

IN THIS CHAPTER

- » Tracking AI from conception to fruition
- » Watching machines fool people and beat the experts
- » Seeing advanced AI capabilities in everyday life

Chapter 1

A Brief History of AI

To fully grasp the role of artificial intelligence (AI) in business, I begin by helping you trace its fascinating history. This background exploration not only illuminates AI's vast advancements, but also highlights its utility in business and marketing.

The earliest conceptions of artificial intelligence date back to Greek mythology, where Talos — an 8-foot-tall giant constructed of bronze — stood guard over the island of Crete to protect it from pirates and other invaders. Talos would throw boulders at ships and patrol the island each day. As the legend goes, Talos was eventually defeated when a plug near his foot was removed, allowing the *ichor* (blood of the gods) to flow out from the single vein in his body.

From that point forward, tales of automated entities flourished in mythology, captivating the minds of scientists, mathematicians, and inventors. Modern science and technology have realized some of these mythological concepts through recent advancements. In this chapter, I introduce you to those advancements, including the Turing test, machine learning, expert systems, and generative AI.

Early Technological Advances

Scientists trace the dawn of automation back to the 17th century and the invention of the *pascaline*, a mechanical calculator. Constructed by French inventor Blaise Pascal between 1642 and 1644, this groundbreaking device featured a *controlled*

carry mechanism that facilitated the arithmetic operations of addition and subtraction by effectively carrying the 1 to the next column. This calculator worked especially efficiently when dealing with large numbers. Following in Pascal's footsteps, Wilhelm Leibniz, a German mathematician, invented a calculator in 1694 that expanded upon the concept of the pascaline by enabling all four basic arithmetic operations — addition, subtraction, multiplication, and division (not just addition and subtraction). These devices first offered a glimpse into the potential for mechanical reasoning.

Fast-forward to the early 1800s, and you encounter the Jacquard system, developed by Joseph-Marie Jacquard of France, which used interchangeable punched cards to dictate the weaving of cloth and the design of intricate patterns. These punched cards laid the groundwork for future developments in computing. Near the mid-1800s, British inventor Charles Babbage unveiled the first computational device known as the *analytical engine*. Employing punch cards, this machine could perform a variety of calculations involving multiple variables, and it featured a reset function when it completed its task. Importantly, it also incorporated temporary data storage for more advanced computations — a feature crucial for any artificial intelligence (AI) system.

By the late 1880s, the development of the tabulating machine — designed by American inventor Herman Hollerith specifically to process data for the 1890 U.S. Census — helped the development of AI reach another milestone. This electro-mechanical device utilized punched cards to store and aggregate data, effectively enhancing the analytical engine's storage capabilities through the inclusion of an accumulator. Remarkably, modified iterations of the tabulating machine remained operational until as recently as the 1980s.

Alan Turing and Machine Intelligence

Many people regard Alan Turing, a British mathematician, logician, and computer scientist, as the founding father of theoretical computer science, and he paved the way for further AI breakthroughs. During World War II, he served at Bletchley Park, the United Kingdom's codebreaking establishment; and he played a pivotal role in decrypting messages encoded by the German *Enigma machine* (a code-generating device). Scholars and historians credit his work at Bletchley Park with both shortening the war and saving millions of lives.

Turing's key innovation at Bletchley was the development of the *Bombe*, a machine that significantly accelerated the code-breaking process used to decode messages from the Enigma machine. The Enigma used a series of rotating disks to

transform plain text messages into encrypted cipher text. The complexity of this encryption device and the coded messages it generated came in part from the fact that Enigma users changed the machine’s settings daily. The United Kingdom and all the Allies found cracking the code within the 24-hour window — before the settings were altered again — exceedingly difficult. The Bombe automated the process of identifying Enigma settings, sorting through various potential combinations far more rapidly than any human could. This automation enabled the British to regularly decode German communications.



REMEMBER

Although the details of this code-breaking device remained classified for many years, the Bombe stands as one of the earliest examples of technology outperforming humans in tasks that traditionally required human intelligence, executing them more efficiently and accurately.

The Turing Test in 1950

Soon after World War II, in a paper published in 1950 titled “Computing Machinery and Intelligence,” Turing introduced the idea of defining a standard by which we can call a machine intelligent. He designed the experiment (now called the *Turing test*) to answer the question, “Can machines think?” The fundamental premise of the experiment said that if a computer can participate in a dialogue with a human in such a way that an observer can’t tell which participant is human and which is computer, then you can consider that computer intelligent.

