unique needs of each learner, ensuring more effective and engaging education. This technology helps educators identify areas where students may struggle, enabling timely interventions and support.

Beyond personalization, AI is streamlining administrative tasks, allowing educators to focus more on teaching. Automated grading systems, for instance, save time by quickly assessing assignments and providing feedback. Additionally, AI-driven platforms facilitate collaborative learning by connecting students with peers and resources worldwide, fostering a more inclusive and interactive educational environment.

AI also plays a critical role in expanding access to education. Virtual tutors and online courses powered by AI make quality learning opportunities available to individuals in remote or underserved areas. By breaking down geographical and economic barriers, AI is helping to create a more equitable education system that benefits learners across the globe.

3.3.1: Pros and Cons

3.3.1.1: Benefits to Students

3.3.1.2: Benefits for Teachers

3.3.1: PROS AND CONS

3.3.1.1: Benefits to Students

3.3.1.2: Benefits for Teachers

This page titled [3.3.1: Pros and Cons](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by Dr. LaTonya Motley - Northland Pioneer College.

3.3.1.1: Benefits to Students

1. **Personalized Learning Experiences:** AI enables tailored educational content that adapts to individual learning styles and paces, leading to more effective knowledge acquisition.
2. **Enhanced Accessibility to Resources:** AI-powered tools provide students with access to diverse learning materials, breaking barriers of location and language.
3. **Improved Student Engagement:** Interactive AI applications make learning more engaging and dynamic, fostering curiosity and active participation in educational activities.
4. **Real-Time Feedback and Support:** AI systems offer instant feedback and assistance, helping students identify areas for improvement and achieve their academic goals efficiently.

3.3.1.2: Benefits for Teachers

1. **Streamlining Administrative Tasks:** AI tools can automate grading, attendance tracking, and scheduling, allowing teachers to focus more on teaching and student engagement.
2. **Enhancing Personalized Learning:** AI-driven platforms analyze student performance to tailor lessons, ensuring each learner receives support suited to their unique needs.
3. **Providing Real-Time Feedback:** AI systems offer instant feedback on assignments, helping students understand mistakes and improve their skills more effectively.
4. **Supporting Professional Development:** AI-powered resources recommend training materials and strategies, enabling teachers to stay updated with the latest educational practices.

3.4: AI IN MARKETING STRATEGIES

Artificial intelligence is revolutionizing the marketing landscape by enabling businesses to analyze vast amounts of data and uncover valuable insights. Through AI-powered tools, marketers can predict consumer behavior, personalize content, and optimize campaigns with unprecedented precision. This allows companies to deliver more targeted and effective messaging to their audiences.

One of the most impactful applications of AI in marketing is the use of chatbots and virtual assistants. These tools provide real-time customer support, answer queries, and even guide users through purchasing decisions. By automating these interactions, businesses can enhance customer satisfaction while reducing operational costs.

AI also plays a crucial role in content creation and distribution. Algorithms can generate tailored recommendations, ensuring that users receive content that aligns with their preferences and interests. Additionally, AI-driven analytics help marketers measure campaign performance and refine strategies, leading to better engagement and higher conversion rates.

3.4.1: Good and Bad

3.4.1.1: AI for Entrepreneurs

3.4.1: GOOD AND BAD

Artificial intelligence is transforming marketing by enabling businesses to deliver highly personalized customer experiences. Through advanced data analysis and predictive algorithms, AI helps marketers understand consumer behavior and tailor campaigns to meet specific needs. This not only improves customer satisfaction but also drives higher conversion rates.

However, the use of AI in marketing raises significant concerns, particularly around data privacy and ethical practices. Collecting and analyzing vast amounts of personal data can lead to potential misuse or breaches, undermining consumer trust. It is crucial for organizations to implement transparent policies and prioritize ethical AI practices to address these challenges effectively.

Artificial intelligence is revolutionizing healthcare by enhancing diagnostic accuracy, streamlining administrative tasks, and personalizing patient care. AI-powered tools can analyze vast amounts of medical data to identify patterns and provide insights that support early disease detection and treatment planning. These advancements not only improve patient outcomes but also reduce the burden on healthcare professionals.

However, the integration of AI in healthcare raises significant concerns, including data privacy, algorithmic bias, and the potential for over-reliance on technology. Ensuring that AI systems are transparent, ethically designed, and rigorously tested is crucial to addressing these challenges. By balancing innovation with responsibility, the healthcare industry can harness AI's potential while safeguarding patient trust and safety.

3.4.1.1: AI for Entrepreneurs

3.4.1.1: AI for Entrepreneurs

Artificial Intelligence (AI) offers entrepreneurs transformative tools to streamline operations and enhance decision-making. From automating routine tasks to providing data-driven insights, AI enables businesses to operate more efficiently and scale faster. Entrepreneurs can leverage AI-powered solutions such as chatbots, predictive analytics, and personalized marketing to stay competitive in dynamic markets.

However, adopting AI also comes with challenges that entrepreneurs must address. Concerns such as data privacy, ethical considerations, and the potential for job displacement require careful planning and implementation. By balancing innovation with responsibility, entrepreneurs can harness AI's potential while mitigating its risks, ensuring sustainable growth for their ventures.

3.5: Privacy and Safety Challenges with AI

Artificial Intelligence (AI) is becoming more integrated into our daily lives, from the phones in our pockets to the websites we visit. While these technologies offer many benefits, they also create new privacy and safety concerns that we need to understand. Think of AI as a powerful assistant that learns from information - including our personal data - to perform tasks and make predictions.

One major challenge is protecting our personal information. AI systems often collect and analyze large amounts of data about us to function effectively. This might include our location, shopping habits, health information, or conversations with digital assistants. Companies developing these systems need strong safeguards to prevent data breaches or misuse of this sensitive information. As users, we should also be aware of what information we're sharing and how it might be used.

Another concern is the creation of misleading content. AI can now generate realistic images, videos, and text that appear authentic but are completely fabricated. These "deepfakes" could be used to spread false information, impersonate people, or create convincing scams. This makes it harder for all of us to trust what we see and hear online. Developing reliable ways to identify AI-generated content and educating people about these possibilities has become increasingly important.

3.6: Summary

Artificial Intelligence (AI) is reshaping industries by enhancing efficiency, personalization, and decision-making. In education, AI supports adaptive learning, automates grading, and offers real-time feedback to students. In healthcare, it aids in diagnosing conditions, monitoring patients, and streamlining clinical workflows.

