8 Artificial Intelligence and Machine Learning: Foundations

Learning Objectives

After reading this chapter, you should be able to:

LO 8.1 Understand the historical context of AI development.

LO 8.2 Know the function of convolutional neural networks.

LO 8.3 Understand the development and function of large language models.

LO 8.4 Know how supervised, unsupervised, and reinforcement learning impact machine learning capabilities.

LO 8.5 Grasp the dynamics of deep learning.

LO 8.6 Understand the possible inputs and outputs of generative artificial intelligence.

This chapter provides a comprehensive explanation of **artificial intelligence (AI)** and **machine learning (ML).** Students will gain insight into the fundamental concepts and principles that underpin AI and ML as the chapter delves into the technological foundations of these technologies. Various types of ML algorithms are explored, including supervised, unsupervised, and reinforcement learning, elucidating their applications in real-world scenarios. Additionally, we elucidate the crucial role of data, models, and algorithms in the development and deployment of AI and ML systems.

The next chapter moves beyond the technological realm and focuses on the business applications of AI and ML. Through both chapters, students will have developed a holistic understanding of AI and ML technologies, their broader business implications, and their specific applications within the finance domain.

8.1 Historical Development of AI and ML

The transformative power of innovation is reshaping our world into a **digital twin.** As depicted in Figure 8.1, every piece of information, every interaction, and every process is being digitized and modeled. These digital twins are used to simulate how the real system will react to failure points, maintenance, and integration with other systems.

8.1.1 Milestones in Artificial Intelligence

This section highlights the significant milestones of human progress, with a particular focus on the latest phase of development. Throughout history, humanity has experienced three major economic revolutions: (1) the Agricultural Revolution, which originated in Great Britain during the 17th century, followed by (2) the Industrial

FIGURE 8.1

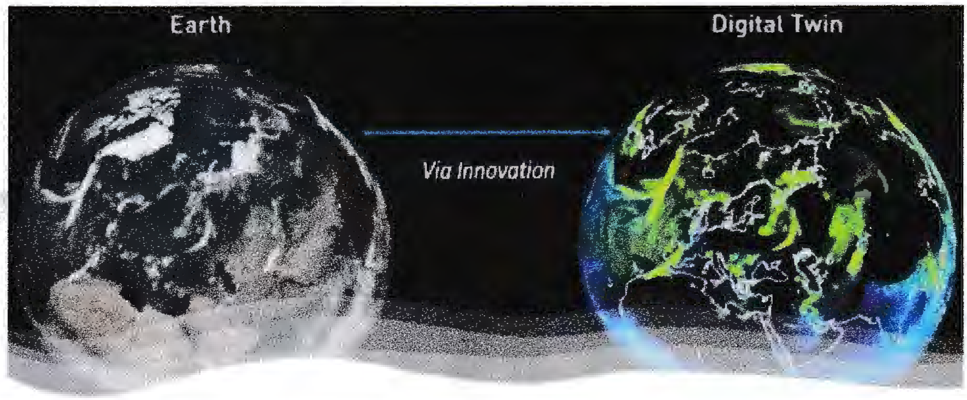

Digital twins replicate physical systems or processes
NVIDIA/YouTube

FIGURE 8.2

Technological revolution milestones

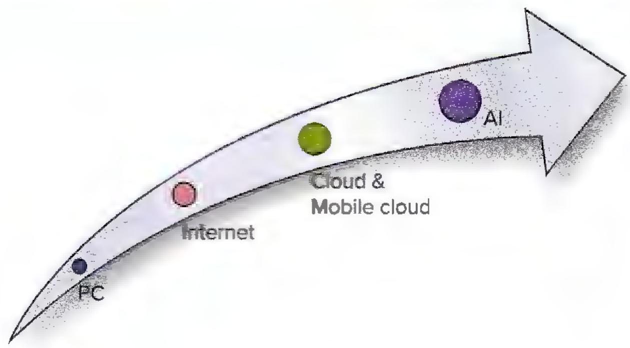

Revolution in 1760. This marked a shift toward new manufacturing methods in Europe and the United States, ushering in the era of machinery and large-scale farming. Now, with the advent of personal computers, PC development, and the exponential growth of the internet, we are in the midst the third revolution: (3) the Technological Revolution, often referred to as the Tech-Rev.

Figure 8.2 shows the major technological events within Tech-Rev. The creation of the personal computer led to the internet, **cloud computing,** and now AI. Each innovation builds from the previous one and magnifies technology's societal and economic impact.

The technological revolution began with the transistor chip. Gordon E. Moore was co-founder of Intel Corporation (INTC). In 1965, he quipped about what is now commonly referred to as Moore's law. This "law" observes that the number of transistors in an integrated circuit roughly doubles every two years. It's not merely a principle but rather an empirical observation tied to the efficiencies gained through experience in production, rooted in the laws of physics. Figure 8.3 illustrates the doubling of transistors on microchips approximately every two years from 1970 to 2020. There is much speculation that although computer power continues to advance, it is at a slower pace due to limitations of physics and heat transfer. However, industry experts point to other developments such as AI chips and quantum computing to drive computing power development in the future.

Skipping over the well-known Internet Revolution leads to a discussion of cloud and mobile cloud technological innovations. Cloud technology encompasses three main components:

- **On-premises resources.** This refers to a computing environment where data and applications are stored and processed on physical hardware situated within the organization's premises. While offering greater control and security compared to cloud-based environments, it can be more costly and challenging to manage.

artificial intelligence (AI)
The simulation of human intelligence processes by machines, especially computer systems.

machine learning (ML)
A subset of AI that enables machines to learn from data and improve from experience without being explicitly programmed.

digital twin
A virtual representation of a physical object or process that can be used to simulate and monitor its behavior in real time.

cloud computing
The delivery of computing services over the internet.

on-premises resources
A computing environment where data and applications are stored on physical hardware situated within the organization's premises.

FIGURE 8.3

Number of transistors on microchips

Moore's Law: The number of Transistors on Microchips has doubled every two yearshttps://ourworldindata.org/moores-law, https://creativecommons.org/licenses/by/4.0/