Turing’s test proposed that a human evaluator assess dialogues between a human and a machine that was designed to generate human-like responses. The evaluator knows that one of the participants is a machine, but not which one. To eliminate any bias from vocal cues, Turing proposed that the test giver limit the interactions to a text-only medium. If the evaluator found it challenging to distinguish between the machine and the human participant, the machine passed the test. The evaluation didn’t focus on the correctness of the machine’s answers, but on how indistinguishable its responses were from a human’s. In fact, the test’s criteria didn’t make any reference to the accuracy of the answers.

The Turing test: 1960s and beyond

In 1966, well after Alan Turing’s death, German-American scientist Joseph Weizenbaum created ELIZA, the first program that some say appeared to pass the Turing test. Many sources refute that it could pass the Turing test, but it was technically capable of making some humans believe that they were talking to human operators. The program worked by studying a user’s typed comments for keywords and then executing a rule that transformed the user’s comments,

resulting in the program returning a new sentence. In effect, the ELIZA, like many programs since then, mimicked an understanding of the world without actually possessing any real-world knowledge.

Taking this development a step further, in 1972, Kenneth Colby, an American psychiatrist, created PARRY, which he described as ELIZA with attitude. Experienced psychiatrists tested PARRY in the early 1970s by using a variation of the Turing test. They analyzed text from real patients and from computers running PARRY. The psychiatrists correctly identified the patients only 52 percent of the time, a statistic consistent with random guessing.



REMEMBER

Even to this day, the Turing test gives the world a concise, easily understandable method of assessing whether a piece of technology has intelligence or not. By limiting the test to text-based interactions that require natural language query (conversational English), anyone could easily understand the nature of the test when Turing first introduced it. And by separating out the accuracy of the response from the question of identification, it focused the test on evaluating what truly makes humans more human.



TIP

Computers have advanced by leaps and bounds since the time that Alan Turing first proposed the Turing test. But consider this timeline regarding the ongoing development of intelligent technology:

- » **As recently as 2021**, chatbots that much of the world had access to struggled to pass the Turing test consistently. Services such as Siri from Apple, Alexa from Amazon, and Google's Assistant could speak to us in natural language but would quickly get stumped with some of the most basic of questions. For example, the question "Describe yourself using only colors and shapes?" may prompt the answer "Okay, I found this on the web for describing colors and shapes. . . ."
- » **As of 2023**, major chat interfaces from the likes of OpenAI, Google, and others, can pass the Turing test. This quick change shows how technological advancements in the field of AI happen in fits and starts, with so much having changed dramatically in just 24 months.

The Dartmouth Conference of 1956

The academic community often considers the Dartmouth Conference of 1956 as the birth of artificial intelligence (AI) as a distinct field of research. Held during the summer of that year at Dartmouth College in Hanover, New Hampshire, the

conference brought together luminaries from various disciplines — computer science, cognitive psychology, mathematics, and engineering — under one roof for an extended period of six to eight weeks. Organized by computer scientists John McCarthy, Marvin Minsky, and Nathaniel Rochester, and mathematician Claude Shannon, the conference aimed to explore “every aspect of learning or any other feature of intelligence,” as stated in the original proposal for the conference.

The Dartmouth Conference of 1956 was groundbreaking for several reasons. It was more than just a summer gathering of intellectuals; it was a seminal event that shaped the trajectory of AI as we know it today. It provided the name, the initial community, the research directions, and the momentum that have fueled decades of innovation in AI.

Specifically, the conference

» **Coined the term *artificial intelligence* (AI):** The conference gave a name to a field that had been, up until that point, loosely defined and interdisciplinary across mathematics, computer science, engineering, and related fields. John McCarthy, one of the organizers, was credited with introducing the term, which helped in shaping the future direction of research by providing a focal point around which scholars could rally.

» **Served as a catalyst for future research:** It set the research agenda for decades to come. During the conference, participants engaged in deep discussions, brainstorming sessions, and even early-stage experiments on foundational topics in the AI field. The participants aimed to discover whether they could program machines to simulate aspects of human intelligence, with research topics such as

- Problem-solving
- Symbolic reasoning
- Neural networks
- Language understanding
- Learning machines

They designed programs to play chess, prove mathematical theorems, and generate rather simplistic sentences.