Meanwhile, marketing uses AI to analyze consumer behavior, deliver targeted content, and improve customer engagement through tools like chatbots and recommendation engines. Across these fields, AI helps process large amounts of data quickly, allowing professionals to focus on more complex, human-centered tasks.

3.7: Assignments

This page titled [3.7: Assignments](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by LibreTexts.

3.7.1: Case Study 1 - AI for Personal Use

Background:

Maria, a 38-year-old mother of two, has recently started using an AI-powered wellness app that tracks her mood, sleep, and physical activity. The app uses natural language processing to analyze her journal entries and offers personalized suggestions for stress management, sleep routines, and nutrition.

Scenario:

After two months of use, Maria notices that the app often suggests she is “at risk for burnout” and recommends lifestyle changes. She’s surprised when one day her insurance provider offers her a policy adjustment based on her “health tech profile.” She didn’t realize the app was sharing data with third parties.

Questions to Explore:

- How might Maria’s data be used in ways she didn’t expect?
- What ethical responsibilities do app developers have in protecting user privacy?
- Should users like Maria be more informed, or should companies do more to limit sharing?

3.7.2: Case Study 2 - AI in Education

Background:

Mr. Davis teaches 10th-grade English in a public school district that recently adopted an AI-powered essay-grading system to save time on grading and standardize feedback. The system scores essays and provides feedback on grammar, structure, and clarity.

Scenario:

After a few months, Mr. Davis notices that essays from multilingual students tend to receive lower scores, even when their content is strong. Students who rely on dialects or nonstandard English are penalized for “language issues.” Mr. Davis begins to worry that the system is reinforcing linguistic bias.

Questions to Explore:

- What risks arise when AI is used to evaluate student work?
- How should educators respond when automated tools show potential bias?
- What balance should be struck between efficiency and fairness?

This page titled [3.7.2: Case Study 2 - AI in Education](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by Dr. LaTonya Motley - Northland Pioneer College.

3.7.3: Case Study 3 - AI for Entrepreneurs

Background:

Raj runs a small e-commerce business selling eco-friendly kitchenware. To improve his marketing strategy, he signs up for an AI tool that predicts which customers are most likely to make a purchase based on their browsing behavior, location, and past purchases.

Scenario:

Raj notices the tool recommends excluding users from certain zip codes and age groups because they “aren’t profitable.” He’s concerned that this might reinforce economic or racial disparities, but the tool is increasing his sales.

Questions to Explore:

- Should Raj follow the tool’s recommendations if they improve business performance?
- What ethical considerations should small businesses weigh when using AI?
- How can entrepreneurs benefit from AI while avoiding discriminatory practices?

3.7.4: Knowledge Check

This page titled [3.7.4: Knowledge Check](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by Dr. LaTonya Motley - Northland Pioneer College.

4: AI and Ethics

This page titled [4: AI and Ethics](#) is shared under a [CC BY-NC-SA 4.0](#) license and was authored, remixed, and/or curated by LibreTexts.

5: Environmental Impacts of AI

This page titled [5: Environmental Impacts of AI](#) is shared under a [CC BY-NC-SA 4.0](#) license and was authored, remixed, and/or curated by LibreTexts.

6: AI in the Workspace

This page titled [6: AI in the Workspace](#) is shared under a [CC BY-NC-SA 4.0](#) license and was authored, remixed, and/or curated by LibreTexts.

7: LEGAL ASPECTS OF AI

7.1: Introduction

7.2: Copyright and Image Generation

7.3: What's the big deal?

7.4: Writing, Film, and Music

7.5: Case Study- Zarya of the Dawn

7.6: Teaching AI Ethics

7.7: Where does all that Data come from?

7.8: Protecting Personal Privacy

7.9: Case Study- AI Defamation

7.10: Teaching AI Ethics (1)

7.1: Introduction

Copyright and intellectual property

Closely related to truth and integrity are ideas of intellectual property and the legal concerns of copyright. Copyright issues have been particularly prevalent in AI image generation. AI image generators like Stable Diffusion, DALL-E 2, and Midjourney are trained on images “scraped” from the internet. [These images are then broken down and run through the AI algorithms, so that they can later be reconstructed.](#)

This has resulted in artists’ “styles” being used in AI image generation without their permission. Many artists believe that this infringes their [intellectual property rights](#), and is an ethical issue. An often-used counterargument is that all art is based on other artists’ work, and therefore the machine is simply replicating those processes. Class action lawsuits have already been filed against some AI image generators on behalf of artists.

Large language models like ChatGPT also incorporate huge amounts of other writers’ work. Where writing is publicly available – such as out of copyright books or CC journalism and articles – it can be incorporated into the dataset. Even when writing is protected by copyright it can become part of the datasets. Prompting a language model to write something *in the style of* another author could be viewed in the same way as an image generator adopting another artist’s style.

[There are also question marks over who owns the copyright to materials produced by AIs such as image generators and language models.](#)

Teaching points

- AI language models and image generators may breach copyright by appropriating others’ work.
- These models give users the ability to write and produce art in the style of other authors and artists without the original creator’s permission.
- Laws around these technologies are still murky, but there are developments happening all the time.

Subject examples

- Visual Arts: Is it art? If a user can generate an image in the style of another artist with just a few prompt words, does the digital output count as “real art”?
- Performing Arts: Some AI models can produce music and lyrics as well as visual art. Is it possible to create a complete AI performing artist? If we can, does it mean we should?
- English: Does producing a piece of writing in the style of another author infringe their intellectual property rights?



I call this one “Melbourne skyline in the style of picasso and van gogh -ar 3:2 -q 2 -v 4” but is it art? And is it legal? Via Midjourney.

Privacy and security

Privacy is a major concern in the development and use of AI systems. As these technologies become more sophisticated and integrated into our lives, there are increasing concerns about the [collection and use of personal data](#), [data breaches](#), and the [lack of transparency](#) in AI decision-making.

One of the most prominent examples of these issues can be found in the use of facial recognition technology. This technology, which is used in a variety of applications such as security, surveillance, and marketing, has been criticised for its potential to violate individuals' privacy and civil rights. For example, [facial recognition systems have been known to have higher error rates for people with darker skin tones](#), and have been used to target and monitor marginalised communities as discussed earlier in “bias”.


Another example of privacy concerns with AI systems is targeted advertising. AI-powered algorithms are used to analyse data on individuals' online activities in order to deliver targeted ads. Whilst this may seem harmless, it raises [concerns about data privacy](#), data breaches, and the use of personal data for commercial gain.