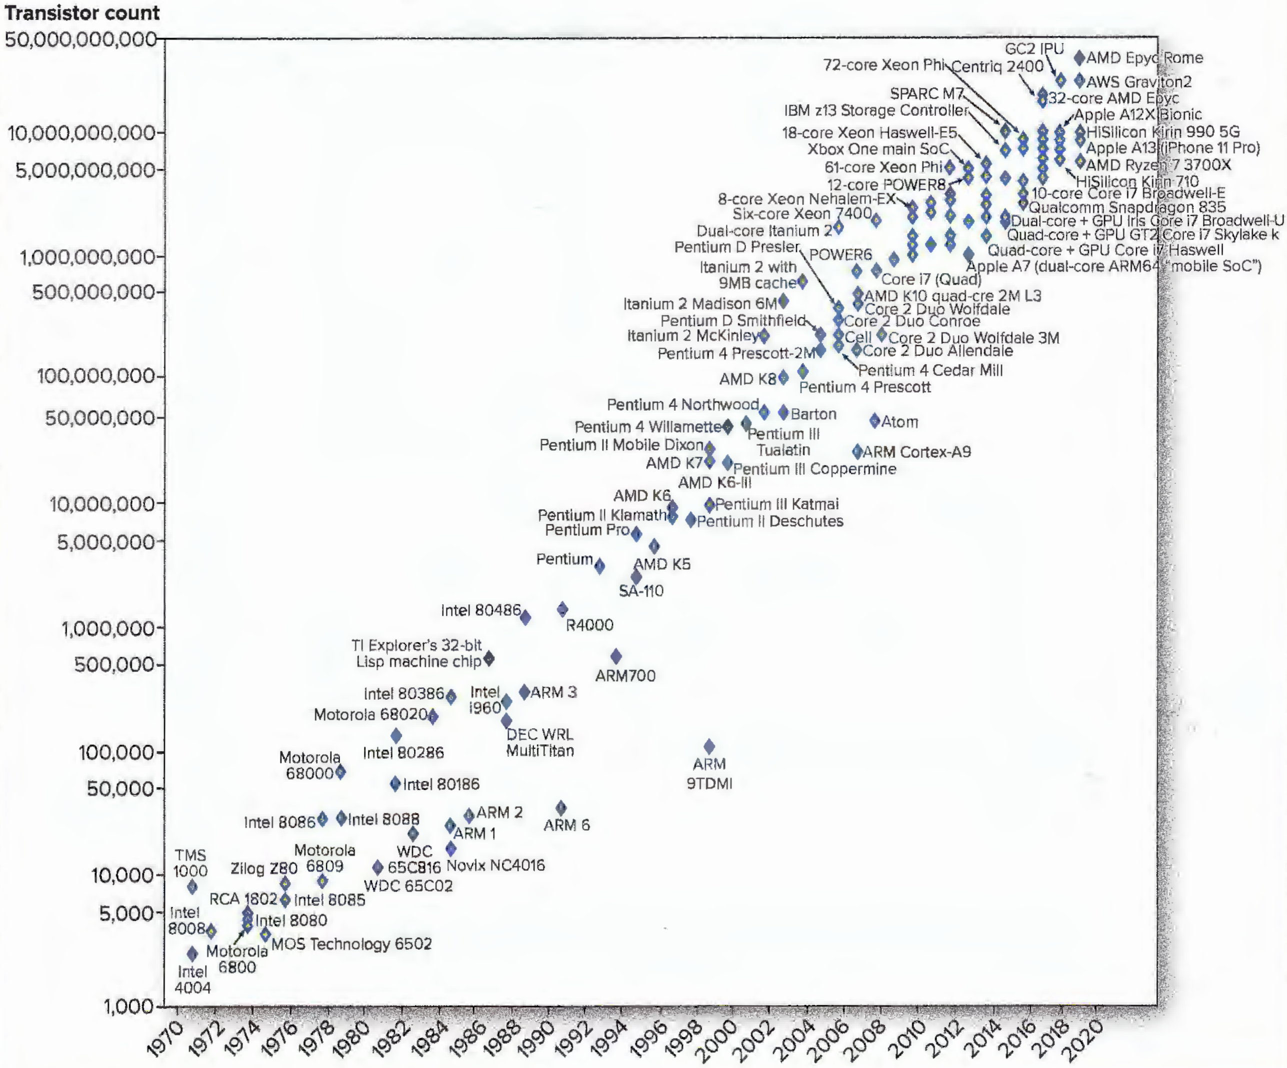

- **Hybrid cloud.** A hybrid cloud integrates on-premises resources with cloud-based resources, enabling organizations to leverage the advantages of both environments. For instance, on-premises resources can be used for high-security or low-latency applications, while cloud-based resources offer scalability and flexibility.

hybrid cloud

Integrates on-premises resources with cloud-based resources.

- **Edge computing.** Edge computing is a distributed computing paradigm that brings computation and data storage closer to the edge of the network. This proximity enhances performance, reduces latency, and boosts security. Edge computing finds applications in real-time processing scenarios such as self-driving cars and augmented reality.

edge computing

Distributed computing that brings computation and data storage closer to the edge of the network.

Table 8.1 summarizes the key distinctions between on-premises, hybrid cloud, and edge computing. Note that the three main components of cloud computing have very different attributes. On-premises resources have high control and security and cost, while edge cloud does not. Hybrid cloud and edge cloud have a high degree of scalability, while on-premises does not.

TABLE 8.1 | **Components of Cloud Technology**

Feature	On-Premises	Hybrid Cloud	Edge Cloud
Data and application location	On-premises	Both on-premises and cloud	Edge of the network
Control and security	High	Medium	Low
Cost	High	Medium	Low
Management complexity	High	Medium	Low
Performance	Good	Excellent	Excellent
Latency	Low	Very low	Very low
Security	Good	Good	Low
Flexibility	Low	High	High
Scalability	Low	High	High

The next frontier of innovation lies in software-driven technology companies, marking a transition dubbed as a "Software 2.0" upgrade. Specifically, this refers to the augmentation of software programmers by machines. Rather than humans manually coding software programs, data and experiences now contribute to software writing, with software itself making future predictions. AI represents the future, democratizing computing by bridging digital and technological divides. It empowers anyone to become a programmer, disrupting the traditional paths of learning languages like Fortran, C++, Java, and Python. With technology now accessible to all, AI serves as a tool for harnessing and connecting.

In the realm of computing, enhancing the power of data centers that support sustainable computing is crucial. For instance, application-specific integrated circuits (ASICs) provide flexibility, versatility, and high performance tailored to specific applications. Developers are increasingly focused on simulating climate, advancing robotics, and bolstering **large language models (LLMs)** and **neural networks.** Meeting these challenges requires further hardware development, supported by innovations like the Omniverse, which connects the digital and physical worlds across various industries such as automotive, logistics, and renewable energy.

large language models (LLMs)

A type of AI that can understand, process, and generate human language.

neural networks

Computing systems inspired by the human brain's network of neurons, used to recognize patterns, and classify data.

AI hallucinations

Occurs when an AI model produces incorrect or misleading information, often with confidence.

The applications of AI and machine learning span wide-ranging fields like digital biology, genomics, transportation, and retail. However, the exponential growth in areas like science and industry, span ray tracing, physical earth and life sciences, quantum physics and chemistry, and computer vision and data processing will also play important roles and challenges. These challenges lead us to the critical consideration of **AI hallucinations,** namely, incorrect or misleading results that AI models generate.

8.1.2 Artificial Intelligence as an Emerging Technology

Alan Turing (1912–1954), an English mathematician, logician, cryptanalyst, and computer scientist, is widely regarded as the pioneer of both computer science and AI. Renowned for his pivotal role in breaking German Enigma codes during World War II, Turing's legacy extends beyond his scientific achievements. Despite enduring persecution for his homosexuality, his intellect and contributions remain seminal. An anecdote often recounted is the discovery of an apple with a few bites taken from it on his table after his death, leading to speculation that this inspired its eventual use as a company logo.

AI is synonymous with the advancement of computer systems, aiming to mimic and replicate human cognitive capabilities. Its objective is to create intelligent machines proficient in

- perception,
- learning,
- prediction,
- problem-solving,
- discovery, inference, and cognition.