» **Provided a collaborative platform for interdisciplinary research:** Researchers who may not have otherwise crossed paths now engaged in meaningful dialogues, forging relationships that would lead to significant collaborations in the years and decades to come. This interdisciplinary nature was crucial for tackling the complex problem of simulating human

intelligence, which requires insights from various fields such as psychology, neuroscience, linguistics, operations research, economics, and more.

» **Attracted critical funding and attention to the developing field of AI:** The visibility and credibility gained from this event led to increased investment in AI research from both governmental and private sectors. This financial backing was essential for the development of labs, academic programs, and research projects that propelled the field forward.

Machine Learning and Expert Systems Emerge

Following the Dartmouth Conference (see the preceding section), two key sub-fields emerged that became the cornerstones of artificial intelligence — machine learning and expert systems. The *expert systems* were rule-based methods that drew upon predefined sets of instructions established by human beings. *Machine learning* (initially referred to as *self-teaching computers*) represented a radical shift in approach that aimed to build systems that learned from data, rather than by following scripted rules.

Meeting machine learning

Arthur Samuel, an American pioneer in the field of computer gaming and artificial intelligence, officially coined the term *machine learning* in 1959. Unlike traditional computing methods that relied on explicit instructions for every operation, machine learning focused on developing algorithms capable of producing results from existing data. These algorithms use statistical techniques to identify patterns, make decisions, or predict future outcomes based on those patterns.

In the 1960s, the Raytheon Company made a significant contribution to the field by developing an early learning machine system that could analyze various types of data, including sonar signals, electrocardiograms, and speech patterns. The machine used a form of *reinforcement learning*, a subset of machine learning in which the algorithm identifies optimal actions through trial and error. In essence, the system was rewarded for correct decisions and punished for incorrect ones. Humans operated and fine-tuned the system, and those humans pushed a goof button to flag and correct any errors. These corrections enabled the machine to adapt and improve its performance over time.

Critical standout features of machine learning include the following:

» **Adaptability:** Instead of relying on humans to manually code solutions to problems, machine learning enables computers to come up with their own solutions by examining large sets of data. This freedom has led to ground-breaking applications across various sectors. For example, machine learning algorithms power large language models and computer vision systems that enable computers to identify and understand objects and people in images and videos.

These systems can

- Generate human-like text.
- Recognize thousands of objects and filter spam e-mails with incredible accuracy.
- Transcribe and translate human speech in real time.

I discuss each of these topics in detail in subsequent chapters (Chapters 4 and 5, for example).

» **Efficient and scalable solutions:** Because developing specific algorithms for each recognition, filtering, or generating task would be both costly and time-consuming, machine learning offers a far more efficient and *scalable solution* (which means that the solution can perform tasks on huge data sets without having a corresponding increase in costs). The data-driven approach to finding solutions has revolutionized the way technologists approach and solve problems, and it has automated complex tasks (such as reviewing social media content for hate speech) that computer scientists once considered beyond the reach of computers.



REMEMBER

Because machine learning continues to evolve, experts expect its impact and relevance across various fields to continue to grow. See Chapter 2 for examples of the effects on areas of business.

Examining expert systems

In the late 1960s, many researchers focused on capturing *domain-specific knowledge*, which laid the foundation for *expert systems*, meaning technology systems or computers that played the role of experts in a specific domain such as drug discovery. Those expert systems were the precursors to modern-day AI systems that now exist. By the 1970s, researchers created some of the first expert systems, including DENDRAL (designed for chemical mass spectrometry) and MYCIN (aimed at diagnosing bacterial infections). These expert systems captured knowledge and reasoning capabilities from human experts to offer advice as diverse as simple medical diagnoses and exploration strategies for mineral mining.

The systems worked well in narrow subject domains, but the cost and difficulty of maintaining and scaling their rule-based knowledge effectively limited their usefulness. Research and development of expert systems went something like this:

- » **In the late 1970s**, a thawing of the AI Winter (see the following section) supported the broader adoption of expert systems in various industries, including healthcare, finance, and manufacturing. During this period, computer scientists developed specific tools to help expand their expert systems while those systems' usefulness grew exponentially.
- » **By the 1990s**, the limitations of expert systems became very evident, particularly their inability to learn from their processing experiences or strengthen their performance without external programming. This shortcoming led to a decline in the development of stand-alone expert systems, and computer scientists began to integrate them into larger, more complex computer systems.
- » **More recently**, ideas at the heart of expert systems have seen a resurgence of sorts, although they often appear in hybrid forms that incorporate machine learning (see the preceding section) and other data-driven techniques. Although not many corporations create and use stand-alone expert systems (after their limitations on explicit knowledge and brittleness became more apparent), the core concepts of capturing and applying human expertise in computational models remains integral to AI. And broader AI solutions incorporate expert systems as a complement to other advanced methods (such as machine learning and natural language processing, or NLP; see the section "More AI Developments in the 1980s" later in the chapter for more).