Teaching points:

- AI systems raise concerns about the collection and use of personal data, data breaches, and the lack of transparency in decision-making.
- Facial recognition technology has been criticized for its potential to violate individuals' privacy and civil rights, and for its higher error rates for people with darker skin tones.
- Targeted advertising raises concerns about data privacy and the use of personal data for commercial gain.

Subject examples:

- Humanities (Legal Studies): What is the impact of AI on data protection laws such as GDPR and the protection of personal data?
- Mathematics: How can we analyse the data sets used by AI systems for potential biases and privacy issues?
- Health and Physical Education: What are the privacy concerns surrounding personal health data and its use in AI-powered healthcare technologies?



Every click and like goes towards powering AI surveillance. Image via Midjourney. Prompt in alt text.

Data collection and “datafication”

The phrase “[data is the new oil](#)” crops up everywhere when you start researching AI. As I wrote about earlier in “bias”, Artificial Intelligence is powered by huge amounts of data. The oil analogy suggests both data as fuel, but also the costly, dangerous, and extractive process of data-mining. In the constant quest for more and more data, the companies that develop AI systems sometimes revert to unethical practices.

“[Datafication](#)” is a term used for turning all parts of our life into a data point to be fed into an AI algorithm. As per the privacy discussion above, this should raise some serious concerns. From location data to health data, shopping habits, likes, clicks, and views, almost every interaction we have with technology is fed into an algorithm somewhere.

As we become commodities, we open ourselves up to exploitation. One major ethical concern with “datafication” is that fact that the users become the products, and that the free-labour of the users is used to generate capital for the platform owners.

“Big Data” also contributes to many of the issues we have described so far, including bias and [discrimination](#). Any data collected by the devices we wear and use or the platforms we subscribe to ultimately becomes part of the algorithm’s “world view”. Unfortunately, because not everyone in the world has access to these technologies, that worldview is by definition missing some very important data.

Teaching points

- Data is the new oil, both in the sense of being fuel, and in that it is costly and damaging to extract.
- “Datafication” is the process of turning every aspect of our lives into data.
- Users become products, and user data becomes capital.
- Big Data doesn’t include marginalised groups, and therefore doesn’t truly represent society.

Subject examples

- Humanities (History): What is the historical and societal context of data collection and the impact of datafication on different communities?
- Design and Technology (Digital Technology): How can we design and develop ethical data collection practices and data privacy measures for AI systems?
- English: Based on recent media, how might we critically analyse datafication and its implications on privacy and data protection?

7.2: Copyright and Image Generation

I'm starting with image generation because it has been the most public and most contentious aspect of AI and copyright. For a quick scan of how problematic the area is, you only have to do a quick search of whether AI artists can copyright their work. Here are some of the results:

- [No, you can't copyright images made by A.I., says the U.S. Copyright Office](#)
- [U.S. Copyright Office Says AI-Generated Images Do Not Qualify For Copyright Protection](#)
- [AI-generated art can be copyrighted, say US officials – with a catch](#)
- [AI-created images lose U.S. copyrights in test for new technology](#)
- [AI art tools Stable Diffusion and Midjourney targeted with copyright lawsuit](#)

In some cases, it's an outright "no". In others, it seems but be a "no, but...". And in the final example, it's a hard no with a side of litigation, as companies like Stability AI (behind Stable Diffusion) and Midjourney find themselves on the receiving end of lawsuits.

7.3: What's the big deal?

So why is AI image generation so contentious?

The primary reason is based on how these models are trained. In order to build an AI image generator, the developer must use millions or even billions of images. Stable Diffusion, for example, was [trained on around 2.3 billion images](#). Many of these images have been “scraped” from the internet without the consent of the original creators.

This leads to problems with attribution, and with the potential for these AI image generators to reproduce art in another artist's style. I'm sure you've seen posts already with AI generated art in the style of Van Gogh or Rembrandt. However, it also applies to living artists and photographers whose work has been scraped from sites like artstation and Flickr.

There's the additional problem of whether the current copyright laws extend to work created “by a machine”. Although AI images are generated by a human controlling the input via the prompt, it has been argued that the actual image is created by the AI, and not the human. This throws a legal roadblock in the way of copyrighting AI art.

US law states that intellectual property can be copyrighted only if it was the product of human creativity, and the USCO only acknowledges work authored by humans at present. Machines and generative AI algorithms, therefore, cannot be authors, and their outputs are not copyrightable.

https://www.theregister.com/2023/03/16/ai_art_copyright_usco/



Jason Allen's Théâtre D'opéra Spatial won first place in the Colorado State Fair, until the artist revealed it was generated in Midjourney and declined the prize. Since then, the artist has not been able to get copyright status for the work.

This page titled [7.3: What's the big deal?](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by LibreTexts.

7.4: Writing, Film, and Music

Although the bulk of the conversation around AI and copyright has centred on the visual arts, the conversation has also arisen in text generation, film, and music.

The same human authorship requirement for copyright protection under the current US law means that an AI-generated writing is likely either a public domain work immediately upon creation or a derivative work of the materials the AI tool was exposed to during training – the text found in the dataset.

The ownership of the rights in such a derivative would depend on various issues, including the dataset for training the AI tool (of which there are many variations, depending on the model used), its components, and the similarity between any particular work in the training set and the AI work.

We're yet to see many real discussions around copyright in film and music, but with AI generated film technologies just around the corner, we surely will. AI can already synthesise realistic versions of existing artist's voices, like the [example](#) below which went viral last week:

7.5: Case Study- Zarya of the Dawn

In February 2023, the US Copyright Office revoked copyright protection for images created using the AI-powered Midjourney image generator for the comic book “Zarya of the Dawn” and issued a new copyright registration covering only the text of the work and the *arrangement* of images and text. The justification for this decision was incomplete information in the original copyright registration, which failed to disclose that the images were created by an AI model. The Copyright Office argued that the images in the work generated by Midjourney were not the product of human authorship and thus not copyrightable. Despite Kashtanova’s claim that she had “guided” the structure and content of each image, the Copyright Office argued that Midjourney, not Kashtanova, originated the “traditional elements of authorship” in the images.

Kashtanova’s attorney argued that every aspect of the work was authored by Kashtanova, with Midjourney serving merely as an assistive tool. However, this argument was rejected by the Copyright Office, which provided additional analogies to explain why Kashtanova was not the creator of the images, including the idea of hiring a human to create images using descriptions and performing a text-based image search on the internet.

Kashtanova reacted to the letter by framing it as an overall win for AI-augmented artists. The ruling protects the comic book’s story and the image arrangement, which “covers a lot of uses for the people in the AI art community.” However, she expressed disappointment that the Copyright Office did not agree to recognise her copyright of the individual images.