AI is indispensable for swift processing and responses, facilitating cost-effective solutions and enhanced user experiences. Moreover, it aids in reducing data storage costs. The surge of AI is propelled by the emergence of vast datasets and heightened computer processing speeds, which enables complex computational tasks. Consequently, AI finds applications across various domains including statistics, computer science, psychology, biology/medicine, and engineering.

Applications are being partially driven by the development of pivotal technologies like natural language processing, computer vision, text-to-speech, robotics motion, machine and deep learning, and auditory systems. AI development can be grouped into three categories.

1. **Narrow AI.** Focusing on specific tasks or problems such as voice assistants and image recognition systems.
2. **General AI.** Resembling human intelligence, capable of understanding, learning, and applying intelligence across diverse tasks and domains.
3. **Artificial superintelligence.** Surpassing human intelligence in all aspects, including creativity, problem-solving, and decision-making.

narrow AI
AI model focused on specific tasks or problems.

general AI
AI model capable of understanding, learning, and applying intelligence across diverse tasks and domains.

artificial superintelligence
A level of intelligence surpassing that of the brightest and most gifted human minds.

As the intelligence and creativity levels of AI increase, so do societal concerns about its impact on humanity. Three major topics of these concerns are regulation, systemic risk, and ethics. *Regulation* encompasses the regulation of technology (RegTech) and the use of technology by supervisors (SupTech). *Systematic risk* comprises exogenous and endogenous risks, including cybersecurity threats and synchronized behavior of AI entities within systems. *Ethics* focuses on issues such as bias in learning datasets and algorithms, particularly concerning racism and sexism. Organizations like the ACLU and the Algorithmic Justice League work against biased algorithms, while independent researchers contribute to ongoing discourse. These issues are discussed in more detail in Chapter 14.

As AI continues to advance, unanswered questionable issues persist, highlighting the need for ongoing research and ethical deliberation. For example,

- Ethical concerns regarding personal data and bias mitigation.
- Regulatory compliance and accountability.
- Performance robustness and hallucinations.
- Data governance and transparency.
- Fairness and nondiscrimination in AI development and content creation.
- The principle of proportionality in AI applications.

8.2 AI Employed Technology

Numerous technological algorithms contribute to the AI ecosystem, but two primary technologies are highlighted, the **convolutional neural network** (CNN) and the LLMs.

convolutional neural network (CNN)
A deep learning algorithm primarily used for analyzing visual data.

8.2.1 Convolutional Neural Network

A CNN is a deep learning network that uses a technique called convolution, a specialized kind of linear mathematics operation, to extract features from images and other gridlike data. CNN stands out as a highly effective tool for image recognition, demonstrating remarkable proficiency across various tasks including object recognition, facial recognition, and image categorization. In addition, it can be quite effective for classifying audio, time-series, and signal data. It is inspired by the biological processes in the visual cortex of animals.

The process is too complex for this text, but essentially, CNNs apply filters to an input image to detect features like edges or shapes within the picture. The filters produce a digitized activation map that reduce dimensionality to reduce processing speed. Images

backpropagation

A machine learning algorithm that trains artificial neural networks by adjusting weights to reduce errors.

within the activation map are classified into categories, such as "tree" or "dog." A CNN is trained on a large data set and uses **backpropagation** to update the weights and biases of image identification. Backpropagation is an algorithm used to train artificial neural networks by error correction through calculating the difference between the input provided and the output it produced. With enough data and computing power, CNNs can achieve superhuman performance on many visual tasks.

Some popular CNN architectures are AlexNet, VGGNet, ResNet, and Inception. These have been used to solve complex problems like identifying thousands of objects or detecting diseases from medical scans. Popular applications are as follows.

- *Automotive.* Image recognition will be critical for the age of driverless cars. Until then, technology is currently improving driver and passenger safety through features like lane line detection.
- *Facial recognition.* Social media platforms analyze photographs and make suggestions on identifying those people present, making it easier to tag friends in photo albums. Face recognition is also useful in identifying a person from surveillance camera footage.
- *Medical imaging.* Computer vision detects anomalies in X-ray and MRI images with better accuracy than the human eye, enabling doctors to better identify cancerous tumors.
- *Retail.* CNNs help drive personalized advertising and purchase recommendations.
- *Document analysis.* This process can identify an individual's handwriting. This is important for crime investigation and legal actions. In addition, just as text and audio can be converted back and forth, handwriting can be converted to text.

The early success of CNN was exemplified by its very high prediction rate, that is, very low error rate. This success was attributed to the depth of the model, which contributes to its high performance despite its computational complexity. Notably, the use of graphics processing units (GPUs) during training makes this level of performance feasible.

8.2.2 Large Language Models

LLMs are trained on immense amounts of data to make them capable of understanding and generating natural language. LLMs can serve many different purposes, depending on how they are trained. Common tasks are answering questions, summarizing documents, translating languages, and completing sentences. They are designed to understand and generate text like a human. Also like humans, they can generate coherent and contextually relevant responses, summarize text, infer from context, translate to other languages, answer questions, and assist in creative writing or code generation tasks.

To do these tasks, billions of parameters enable them to capture intricate patterns in language. LLMs are revolutionizing applications in various fields, from chatbots and virtual assistants to content generation, research assistance, and language translation. LLMs are easily accessed through interfaces such as Open AI's Chat GPT-3 and GPT-4, Meta's Llama models, and Google's BERT/RoBERTa and PaLM models.

The training process involves the program learning to predict the next word in a sentence based on the context provided by the preceding words. For example, "A dog likes to chew on a ______________." Note that the English language encompasses more than 100,000 words, with an average sentence consisting of about 10 words. The model predicts the next word through attributing a probability score to the numerous possibilities. It is the training process that helps the LLM to determine the probability scores. It also must learn grammar, semantics, and conceptual relationships.

If the underlying neural network is large and has been trained on enough data, the LLM becomes really good at predicting the next word. Fortunately, there is an abundance

FINTECH IN THE NEWS

Nobel Prize in Physics Awarded for the Building Blocks of Machine Learning

Geoffrey Hinton and John Hopfield won the 2024 Nobel Prize in Physics for inventions that enable machine learning with artificial neural networks. The prize is awarded by the Royal Swedish Academy of Sciences and comes with an 11 million Swedish kronor award, equivalent to about $1 million.

Artificial neural networks are inspired by the abilities of the human brain and the network of neurons that enable learning. Neural networks are programs that help machines learn. ML aims to mimic the human ability to recognize images and speech and associate them with memories and past experiences. ML does this by feeding computers incredible amounts of data so they can master tasks.

The scientists' work used fundamental concepts from statistical physics to help train artificial neural networks to recognize patterns in large data sets. Thus, computers can reconstruct images and perform tasks like identifying specific elements in pictures. Such machine learning contributes to aspects of our daily lives such as facial recognition and language translation.

Hinton also won the Association for Computing Machinery's Turing Award in 2018 for breakthroughs in artificial intelligence. He is often referred to as one of the "godfathers of AI." Not just an academic, he also worked for Google for more than a decade. He quit Google in 2023 in order to speak more freely about the risks of AI development. He has expressed alarm over the ways AI could harm humanity and reportedly expressed his worry to the Nobel committee when he was informed of the prize.

of text on the internet in books, research papers, and more (in the billions of pages), and after the LLM is exposed to all of this material, it will figure out the answer. This is called self-supervised learning. Once trained, LLMs can generate text by autonomously predicting the next word based on the input they receive and drawing on the patterns and knowledge it has acquired. The result is coherent and contextually relevant language generation that can be harnessed for content generation tasks.