REMEMBER

The introduction of expert systems was an important moment in the history of artificial intelligence. Expert systems development pioneered knowledge engineering techniques that computer scientists still use to train AI systems today. But most AI tools now depend more on machine learning (which is much more scalable, or easily expanded), rather than explicitly programmed rules that require human involvement.

An AI Winter Sets In

After the hype of artificial intelligence in the 1960s and early 1970s, the limitations of early AI became clear, leading to a period of reduced funding and interest, which was coined the *AI winter*. The Lighthill report, compiled for the British Science Research council and originally published in 1973, helped bring about this AI winter. The report criticized the lack of practical applications and questioned the

potential of AI research. These criticisms led to reduced government funding in several countries, including the United Kingdom.

But even during this period of reduced funding, research continued that advanced core technical capabilities such as probabilistic reasoning, neural networks, and intelligent agents. Even in this period of reduced optimism, diligent computer scientists still drove key advancements before machine learning unlocked its next era of rapid progress in the 1980s.



TIP

The lessons of the AI winter of the 1970s have continued to inform the ethics debate around realistic versus overhyped claims in the AI world. This debate matters more than ever while differing opinions on the promise and perils of AI collide around the world.

The Stanford Cart: From the '60s to the '80s

You can't have a conversation about the history of artificial intelligence (AI) without discussing the story of the Stanford Cart, a remote controlled four-wheeled cart first developed in the 1960s that later came equipped with a camera and onboard computer for vision and control. This seminal project in the history of AI and robotics was one of the earliest attempts to create a self-driving vehicle. The cart, which was developed over a 20-year period, served as a platform for research into computer vision, path planning, and autonomous navigation.

The evolution of the Stanford Cart project not only mirrored the evolution of AI and robotics over its 20-year time span, but it also shaped the trajectory of AI and robotics, as well. The project remains a testament to the enduring impact of focused research and iterative development in the field of AI.

The stages of the Stanford Cart's evolution include

- » **Remote control:** In the 1960s, the first version of the cart simply allowed for remote control capabilities. Starting the cart's development this way made perfect sense because the cart served as a research platform for investigating the problem of controlling a Moon rover remotely from Earth.
- » **Self-navigation:** The early 1970s saw the cart get a camera and an onboard computer, which allowed it to navigate an obstacle course by taking photographs and then computing the best path forward based on those images. Later in the 1970s, more advanced computer vision algorithms allowed the cart to navigate complex environments more quickly while the image processing capabilities accelerated as well.

» **Real-time complex navigation:** By the 1980s, the cart could follow roads and avoid obstacles in real time, largely due to improvements in both hardware and software, especially in broad increases of computer processing power. This capability marked a significant milestone in the development of autonomous vehicles, which entered commercial production decades later. Increased processing power allowed for faster and more complex computations, while advanced algorithms enabled the cart to make split-second decisions.



REMEMBER

As one of the first practical applications of AI in robotics, the Stanford Cart demonstrated how computers could interact with the real world. The computer components that allowed visual input and analysis demonstrated the potential benefits of sophisticated image recognition and scene interpretation. And today's robotics and autonomous systems for path planning and obstacle avoidance use various algorithmic techniques that the Stanford Cart first introduced.

More AI Developments in the 1980s

Arguably, the 1980s stand as a critical decade in the development of artificial intelligence, characterized by groundbreaking advancements in various subfields, especially in machine learning, neural networks, and natural language processing. This period saw foundational advancements that set the stage for the AI technologies of today.

This decade's significant developments include

» **Backpropagation:** The introduction and popularization of the backpropagation algorithm for training neural networks. Before backpropagation, training complex neural networks took a lot of computational power and was less effective. The *backpropagation algorithm* streamlined the training process by efficiently calculating the error between predicted and actual outcomes, and then distributing this error back through the network to adjust the *internal weights* (which effectively transform the input data within the network's hidden layers). This innovation facilitated the training of multi-layer neural networks and paved the way for more complex architectures and applications.