Kashtanova believes the output of a Generative AI model depends directly on the creative input of the artist and is not random: individual images produced by Midjourney are a direct expression of her creativity and thus copyrightable.

The decision means that AI-generated imagery, without human-authored elements, cannot currently be copyrighted in the United States. The Copyright Office’s ruling on the matter will likely hold unless it’s challenged in court, revised by law, or re-examined in the future. The decision may eventually be reconsidered as the result of a cultural shift in how society perceives AI-generated art, allowing for a new interpretation by different members of the US Copyright Office in the future. AI-powered artwork is still a novel and poorly understood technology, but it is rapidly becoming a feature of digital art.

[Cover of Zarya of the Dawn. Image via https://en.Wikipedia.org/wiki/Zarya_of_the_Dawn](https://en.Wikipedia.org/wiki/Zarya_of_the_Dawn)

7.6: Teaching AI Ethics

1. Legal Studies: What is the current legal stance on AI copyright [in your country](#)?
2. English and Literature: How does the use of AI to generate content challenge traditional [notions of authorship and creativity](#), and what are the ethical implications of this challenge?
3. Computer Science: What [ethical considerations](#) should be taken into account when designing AI systems for copyright enforcement?
4. Philosophy: What [ethical theories](#) can be applied to the use of AI in copyright enforcement, and how do these theories inform debates over the balance between copyright protection and free expression?
5. Business and Economics: How does the use of AI [affect the market for creative works](#), and what ethical issues arise from these effects?
6. Media Studies: What implications might AI have on the [creation](#) and copyrighting of film?
7. Visual Arts: How does the use of AI in creating art raise questions about [copyright and intellectual property rights](#)? What ethical considerations arise when using AI to produce artwork?
8. Music: Is it possible to [copyright an artist's voice](#)?

7.7: Where does all that Data come from?

Developers of large language models, such as ChatGPT, often scrape their training data indiscriminately from the web without paying any attention to individual rights. These models are trained on vast swathes of internet data, and often include personal information that has been collected without consent or used in violation of privacy laws. This has raised concerns about the ethical implications of developing AI models that rely on data collected without regard for individual privacy rights.

The lack of transparency and accountability around the collection and use of personal data in AI development has been a longstanding issue. The vast amount of data required to train these models means that personal information is often collected without explicit consent or knowledge of the individuals affected. Critics argue that developers of large language models prioritise the creation of powerful algorithms over individual privacy rights, and that [the industry is not sufficiently regulated](#).

These concerns have landed [OpenAI in trouble with European regulators](#), particularly under the General Data Protection Regulation (GDPR) laws. The Italian regulator recently issued a temporary emergency decision demanding that OpenAI stop using the personal information of millions of Italians included in its training data, citing a lack of legal justification for using people's personal information in ChatGPT. The GDPR rules protect the data of over 400 million people across Europe and apply to personal data that is freely available online. The decision by the Italian regulator highlights the growing concerns around the development of large AI models and the use of personal information in training data.

In the US, the [federal privacy commission is also investigating OpenAI](#) following a claim made against the company that it has been unlawfully using personal and private data.



This page titled [7.7: Where does all that Data come from?](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by LibreTexts.

7.8: Protecting Personal Privacy

As I covered in the [first post in this series](#), AI systems have the potential to perpetuate and amplify biases in data, leading to discrimination against certain groups or individuals. This is a serious concern when it comes to privacy, as these biases can lead to the exclusion or mistreatment of individuals based on their personal characteristics. It can lead to members of the public being surveilled based on skin colour, place of residence, or other factors which are part of the data used when training the models. These concerns extend into many areas of the AI industry including facial and affect recognition, which I'll talk about in a later post.

The storage of personal data in AI training data is also a significant privacy concern. In the creation of these models, personal data has been collected without explicit consent or knowledge of the individuals affected, and there may be inadequate protections in place to ensure that this data is used ethically and responsibly. Data breaches and cyber attacks also a huge risk for AI systems. [Several weeks ago, OpenAI experienced a breach](#) due to a bug in one of their code libraries which revealed the first and last names and email addresses of ChatGPT Plus subscribers, along with financial details.

This page titled [7.8: Protecting Personal Privacy](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by LibreTexts.

7.9: Case Study- AI Defamation

As covered in [my article on truth and academic integrity](#), Artificial intelligence has the potential to generate false information, leading to serious privacy concerns. In a recent case in [Hepburn Shire, Australia](#), OpenAI once again faces the possibility of legal action for defamation. ChatGPT incorrectly described regional mayor Brian Hood as a guilty party in a foreign bribery scandal. The mayor was actually a whistle-blower who had reported the bribe payments.

ChatGPT's errors arose from its indiscriminate data-scraping, as well as the inability of these models to distinguish between true and false claims. As a result, it generated convincing but incorrect information. Although OpenAI, the company that created ChatGPT, has taken some steps to protect people's privacy, such as removing personal information from training data, such actions may not be sufficient to prevent the spread of false information.

This case highlights the legal challenges associated with suing AI companies for defamation, particularly given the issue of jurisdiction. Although the legal implications of AI technologies like ChatGPT are still uncharted, the case demonstrates the need for more cooperative efforts between AI developers, social media companies, and government agencies to mitigate the risk of generating misleading information.

When personal user data – even publicly available data, like the original news story about Brian Hood's involvement in the bribery case – is combined with a language model's capacity for generating falsehoods, we have a recipe for damaging output.

7.10: Teaching AI Ethics (1)

- Legal Studies: How do [GDPR laws](#) protect user privacy when it comes to AI? What legal consequences can organisations face for non-compliance or data breaches?
- English and Literature: [How has literature portrayed Artificial Intelligence](#) and our concerns about these technologies? What can we learn from these portrayals about public sentiment towards privacy concerns?
- Mathematics: [How do encryption algorithms protect personal data?](#) What is the role of probability in securing data privacy and preventing data breaches?
- Computer Science: What [methods can be implemented](#) to ensure data privacy and security in AI systems? How can anonymisation techniques help in protecting user privacy?
- Philosophy: How can we [balance the benefits of AI and the need for privacy](#)? What ethical theories can be applied to address privacy concerns in AI development and usage?
- Science: How can [scientific research contribute](#) to the development of more secure AI systems? What role do scientists and researchers play in addressing privacy concerns in AI?
- Business and Economics: How do privacy concerns and regulatory compliance affect the market and business outcomes? [What economic incentives exist for companies to invest in privacy](#) and data protection technologies for AI?
- Health and PE: [How can AI be used in healthcare while maintaining patient privacy?](#) What are the potential ethical implications of privacy breaches in the healthcare system?