Training an LLM takes an enormous amount of computing power. Using a single GPU or computer, training ChatGPT-3's LLM would require approximately 334 years. However, leveraging highly parallelizable computers, this process can be completed in roughly a month with the assistance of thousands of GPUs. Most LLMs take about three months to train.

8.3 Machine Learning and Generative AI

Many technological innovations have led to artificial intelligence. Figure 8.4 gives them some context. This section discusses machine learning and AI.

8.3.1 Machine Learning

Machine learning functions as a subset of AI, where ML applications develop algorithms and models empowering computer systems to learn and make estimates or decisions autonomously, without explicit programming. ML algorithms discern patterns, relationships, and insights to execute tasks by employing statistical techniques to learn patterns and make decisions or predictions.

FIGURE 8.4

Layers of technology leading to artificial intelligence

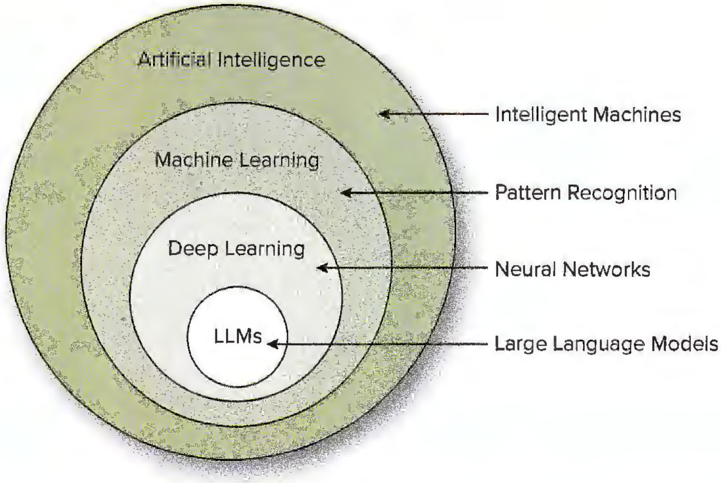

FIGURE 8.5

Stages of supervised learning

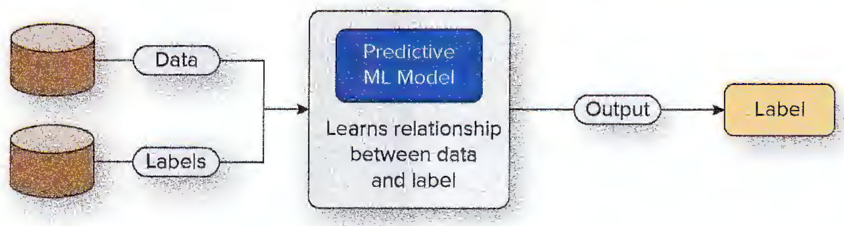

supervised learning (SL)

An ML approach where the model is trained on labeled data to make predictions.

predictive ML

A statistical technique that uses machine learning algorithms to analyze historical data and predict future outcomes.

unsupervised learning (USL)

An ML approach where the model identifies patterns and relationships in unlabeled data.

reinforcement learning (RL)

A type of ML where an agent learns to make decisions by receiving rewards or penalties.

The three primary types of ML are

- Supervised learning,
- Unsupervised learning, and
- Reinforcement learning.

These learning models encompass distinct problem types, data structures, algorithms, feedback mechanisms, and applications. **Supervised learning (SL)** uses a training data set, which includes inputs and correct outputs, to teach models to yield the desired output. The algorithm adjusts until the error has been sufficiently minimized. One form of SL uses a classifying algorithm to assign test data into specific categories. Another form of SL often employs established statistical models like linear regression, logistical regression, and polynomial regression to make projections. SL yields highly accurate, efficient, and labeled data, and is commonly referred to as **predictive ML.**

SL progresses through three stages, as shown in Figure 8.5: input (data or labeled data), followed by the exploration of relationships between variables in the data, and finally output and labeled output.

Unsupervised learning (USL) uses algorithms to analyze and cluster unlabeled data sets without human supervision. For training, models are given unlabeled data and allowed to discover patterns and insights on their own. USL can be applied to unlabeled data such as photo images, audio/video recordings, articles, tweets, or news. Figure 8.6 outlines the distinctions between SL and USL.

One learning technique, **reinforcement learning (RL),** mimics the human trial and error process. It trains the algorithm to make decisions to achieve the most optimal results over time by engaging with its environment, executing actions, learning from them, and navigating through errors and pathways. Netflix and Amazon have leveraged RL applications to refine their suggestions to users.

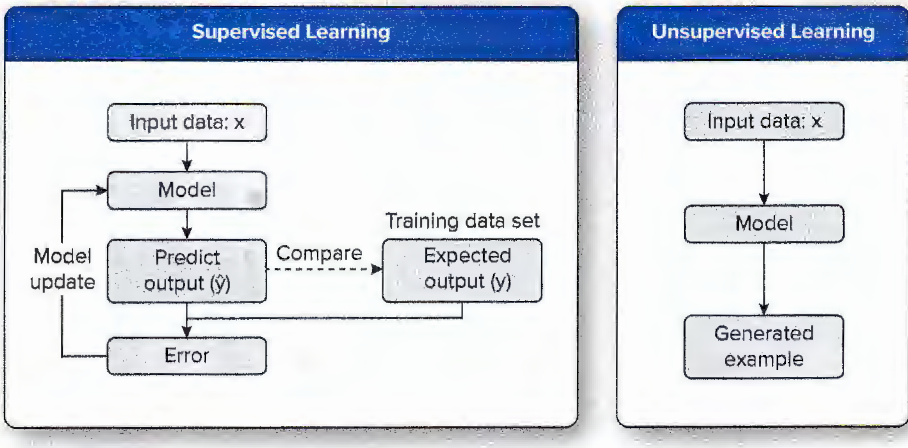

FIGURE 8.6

Differences between supervised learning and unsupervised learning

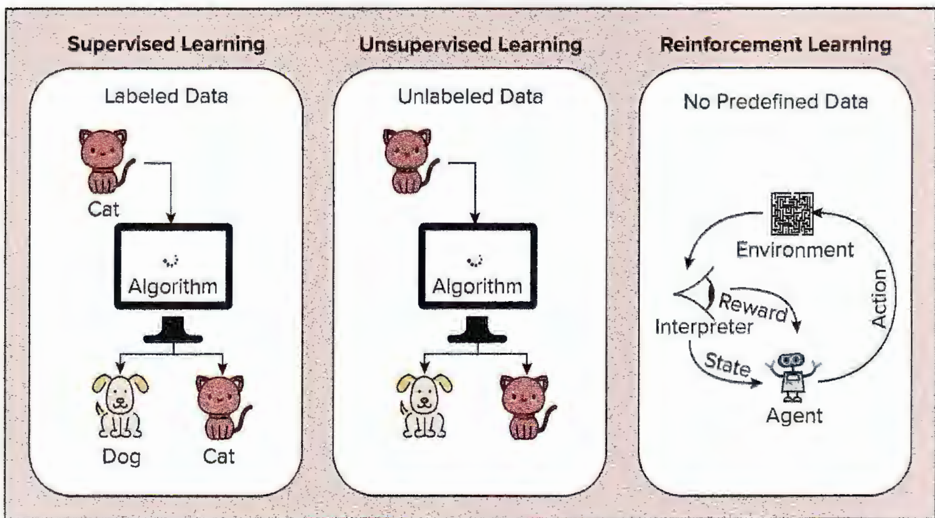

FIGURE 8.7

Data types used for machine learning

www.medium.com

The three types of learning processes use data differently. As shown in Figure 8.7, SL uses data for training with labeled data inputs and known outputs to refine its model. Once trained, the model examines new data to make predictions to label the output. In the figure, a picture of a cat is labeled and processed so the output of a cat is correct. In the next panel, USL trains itself with unlabeled data to learn categories. Here, the input of a cat is not labeled, and the model categorizes this new data. In the right panel, RL learns real time through trial and error by interacting directly with its environment. The figure shows it is an iterative process of learning.

8.3.2 Machine Learning Examples

Machine learning applications exhibit diversity based on the learning paradigm employed. In particular, supervised learning finds utility in risk assessment and prediction. For instance, consider the scenario of loan approval or rejection. The inputs into such a decision might include income, wealth, length of current job, home ownership, and other criteria. Figure 8.8 illustrates that inputs like the amount of credit are coded along with debt duration. In addition, customer information like age, gender, income, and the like is added. If a bank possesses customer data, SL algorithms can utilize such variables (the X-axis) and credit history (the Y-axis) to construct a predictive model. Leveraging this model, the bank can determine whether to approve or decline a loan application.

FIGURE 8.8

Using supervised learning to assess loan applications
Microsoft Excel

```
$ Account . Balance                      : int   1 1 2 1 1 1 1 1 4 2 ...
$ Duration . of . Credit . . month .     : int   18 9 12 12 12 10 8 6 18 24 ..
$ Payment . Status . of . Previous . Credit : int   4 4 2 4 4 4 4 4 4 2 ...
$ Purpose                                : int   2 0 9 0 0 0 0 0 3 3 ...
$ Credit . Amount                        : int   1049 2799 841 2122 2171 2241
$ Value . Savings . Stocks               : int   1 1 2 1 1 1 1 1 1 3 ...
$ Length . of . current . employment     : int   2 3 4 3 3 2 4 2 1 1 ...
$ Instalment . per . cent                : int   4 2 2 3 4 1 1 2 4 1 ...
$ Sex . . . Marital . Status             : int   2 3 2 3 3 3 3 3 2 2 ...
$ Guarantors                             : int   1 1 1 1 1 1 1 1 1 1 ...
$ Duration . in . Current . address      : int   4 2 4 2 4 3 4 4 4 4 ...
$ Most . valuable . available . asset    : int   2 1 1 1 2 1 1 1 3 4 ...
$ Age . . years .                        : int   21 36 23 39 38 48 39 40 65 23
$ Concurrent . Credits                   : int   3 3 3 3 1 3 3 3 3 3 ...
$ Type . of . apartment                  : int   1 1 1 1 2 1 2 2 2 1 ...
$ No. of . Credits . at . this . Bank    : int   1 2 1 2 2 2 2 1 2 1 ...
$ Occupation                             : int   3 3 2 2 2 2 2 2 1 1 ...
$ No . of . dependents                   : int   1 2 1 2 1 1 2 1 1 1 ...
$ Telephone                              : int   1 1 1 1 1 1 1 1 1 1 ...
$ Foreign . Worker                       : int   1 1 1 2 2 2 2 2 1 1 ...
```

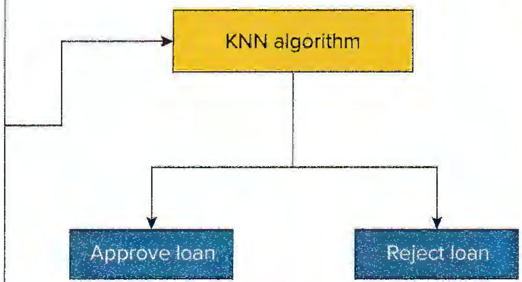

Now consider the unsupervised learning approach in evaluating a movie's popularity. Much data is available from prior movies on the popularity of the director, actors, and more. Note that an actor might be associated with both a blockbuster film and a bust. Thus, even past data is not always clear. Nevertheless, the metric for movie popularity is not explicitly defined. In such cases, USL algorithms utilize the available data to discern varying levels of movie popularity autonomously.

Finally, a reinforcement learning application occurs in auto-stock trading. Prior to RL use in stock trading, SL was used. As an example, consider that with human traders, an SL algorithm buys the S&P500 Index whenever it closes above its past 30 days high and sells the position whenever it closes below the past 30 day low. For any technical approach like this to be systematic, the trader specifies precise conditions to unambiguously define whether a particular pattern had been formed and whether to buy, sell, or hold. Indeed, EquBot and Horizons have launched supervised machine learning exchange traded funds (ETFs) that automate the compiling of market information from thousands of U.S. companies. They process market signals, quarterly news articles, and social media postings to make stock price predictions. The human traders help to choose thresholds, account for latencies, estimate fees, and so on.

To be fully automated, an AI-driven trading model must do more than predict prices. It must make buy, sell, or hold decisions, too. These decisions include timing and dollar value sizes. Reinforcement learning allows for complete optimization that maximizes the rewards to trading. The RL agent adjusts the parameters without a human trader to identify the optimal decision.

8.3.3 Deep Learning

deep learning

A type of neural network with many layers that can learn increasingly abstract representations of data.

Deep learning (DL) constitutes a subset within the realm of machine learning, characterized by the use of multiple layers of neural networks. These algorithms demand substantial quantities of labeled data and considerable computational power for training. By adjusting weights and biases, DL models undertake computations to generate outputs. The difference between regular ML and DL might be characterized as such: ML uses algorithms to learn patterns from data, while DL employs artificial neural networks to learn and make intelligent decisions on its own. Once trained, DL handles unstructured data much better than ML.

DL comprises two very different methods for understanding and retrieving knowledge. Consider this analogy of two students and how they learn. One student, Gianna, learns every in-depth detail about each topic. This process takes a lot of time and brain storage space. This is like the **generative DL** model of learning. The other student, Nora, gets the general idea of a topic and then learns the differences and patterns between the subtopics. She learns by learning the differences. This is like the **discriminative DL** model of learning. Note that the mathematical processes and algorithms would need to be different between the two methods for accessing information, providing answers, and completing tasks.

generative DL

A model that aims to learn the underlying patterns or distributions of data in order to generate new, similar data.

discriminative DL

A model that separates the data points into different classes and learns the boundaries using probability estimates and maximum likelihood.

Discriminative models divide the data inputs into categories by learning the boundaries as a part of the process. Generative models have data categories and input labels. Both approaches are very different, which makes them suited for specific tasks. Discriminative models are more effective and work well for DL supervised tasks. Generative models are better for DL unsupervised tasks.