8: AI and Society

This page titled [8: AI and Society](#) is shared under a [CC BY-NC-SA 4.0](#) license and was authored, remixed, and/or curated by LibreTexts.

9: The Convergence of Thinking

This page titled [9: The Convergence of Thinking](#) is shared under a [CC BY-NC-SA 4.0](#) license and was authored, remixed, and/or curated by LibreTexts.

10: AI and Politics

This page titled [10: AI and Politics](#) is shared under a [CC BY-NC-SA 4.0](#) license and was authored, remixed, and/or curated by LibreTexts.

11: AI and Humanity

This page titled [11: AI and Humanity](#) is shared under a [CC BY-NC-SA 4.0](#) license and was authored, remixed, and/or curated by LibreTexts.

12: Summary

This page titled [12: Summary](#) is shared under a [CC BY-NC-SA 4.0](#) license and was authored, remixed, and/or curated by LibreTexts.

DETAILED LICENSING

Overview

Title: Ethics and Current Issues in Technology

Webpages: 61

Applicable Restrictions: Noncommercial

All licenses found:

- CC BY-NC-SA 4.0: 37.7% (23 pages)
- CC BY 4.0: 29.5% (18 pages)
- CC BY-NC 4.0: 16.4% (10 pages)
- CC BY-SA 4.0: 8.2% (5 pages)
- Public Domain: 3.3% (2 pages)
- Undeclared: 3.3% (2 pages)
- Unknown License: 1.6% (1 page)

By Page

- Ethics and Current Issues in Technology - *Unknown License*
 - Front Matter - CC BY-NC-SA 4.0
 - TitlePage - CC BY-NC-SA 4.0
 - InfoPage - Public Domain
 - Table of Contents - Undeclared
 - Licensing - CC BY-NC-SA 4.0
 - 1: Introduction - CC BY-NC-SA 4.0
 - 1.1: Chapter Introduction - CC BY-SA 4.0
 - 1.2: Machines, AI, and Humans - CC BY-SA 4.0
 - 1.3: Types of AI - CC BY-SA 4.0
 - 1.4: Applications of AI - CC BY-SA 4.0
 - 1.5: Emerging Technology Trends - CC BY-SA 4.0
 - 2: The History of AI - CC BY-NC-SA 4.0
 - 2.1: The Early Foundations (Pre-1950s) - CC BY-NC-SA 4.0
 - 2.2: The Birth of AI (1950s – 1960s) - CC BY-NC-SA 4.0
 - 2.3: The AI Winters and Specialization (1970s – 1980s) - CC BY-NC-SA 4.0
 - 2.4: Machine Learning and the Data Revolution (1990s – 2000s) - CC BY-NC-SA 4.0
 - 2.5: The Deep Learning Era (2010s) - CC BY-NC-SA 4.0
 - 2.6: The Age of Generative AI (2020s – Present) - CC BY-NC-SA 4.0
 - 3: The Current State of AI - CC BY 4.0
 - 3.1: AI in Everyday Life - CC BY 4.0
 - 3.2: AI Transforming Healthcare - CC BY 4.0
 - 3.2.1: Benefits and Concerns - CC BY 4.0
 - 3.3: AI Transforming Education - CC BY 4.0
 - 3.3.1: Pros and Cons - CC BY 4.0
 - 3.3.1.1: Benefits to Students - CC BY 4.0
 - 3.3.1.2: Benefits for Teachers - CC BY 4.0
 - 3.4: AI in Marketing Strategies - CC BY 4.0
 - 3.4.1: Good and Bad - CC BY 4.0
 - 3.4.1.1: AI for Entrepreneurs - CC BY 4.0
 - 3.5: Privacy and Safety Challenges with AI - CC BY 4.0
 - 3.6: Summary - CC BY 4.0
 - 3.7: Assignments - CC BY 4.0
 - 3.7.1: Case Study 1 - AI for Personal Use - CC BY 4.0
 - 3.7.2: Case Study 2 - AI in Education - CC BY 4.0
 - 3.7.3: Case Study 3 - AI for Entrepreneurs - CC BY 4.0
 - 3.7.4: Knowledge Check - CC BY 4.0
 - 4: AI and Ethics - CC BY-NC-SA 4.0
 - 5: Environmental Impacts of AI - CC BY-NC-SA 4.0
 - 6: AI in the Workspace - CC BY-NC-SA 4.0
 - 7: Legal Aspects of AI - CC BY-NC-SA 4.0
 - 7.1: Introduction - CC BY-NC 4.0
 - 7.2: Copyright and Image Generation - CC BY-NC 4.0
 - 7.3: What's the big deal? - CC BY-NC 4.0
 - 7.4: Writing, Film, and Music - CC BY-NC 4.0
 - 7.5: Case Study- Zarya of the Dawn - CC BY-NC 4.0
 - 7.6: Teaching AI Ethics - CC BY-NC 4.0
 - 7.7: Where does all that Data come from? - CC BY-NC 4.0

- 7.8: Protecting Personal Privacy - *CC BY-NC 4.0*
- 7.9: Case Study- AI Defamation - *CC BY-NC 4.0*
- 7.10: Teaching AI Ethics (1) - *CC BY-NC 4.0*
- 8: AI and Society - *CC BY-NC-SA 4.0*
- 9: The Convergence of Thinking - *CC BY-NC-SA 4.0*
- 10: AI and Politics - *CC BY-NC-SA 4.0*
- 11: AI and Humanity - *CC BY-NC-SA 4.0*
- 12: Summary - *CC BY-NC-SA 4.0*
- Back Matter - *CC BY-NC-SA 4.0*
 - Index - *Public Domain*
 - Glossary - *CC BY-NC-SA 4.0*
 - Detailed Licensing - *CC BY-NC-SA 4.0*
 - Detailed Licensing (1) (1) - *Undeclared*

GLOSSARY

S

Sample Word 1 | Sample Definition 1

This page titled [Glossary](#) is shared under a [CC BY-NC-SA 4.0](#) license and was authored, remixed, and/or curated by LibreTexts.

INDEX

A

academic integrity
7.9: Case Study- AI Defamation

academic performance
3.7.2: Case Study 2 - AI in Education

academic support
3.3.1.1: Benefits to Students

adaptive learning
3.2.1: Benefits and Concerns
3.6: Summary

administrative tasks
3.3: AI Transforming Education
3.3.1.2: Benefits for Teachers

AI
1: Introduction
1.2: Machines, AI, and Humans
1.3: Types of AI
1.4: Applications of AI
3: The Current State of AI
3.1: AI in Everyday Life
3.2: AI Transforming Healthcare
3.5: Privacy and Safety Challenges with AI
3.6: Summary
3.7.2: Case Study 2 - AI in Education
7: Legal Aspects of AI
7.1: Introduction
7.2: Copyright and Image Generation
7.3: What's the big deal?
7.4: Writing, Film, and Music
7.5: Case Study- Zarya of the Dawn
7.8: Protecting Personal Privacy

AI algorithms
3.2.1: Benefits and Concerns

AI art
7.2: Copyright and Image Generation
7.3: What's the big deal?

AI development
2.2: The Birth of AI (1950s - 1960s)

AI ethics
7.7: Where does all that Data come from?

AI in education
3.3.1.1: Benefits to Students

AI technologies
7.9: Case Study- AI Defamation

AI tool
3.7.3: Case Study 3 - AI for Entrepreneurs

AI tools
3.3.1.2: Benefits for Teachers

AI Winter
2.3: The AI Winters and Specialization (1970s - 1980s)

AI_Winter
2: The History of AI

Alexa
3.1: AI in Everyday Life

AlexNet
2.5: The Deep Learning Era (2010s)

algorithm
2.4: Machine Learning and the Data Revolution (1990s - 2000s)

algorithm bias
7.1: Introduction

algorithms
1.2: Machines, AI, and Humans

analysis
3: The Current State of AI

analytics
3.4: AI in Marketing Strategies

applications_of_AI
1: Introduction

art
7.3: What's the big deal?

Artificial Intelligence
1.1: Chapter Introduction
1.3: Types of AI
2.1: The Early Foundations (Pre-1950s)
2.2: The Birth of AI (1950s - 1960s)
2.3: The AI Winters and Specialization (1970s - 1980s)
2.4: Machine Learning and the Data Revolution (1990s - 2000s)
2.5: The Deep Learning Era (2010s)
2.6: The Age of Generative AI (2020s - Present)
3.1: AI in Everyday Life
3.2.1: Benefits and Concerns
3.3: AI Transforming Education
3.4: AI in Marketing Strategies
3.4.1: Good and Bad
3.4.1.1: AI for Entrepreneurs
3.5: Privacy and Safety Challenges with AI
3.6: Summary
7.6: Teaching AI Ethics
7.9: Case Study- AI Defamation

artificial_intelligence
2: The History of AI

ArtificialIntelligence
3.2: AI Transforming Healthcare

artists rights
7.1: Introduction

assessment
3.7.2: Case Study 2 - AI in Education

assignments
3.3.1.2: Benefits for Teachers

attendance tracking
3.3.1.2: Benefits for Teachers

attention mechanisms
2.6: The Age of Generative AI (2020s - Present)

attribution
7.3: What's the big deal?

authorship
7.4: Writing, Film, and Music
7.5: Case Study- Zarya of the Dawn
7.6: Teaching AI Ethics

automate grading
3.3.1.2: Benefits for Teachers

automated grading systems
3.3: AI Transforming Education

automating tasks
3.4.1.1: AI for Entrepreneurs

automation
1.4: Applications of AI

autonomous technology
1.4: Applications of AI

autonumheader:yes1
3: The Current State of AI
3.3.1: Pros and Cons

autonumheader:yes3
1: Introduction
1.1: Chapter Introduction
1.2: Machines, AI, and Humans
1.3: Types of AI
1.4: Applications of AI
1.5: Emerging Technology Trends
10: AI and Politics
11: AI and Humanity
12: Summary
4: AI and Ethics
5: Environmental Impacts of AI
6: AI in the Workspace
8: AI and Society
9: The Convergence of Thinking
Ethics and Current Issues in Technology

B

BERT
2.6: The Age of Generative AI (2020s - Present)

bias
3.2.1: Benefits and Concerns

biases
7.8: Protecting Personal Privacy

browsing behavior
3.7.3: Case Study 3 - AI for Entrepreneurs

burnout risk
3.7.1: Case Study 1 - AI for Personal Use

business
7.6: Teaching AI Ethics

C

cancer
3.2: AI Transforming Healthcare

case study
7: Legal Aspects of AI

cautionary_approach
1: Introduction

chatbots
1.3: Types of AI
3.4: AI in Marketing Strategies

ChatGPT
7.7: Where does all that Data come from?
7.9: Case Study- AI Defamation

chess
2.4: Machine Learning and the Data Revolution (1990s - 2000s)

class action lawsuits
7.1: Introduction

cognitive capabilities
2.1: The Early Foundations (Pre-1950s)

collaborative learning
3.3: AI Transforming Education

comic book
7.5: Case Study- Zarya of the Dawn

complex tasks
1.2: Machines, AI, and Humans

computational power

2.4: Machine Learning and the Data Revolution (1990s – 2000s)

computer science

7.6: Teaching AI Ethics

computer vision

2.5: The Deep Learning Era (2010s)

ComputerScience

7.10: Teaching AI Ethics (1)

connected devices

1.5: Emerging Technology Trends

consciousness

2.1: The Early Foundations (Pre-1950s)

consent

7.8: Protecting Personal Privacy

consumer behavior

3.4: AI in Marketing Strategies
3.4.1: Good and Bad

content creation

3.4: AI in Marketing Strategies

copyright

7: Legal Aspects of AI
7.1: Introduction
7.2: Copyright and Image Generation
7.3: What's the big deal?
7.4: Writing, Film, and Music
7.5: Case Study- Zarya of the Dawn
7.6: Teaching AI Ethics

copyright infringement

7.1: Introduction

copyright protection

7.2: Copyright and Image Generation

creative input

7.5: Case Study- Zarya of the Dawn

creative works

7.6: Teaching AI Ethics

creativity

7.6: Teaching AI Ethics

creators

7.3: What's the big deal?

customer prediction

3.7.3: Case Study 3 - AI for Entrepreneurs

customer satisfaction

3.4.1: Good and Bad

customer support

3.4: AI in Marketing Strategies

D

Dartmouth Conference

2.2: The Birth of AI (1950s – 1960s)