For example, a discriminative AI model could be trained on a data set of images of cats and dogs and then used to classify new images as either cats or dogs as shown in Figure 8.9. A generative AI model could be trained on a data set of images of dogs and then used to generate new images of dogs.

8.3.4 Generative AI

This section discusses **generative artificial intelligence (GenAI),** addressing key questions such as: What exactly is GenAI? How does GenAI operate? And what categories and applications does GenAI encompass?

generative artificial intelligence (GenAI)

A type of AI that creates new content, such as text, images, videos, and music, based on a prompt.

GenAI functions as a pattern-matching mechanism that generates novel content by leveraging statistical models trained on existing data. Operators employ prompts to guide GenAI in crafting fresh content. The operational mechanics of GenAI are summarized in Figure 8.10.

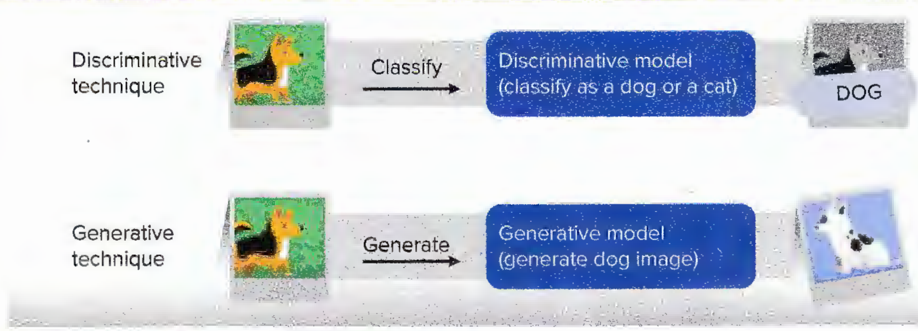

FIGURE 8.9

DL discriminative and generative models

McGraw Hill LLC

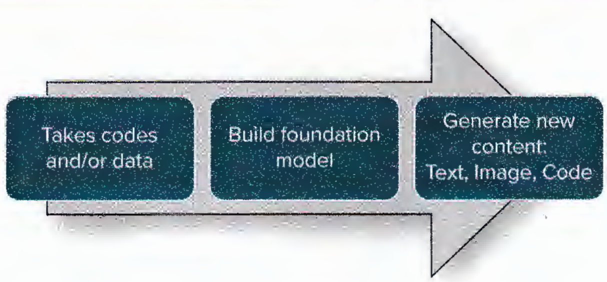

FIGURE 8.10

Operational mechanics of GenAI

Consider the example in Figure 8.11. A user initiates a prompt, like "How's it going?" The input is transformed by encoding and decoding components to a function the AI can understand. An appropriate response is determined, and the new content is output to the user in either text, image, or code. In this case, the text is "I'm doing alright, thanks for asking. How are you?" Note how the form of the input is mimicked in the output. Specifically, both use contractions.

The figure is an example of a text output. However, there are three types of GenAI outputs: text, images, and video.

The first GenAI produced text. Open AI released ChatGPT in 2022, which was a conversational chatbot based on generative pretrained transformer models. The platform reached 1 million users within five days. A newer version, GPT-4, was released the following year. It is reportedly more accurate and has advanced reasoning capabilities. Google released its first GenAI model as Bard. Later, the revised version was renamed Gemini. It combines natural language processing with advanced multimodal capabilities to generate high-quality text and integrate seamlessly with Google's suite of tools.

AI image generators take a text prompt and then turn it into a matching image. Several popular GenAI image generators are available, including DALL-E3, Adobe Firefly, and Generative AI by Getty. Midjourney also provides accurate text with its AI images, which is useful for search engine discoverability. GenAI can generate very creative images.

A GenAI video generator is a tool that enables the user to generate videos from text. The most popular one is Synthesia. It has 60 templates with 160 lifelike avatars and is commonly used to make training videos. Other popular video generators, like Colossyan, Hour One, D-ID, Elai, HeyGen, Runway, Pictory, Deepbrain AI, InVideo, and Fliki, have varying degrees of realism.

Some GenAI models accept images as an input and can generate text about them, create other images, or generate video, as shown in Figure 8.12. For example, CornellLab's Merlin app will identify the bird in a photo. That is, it reads the photo and outputs text. DALL-E can receive an image and manipulate it as requested. Programs like Canva and Monica will animate an uploaded photo.

What delineates GenAI and predictive machine learning? GenAI engages in more intricate computations as it lacks a predefined model for output prediction. Instead, GenAI scrutinizes input and explores potential patterns within the data, essentially learning from patterns inherent in unstructured data, as illustrated in Figure 8.13. In contrast to predictive ML where the dependent variable is numeric and known, GenAI produces non-numeric outputs, using natural language data such as Yelp restaurant reviews,

FIGURE 8.11

GenAI interaction example

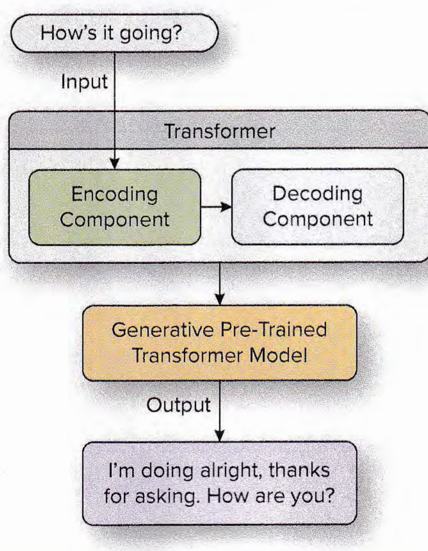

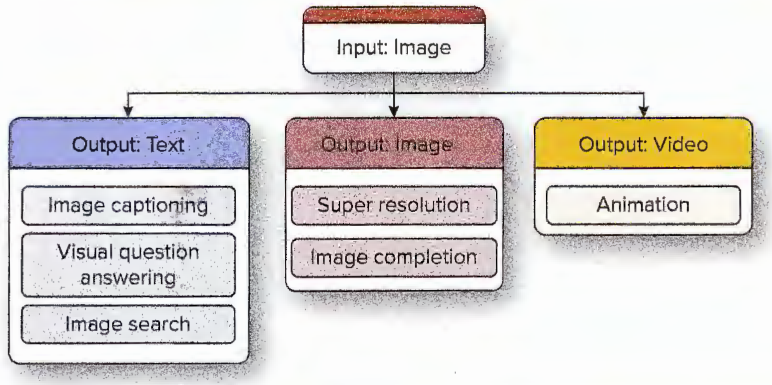

FIGURE 8.12

GenAI image processing

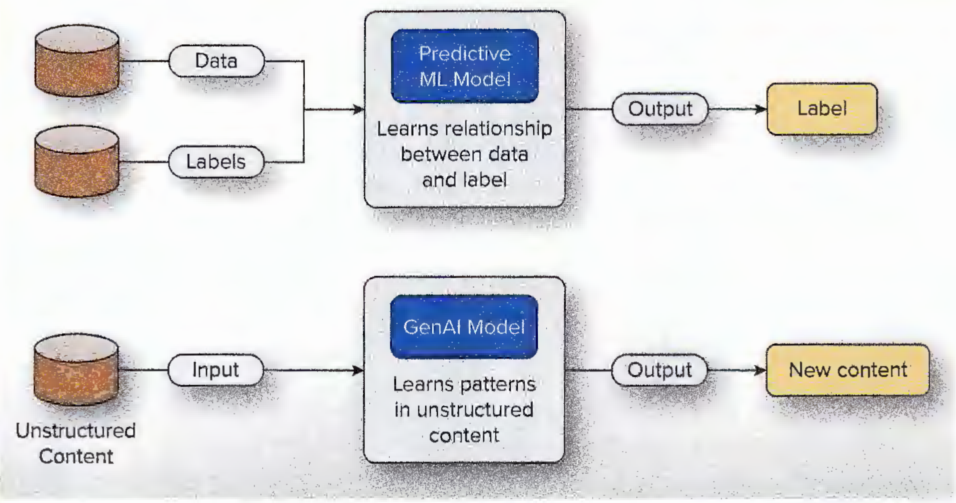

FIGURE 8.13

GenAI vs. predictive machine learning

images, and audio to generate plausible outputs. For instance, if a customer's restaurant review is condensed into a paragraph, GenAI can transform this text into an opinion about the restaurant.

There are numerous applications for AI. Figure 8.14 illustrates many capabilities that are applicable to all areas of society. The next chapter focuses on AI and ML applications, particularly those in finance.

8.3.5 Concerns About Artificial Intelligence

AI is purported to have the potential to be a more powerful positive force for society than any other technology in history. It has already shown immense potential to improve efficiency, aid research and discovery, enhance decision-making, and solve complex problems—and we have barely scratched the surface of its potential. It also stands to reason that it could have the same potential to negatively impact society. *What are the significant ethical concerns to be considered?*

Consider these potential negative impacts on society.

- *Privacy.* From facial recognition to smart home devices, the potential for invasions of privacy abounds.
- *Data bias.* AI systems are only as good as their training data. Biased data will create biased AI outputs. For example, Amazon stopped using a hiring algorithm because it chose candidates with words like "executed" or "captured" on the resumes, which were more common for men.

FIGURE 8.14

GenAI applications

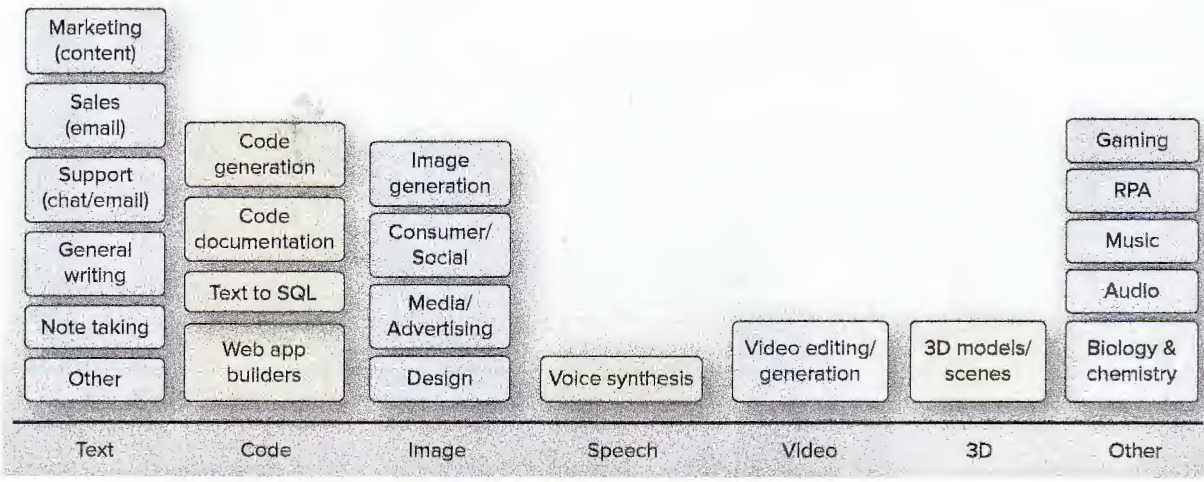

- *Accountability.* AI systems are making more autonomous decisions that impact people's lives. It is critical to determine clear lines of responsibility when errors occur. For example, who is accountable when an autonomous vehicle has an accident? Or when faulty AI healthcare diagnoses are made and decisions about medications or therapies are carried out?
- *Job displacement.* New technologies lead to job losses in old industries and job creation in new ones.
- *Misinformation.* AI users might create and spread misinformation to impact elections through social media, deepfake videos, and other means.
- *Security.* There are two issues of AI security. The first is that AI systems have been hacked, leading to misbehavior in autonomous vehicles or security camera footage. The second issue is that AI can be used to infiltrate personal and company security.

The capability of AI is developing rapidly. Superintelligent machines are likely to exist sometime in the future. Anyone who has seen movies such as *2001: A Space Odyssey, The Terminator, The Matrix, I Robot, Ex Machina,* and even the animated film *WALL-E,* must wonder if it is possible to lose control over superintelligent machines that might prioritize self-preservation over human safety.

Summary

The technical revolution has evolved through the personal computer to the internet, to cloud computing, and now artificial intelligence. AI is synonymous with the advancement of computer systems, aiming to mimic and replicate human cognitive capabilities. Applications are being partially driven by the development of pivotal technologies like natural language processing, computer vision, text-to-speech, robotics motion, machine and deep learning, and auditory systems. These advancements have led to convolutional neural networks and LLMs. A convolutional neural network is a deep learning network to extract features from images and other grid-like data. LLMs are trained on immense amounts of data to make them capable of understanding and generating natural language. These technologies make machine learning and AI possible. AI is purported to have the potential to be a more powerful positive force for society than any other technology in history. It also stands to reason that it could have the same potential to negatively impact society.

Keywords

artificial intelligence (AI) - The simulation of human intelligence processes by machines, especially computer systems. *p. 139*

AI hallucinations - Occur when an AI model produces incorrect or misleading information, often with confidence. *p. 142*

artificial superintelligence - A level of intelligence surpassing that of the brightest and most gifted human minds. *p. 143*

backpropagation - A machine learning algorithm that trains artificial neural networks by adjusting weights to reduce errors. *p. 144*

cloud computing - The delivery of computing services over the internet. *p. 140*

convolutional neural network (CNN) - A deep learning algorithm primarily used for analyzing visual data. *p. 143*

deep learning - A type of neural network with many layers that can learn increasingly abstract representations of data. *p. 148*

digital twin - A virtual representation of a physical object or process that can be used to simulate and monitor its behavior in real time. *p. 139*

discriminative DL - A model that separates the data points into different classes and learns the boundaries using probability estimates and maximum likelihood. *p. 149*

edge computing - Distributed computing that brings computation and data storage closer to the edge of the network. *p. 141*

general AI - AI model capable of understanding, learning, and applying intelligence across diverse tasks and domains. *p. 143*

generative artificial intelligence (GenAI) - A type of AI that creates new content, such as text, images, videos, and music, based on a prompt. *p. 149*