Dartmouth_Conference

2: The History of AI

data analysis

1.2: Machines, AI, and Humans
3.4: AI in Marketing Strategies
3.4.1: Good and Bad

data breaches

3.5: Privacy and Safety Challenges with AI
7.8: Protecting Personal Privacy

data monitoring

1.5: Emerging Technology Trends

data privacy

3.2.1: Benefits and Concerns
3.4.1: Good and Bad
3.7.1: Case Study 1 - AI for Personal Use
7.1: Introduction

data protection

3.5: Privacy and Safety Challenges with AI
7.7: Where does all that Data come from?
7.8: Protecting Personal Privacy

data revolution

2.4: Machine Learning and the Data Revolution (1990s – 2000s)

data scraping

7.9: Case Study- AI Defamation

dataBreaches

7.10: Teaching AI Ethics (1)

dataPrivacy

7.10: Teaching AI Ethics (1)

dataSecurity

7.10: Teaching AI Ethics (1)

decision making

1.2: Machines, AI, and Humans

deep blue

2.4: Machine Learning and the Data Revolution (1990s – 2000s)

deep learning

1.3: Types of AI
2.5: The Deep Learning Era (2010s)

DENDRAL

2.3: The AI Winters and Specialization (1970s – 1980s)

derivative works

7.4: Writing, Film, and Music

diagnosticaccuracy

3.2: AI Transforming Healthcare

digital art

7.5: Case Study- Zarya of the Dawn

digital assistants

3.5: Privacy and Safety Challenges with AI

discrimination

7.8: Protecting Personal Privacy

diverse learning materials

3.3.1.1: Benefits to Students

drone

1.4: Applications of AI

E

early_foundations

2: The History of AI

earlydetection

3.2: AI Transforming Healthcare

ecommerce

3.7.3: Case Study 3 - AI for Entrepreneurs

economics

7.6: Teaching AI Ethics

education

3: The Current State of AI
3.2.1: Benefits and Concerns
3.3: AI Transforming Education
3.6: Summary
3.7.2: Case Study 2 - AI in Education

educators

3.2.1: Benefits and Concerns

effective education

3.3: AI Transforming Education

efficiency

3.4.1.1: AI for Entrepreneurs
3.6: Summary
3.7.2: Case Study 2 - AI in Education

emerging technology trends

1.1: Chapter Introduction

emerging_technology_trends

1: Introduction

emotional intelligence

1.2: Machines, AI, and Humans

encryptionAlgorithms

7.10: Teaching AI Ethics (1)

energy harvesting

1.5: Emerging Technology Trends

enhanced accessibility

3.3.1.1: Benefits to Students

entertainment

3.1: AI in Everyday Life

entrepreneurs

3.4.1.1: AI for Entrepreneurs
3.7.3: Case Study 3 - AI for Entrepreneurs

equitable learning

3.2.1: Benefits and Concerns

ethical concerns

7.1: Introduction

ethical considerations

3.4.1.1: AI for Entrepreneurs

ethical implications

7.6: Teaching AI Ethics

ethical responsibilities

3.7.1: Case Study 1 - AI for Personal Use

ethical use

7.8: Protecting Personal Privacy

ethics

3: The Current State of AI
7: Legal Aspects of AI

everyday technology

1.2: Machines, AI, and Humans

expert systems

2.3: The AI Winters and Specialization (1970s – 1980s)

expert_systems

2: The History of AI

F

false information

3.5: Privacy and Safety Challenges with AI
7.9: Case Study- AI Defamation

feedback

3.7.2: Case Study 2 - AI in Education

film

7: Legal Aspects of AI
7.4: Writing, Film, and Music

founding principles

2.2: The Birth of AI (1950s – 1960s)

Frankenstein

2.1: The Early Foundations (Pre-1950s)

G

garry kasparov

2.4: Machine Learning and the Data Revolution (1990s – 2000s)

GDPR

7.10: Teaching AI Ethics (1)
7.7: Where does all that Data come from?

general AI

1.3: Types of AI

generative AI

2.6: The Age of Generative AI (2020s – Present)
7.3: What's the big deal?
7.5: Case Study- Zarya of the Dawn

golem

2.1: The Early Foundations (Pre-1950s)

GPU computing

2.5: The Deep Learning Era (2010s)

grading

3.7.2: Case Study 2 - AI in Education

H

haptic devices

1.5: Emerging Technology Trends

health tech profile

3.7.1: Case Study 1 - AI for Personal Use

healthcare

3: The Current State of AI
3.2: AI Transforming Healthcare
3.4.1: Good and Bad
3.6: Summary

history of artificial intelligence

1.1: Chapter Introduction

human behavior

1.2: Machines, AI, and Humans

human reasoning

2.2: The Birth of AI (1950s - 1960s)

humans

1: Introduction

I

identity

2.1: The Early Foundations (Pre-1950s)

image classification

2.5: The Deep Learning Era (2010s)

image generation

7: Legal Aspects of AI
7.1: Introduction
7.2: Copyright and Image Generation
7.3: What's the big deal?

imagery

7.5: Case Study- Zarya of the Dawn

imagingtechnologies

3.2: AI Transforming Healthcare

impact on business

1.1: Chapter Introduction

individual consent

7.7: Where does all that Data come from?

individual learning styles

3.3.1.1: Benefits to Students

industries

3.6: Summary

innovation

2.2: The Birth of AI (1950s - 1960s)

intellectual property

7.1: Introduction
7.3: What's the big deal?
7.5: Case Study- Zarya of the Dawn

intelligent agents

1.4: Applications of AI

intelligent devices

1.1: Chapter Introduction

intelligent machines

1.2: Machines, AI, and Humans

interactive learning

3.3.1.1: Benefits to Students

Internet of Things

1.5: Emerging Technology Trends

IoT

1.5: Emerging Technology Trends

J

John McCarthy

2.2: The Birth of AI (1950s - 1960s)

John_McCarthy

2: The History of AI

K

knowledge representation

2.3: The AI Winters and Specialization (1970s - 1980s)

L

language models

7.7: Where does all that Data come from?

large language models

2.6: The Age of Generative AI (2020s - Present)

learning methods

1.1: Chapter Introduction

learning patterns

3.3: AI Transforming Education

learning styles

3.2.1: Benefits and Concerns

legal challenges

7.9: Case Study- AI Defamation

legal issues

7.2: Copyright and Image Generation
7.3: What's the big deal?

legal studies

7.6: Teaching AI Ethics

LegalStudies

7.10: Teaching AI Ethics (1)

Lighthill Report

2.3: The AI Winters and Specialization (1970s - 1980s)

Lighthill_Report

2: The History of AI

litigation

7.2: Copyright and Image Generation

location data

3.7.3: Case Study 3 - AI for Entrepreneurs

Logic Theorist

2.2: The Birth of AI (1950s - 1960s)

M

machine learning

1.2: Machines, AI, and Humans
1.3: Types of AI
2.4: Machine Learning and the Data Revolution (1990s - 2000s)
2.5: The Deep Learning Era (2010s)
2.6: The Age of Generative AI (2020s - Present)

machine_learning

2: The History of AI

machinelearning

3.2: AI Transforming Healthcare

machines

1: Introduction

manufacturing

1.5: Emerging Technology Trends

marketing

3: The Current State of AI
3.4: AI in Marketing Strategies
3.4.1: Good and Bad
3.6: Summary

marketing strategy

3.7.3: Case Study 3 - AI for Entrepreneurs

Marvin_Minsky

2: The History of AI

Mary Shelley

2.1: The Early Foundations (Pre-1950s)

medical technology

1.4: Applications of AI

medicaldata

3.2: AI Transforming Healthcare

Midjourney

7.2: Copyright and Image Generation
7.5: Case Study- Zarya of the Dawn

misleading content

3.5: Privacy and Safety Challenges with AI

misleading information

7.9: Case Study- AI Defamation

mood tracking

3.7.1: Case Study 1 - AI for Personal Use

music

7: Legal Aspects of AI
7.4: Writing, Film, and Music

MYCIN

2.3: The AI Winters and Specialization (1970s - 1980s)

mythology

2.1: The Early Foundations (Pre-1950s)

N

nanobot

1.4: Applications of AI

narrow AI

1.3: Types of AI

narrow domains

2.3: The AI Winters and Specialization (1970s - 1980s)

natural language processing

1.1: Chapter Introduction
1.3: Types of AI
1.4: Applications of AI
2.5: The Deep Learning Era (2010s)
3.7.1: Case Study 1 - AI for Personal Use

neural networks

2.5: The Deep Learning Era (2010s)
2.6: The Age of Generative AI (2020s - Present)

NLP

1.3: Types of AI

O

online courses

3.3: AI Transforming Education

OpenAI

7.7: Where does all that Data come from?
7.9: Case Study- AI Defamation

ownership rights

7.4: Writing, Film, and Music

P

past purchases

3.7.3: Case Study 3 - AI for Entrepreneurs

patientoutcomes

3.2: *AI Transforming Healthcare*

patientPrivacy

7.10: *Teaching AI Ethics (1)*

personal data

7.8: *Protecting Personal Privacy*

personal information

7.7: *Where does all that Data come from?*

personal privacy

7: *Legal Aspects of AI*

personalization

3.4: *AI in Marketing Strategies*

3.6: *Summary*

personalized customer experiences

3.4.1: *Good and Bad*

personalized learning

3.2.1: *Benefits and Concerns*

3.3: *AI Transforming Education*

3.3.1.1: *Benefits to Students*

3.3.1.2: *Benefits for Teachers*

physical activity

3.7.1: *Case Study 1 - AI for Personal Use*

predictive algorithms

3.4.1: *Good and Bad*

predictive analytics

3.4.1.1: *AI for Entrepreneurs*

privacy

3: *The Current State of AI*

3.5: *Privacy and Safety Challenges with AI*

7.8: *Protecting Personal Privacy*

privacy concerns

7.9: *Case Study- AI Defamation*

privacy rights

7.7: *Where does all that Data come from?*

productivity

3.1: *AI in Everyday Life*

profitability

3.7.3: *Case Study 3 - AI for Entrepreneurs*

public domain

7.4: *Writing, Film, and Music*

public school

3.7.2: *Case Study 2 - AI in Education*

R

recommendation systems

3.1: *AI in Everyday Life*

regulatoryCompliance

7.10: *Teaching AI Ethics (1)*

robotics

1.1: *Chapter Introduction*

robots

1.4: *Applications of AI*

S

safety

3: *The Current State of AI*

scheduling

3.3.1.2: *Benefits for Teachers*

sequence modeling

2.6: *The Age of Generative AI (2020s - Present)*

Siri

3.1: *AI in Everyday Life*

sleep analysis

3.7.1: *Case Study 1 - AI for Personal Use*

small business ethics

3.7.3: *Case Study 3 - AI for Entrepreneurs*

smart devices

1.5: *Emerging Technology Trends*

smart home devices

3.1: *AI in Everyday Life*

software algorithms

1.1: *Chapter Introduction*

solutions

3: *The Current State of AI*

Stability AI

7.2: *Copyright and Image Generation*

strategic thinking

2.4: *Machine Learning and the Data Revolution (1990s - 2000s)*

streamlining

3.3.1.2: *Benefits for Teachers*

student engagement

3.3.1.1: *Benefits to Students*

student performance

3.3.1.2: *Benefits for Teachers*

super AI

1.3: *Types of AI*

surveillance

7.8: *Protecting Personal Privacy*

T

tailored content

3.3.1.1: *Benefits to Students*

Talos

2.1: *The Early Foundations (Pre-1950s)*

teaching

3.7.2: *Case Study 2 - AI in Education*

7: *Legal Aspects of AI*

technology

1.5: *Emerging Technology Trends*

technology ethics

3.5: *Privacy and Safety Challenges with AI*

technology trends

2.5: *The Deep Learning Era (2010s)*

text generation

7.4: *Writing, Film, and Music*

tools

3.4.1.1: *AI for Entrepreneurs*

transformers

2.6: *The Age of Generative AI (2020s - Present)*

transparency

7.7: *Where does all that Data come from?*

treatment planning

3.4.1: *Good and Bad*

Turing Test

2.1: *The Early Foundations (Pre-1950s)*

Turing_Test

2: *The History of AI*

types of AI

1.1: *Chapter Introduction*

types_of_AI

1: *Introduction*

U

U.S. Copyright Office

7.2: *Copyright and Image Generation*

UAV

1.4: *Applications of AI*

user awareness

3.5: *Privacy and Safety Challenges with AI*

user consent

3.7.1: *Case Study 1 - AI for Personal Use*

userPrivacy

7.10: *Teaching AI Ethics (1)*

V

virtual assistants

3.1: *AI in Everyday Life*

3.4: *AI in Marketing Strategies*

visual arts

7.4: *Writing, Film, and Music*

voice commands

3.1: *AI in Everyday Life*

W

wearable

1.5: *Emerging Technology Trends*

wellness app

3.7.1: *Case Study 1 - AI for Personal Use*

writing

7: *Legal Aspects of AI*

X

XCON

2.3: *The AI Winters and Specialization (1970s - 1980s)*