generative DL - A model that aims to learn the underlying patterns or distributions of data in order to generate new, similar data. *p. 149*

hybrid cloud - Integrates on-premises resources with cloud-based resources. *p. 141*

large language models (LLMs) - A type of AI that can understand, process, and generate human language. *p. 142*

machine learning (ML) - A subset of AI that enables machines to learn from data and improve from experience without being explicitly programmed. *p. 139*

narrow AI - AI model focused on specific tasks or problems. *p. 143*

neural networks - Computing systems inspired by the human brain's network of neurons, used to recognize patterns, and classify data. *p. 142*

on-premises resources - A computing environment where data and applications are stored on physical hardware situated within the organization's premises. *p. 140*

predictive ML - A statistical technique that uses machine learning algorithms to analyze historical data and predict future outcomes. *p. 146*

reinforcement learning (RL) - A type of ML where an agent learns to make decisions by receiving rewards or penalties. *p. 146*

supervised learning (SL) - An ML approach where the model is trained on labeled data to make predictions. *p. 146*

unsupervised learning (USL) - An ML approach where the model identifies patterns and relationships in unlabeled data. *p. 146*

Questions and Activities

8.1 Summarize the three economic revolutions.

8.2 Present the different phases of the tech revolution.

8.3 Explain the key technology developments preceding and contributing to AI.

8.4 Summarize the three main components of cloud computing.

8.5 Define function and capabilities of LLMs and convolutional neural networks.

8.6 Describe the relationship among large language models, deep learning, machine learning, and artificial intelligence.

8.7 Explain the similarities and differences between supervised learning, unsupervised learning, and reinforcement learning.

8.8 Describe the types of GenAI inputs and outputs and provide some applications for them.

8.9 What are the major concerns about AI and their ramifications?

Academic Analysis

Aman Saggu and Lennart Ante, "The Influence of ChatGPT on Artificial Intelligence Related Crypto Assets: Evidence from a Synthetic Control Analysis," *Finance Research Letters* 55, part B (2023).

The launch of OpenAI's large language model, ChatGPT, significantly heightened investor interest in artificial intelligence (AI) technologies, including AI-related crypto assets not directly linked to ChatGPT. By employing the synthetic difference-in-difference methodology, we observed notable "ChatGPT effects." AI-related crypto assets saw average returns ranging from 10.7 to 15.6 percent over the one-month period and from 35.5 to 41.3 percent over the two-month period following the ChatGPT launch. Additionally, Google search volumes—used as a proxy for AI attention—proved to be crucial pricing indicators for AI-related crypto assets post-launch. Our findings suggest that investors perceived AI-related cryptoassets as having increased potential or value after the launch, leading to higher market valuations.

Questions

1. What is the "ChatGPT effect"?
2. How did ChatGPT impact investor interest in AI and AI-related crypto assets?
3. Describe the AI-related crypto assets returns post-launch.

Yosef Bonaparte, *Innovation and Entrepreneurship from the Stone Tool to Artificial Intelligence* (Cambridge Scholars Publishing, 2024). ISBN: 1-5275-5965-3; ISBN13: 978-1-5275-5965-3

This book delves into the critical role of innovation and entrepreneurship for nations and individuals. Innovation transcends mere convenience; it reconfigures power dynamics, granting strategically innovative nations a significant edge.

The book is divided into three sections. Section 1 examines the roots of innovation and entrepreneurship, exploring the three major economic revolutions—agricultural, industrial, and technological—as well as current revolutions, the rise and fall of nations, and the historical contexts of various empires. Section 2 shifts focus to the latest frontier in innovation, artificial intelligence, analyzing its financial impact on stocks and exchange-traded funds. Section 3 addresses the realm of academia and the classroom, emphasizing the crucial role of educational innovation.

Questions

1. How does Darwin's theory of adaptability influence your perspective on financial valuation, particularly in the context of innovation and entrepreneurship?
2. What role do you see AI and blockchain technologies playing in the current economic revolution, and how do they compare to previous revolutions in terms of their impact on global financial systems?
3. Elaborate on how innovation has historically shaped the rise and fall of nations, and how this pattern might apply to contemporary global superpowers.

M. Dowling et al. "ChatGPT for (Finance) Research: The Bananarama Conjecture," *Finance Research Letters* (2023).

Based on ratings by finance journal reviewers of generated output, we demonstrate that the recently released AI chatbot ChatGPT can significantly assist with finance research. These results should, in principle, be generalizable across research domains. ChatGPT offers clear advantages in idea generation and data identification. However, the technology is less effective at literature synthesis and developing appropriate testing frameworks. Importantly, we show that the inclusion of private data and researcher domain expertise are key factors in determining the quality of the output. We conclude by discussing the implications, particularly the ethical considerations, arising from this new technology.

Questions

1. How can ChatGPT assist in the research process?
2. What are the limitations of ChatGPT's role in research synthesis?
3. Describe the ethical considerations illustrated in the paper.

Case Study

Algorithmic Bias in Mortgage Lending

AI/ML models, trained on imperfect historical data, can inadvertently perpetuate and amplify societal inequalities. The case of algorithmic bias in mortgage lending leads to "digital redlining" that impacts fundamental rights like access to homeownership and wealth creation. Let this serve as a powerful reminder that technological advancement, while promising efficiency, carries significant ethical responsibilities. The opacity of "black box" algorithms further complicates detection, accountability, and redress. This case underscores the critical need for a multidisciplinary approach to AI development and deployment. Solutions to algorithmic bias like this come not just as technical expertise, but also a deep understanding of ethics, law, sociology, and public policy. The challenge is to build AI systems that are not only accurate and efficient but also fair, transparent, and accountable, ensuring that the promise of AI serves to advance economic justice and social equity.

Connect provides

- The mechanics of algorithmic mortgage lending and bias
- Societal ramifications
- Legal framework and accountability
- Additional readings
- Questions

Even super-smart algorithms can fail history class, inadvertently redlining neighborhoods with lines of code instead of red ink.

Appendix AI Chrome Extensions

- **Glasp YouTube Summary:** Auto-generate summaries of your favorite YouTube video via AI. https://glasp.co/youtube-summary
- **Compose AI:** Cut your writing by 40 percent with the brilliant power of AI auto-completion and text generation. This extension will save you countless hours with emails and writing tasks. https://compose.ai
- **ChatGPT Prompt Genius:** Browse, Import, Store https://chrome.google.com/webstore/detail/chatgpt-prompt-genius
- **Wordtune:** Communicate your ideas better with clear and compelling writing. Change your writing tone, expand on your paragraphs, make your writing more concise. https://wordtune.com
- **Scribe:** Auto-generate step-by-step guides just by pressing record. Watches your workflow process and then uses the power of AI to create your guide on autopilot. Perfect for building SOPs. https://chrome.google.com/webstore/detail/scribe-ai-documentation-s
- **SciSpace Copilot:** Your all-knowing personal research assistant. SciSpace Copilot explains complex topics, formulas, and math in-tab. You can also ask follow-up questions to get clearer answers. https://chrome.google.com/webstore/detail/scispace-copilot
- **Perplexity:** Summarize any web page. Give answers to any question. Research co-pilot for web browsing. https://chrome.google.com/webstore/detail/perplexity-ask-ai