» **Deep learning:** A subfield of machine learning that uses neural networks that have three or more layers. Researchers such as Geoffrey Hinton, Yann LeCun, and Yoshua Bengio (operating at various universities) were instrumental during this period because they laid the groundwork for this subfield. These layered neural networks found use in a range of applications, from image and voice recognition to natural language understanding, which would later fuel innovations in automating various business processes.

- » **Natural language processing (NLP):** Initially, programmers largely based NLP systems on handcrafted rules. However, the 1980s saw a significant shift toward statistical models, making these systems more robust and scalable. The decade set the stage for machine learning–based approaches that have come to dominate the NLP landscape, enabling more complex applications such as chatbots, translation services, and sentiment analysis tools.
- » **Robotics:** The decade also marked the beginning of significant advancements in robotics, much of which was built on the foundational concepts of AI. The Stanford Cart project, for example (see the preceding section), served as a crucial catalyst for research into autonomous systems.

Rapid Advancements of AI in the 1990s and Beyond

The remarkable journey of artificial intelligence (AI) goes from its mythological inspirations (Talos, the bronze giant in Greek mythology who protected Crete) to groundbreaking inventions such as Pascal’s calculator (discussed in the section “Early Technological Advances,” earlier in this chapter) and projects such as the Stanford Cart (see the section “The Stanford Cart: From the ’60s to the ’80s,” earlier in this chapter). The progress made since the early 2010s alone transformed the AI landscape and altered the way people think about technology’s role in various domains, including business and society at large.

Beginning in the 1990s, rapid advancements in existing branches of AI research brought expansion of capabilities to machine learning and deep learning. Other advancements in AI research brought new depth to the capability of AI to demonstrate seemingly intuitive thinking and to generate human–like original content.

Watching machine learning grow up

Between the 1990s and the early 2000s, machine learning emerged as a dominant force in AI development. (See the section “Meeting machine learning,” earlier in this chapter, for an introduction to machine learning.) This field of AI uses algorithms to analyze huge data sets to uncover patterns and make predictions without built-in, explicitly programmed rules. Spurred on by significant increases in computing power and data availability, machine learning delivered new use cases in the realm of *computer vision* (where computers derive information from images, videos, and other input) and *recommender systems* (information filtering systems that suggest items most pertinent to the user).

These AI advancements came about in part because the AI engines had access to large data sets. The models used to analyze these data sets mimicked more human-like pattern recognition and decision making by using statistical relationships between the data objects. These developments illustrated how quickly an AI system could learn (extrapolate) from data on its own, rather than having a programmer code specific and explicit instructions for that system. Machine learning is at the heart of AI to this day.

Playing a pivotal chess match

The 1990s saw a pivotal moment in the history of AI that captured the imagination of people around the world. IBM's Deep Blue, a chess-playing computer, defeated the reigning world chess champion, Garry Kasparov, in 1997. Even though Deep Blue didn't have the benefit of a modern neural network at the time and instead relied on brute-force heuristic search techniques and specialized chess algorithms, it did incorporate basic machine learning techniques to evaluate board positions and enhance its game play. Deep Blue's chess win was another momentous advance for AI and machine learning; it

- » Proved that a machine can outperform a human in a task that required complex decision-making over many steps.
- » Triggered huge debates about the future of AI and its potential impact on all facets of life. Those debates have only accelerated today with the much more recent introduction of generative artificial intelligence (see the section "Creating content with generative AI," later in the chapter).
- » Supported Kasparov's perspective that machines and humans working together can accomplish much more than either of them working alone. He introduced the term, *advanced chess* for a form of chess in which humans partner with computer systems to play chess, emphasizing that human intuition and machine calculations together were an almost unbeatable combination.



REMEMBER

Kasparov's idea of advanced chess had a lasting impact on how we think about AI today, and many AI researchers consider advanced chess a precursor to modern theories around AI serving as an assistant to a human operator in various domains. (Satya Nadella, Microsoft CEO, has referred to this assistance more popularly as AI co-piloting.) In subsequent chapters, I delve into the role of AI as a complementary tool for humans in the realm of business and marketing, and in those discussions, you can clearly trace the philosophical roots of this cooperative approach to Kasparov's insights.

Tracking the deep learning revolution

In recent years, the advent of deep learning has significantly elevated the capabilities and accuracy of AI systems. Building on the foundations laid by traditional machine learning, deep learning employs neural networks that have multiple layers — often referred to as *deep neural networks* — to achieve unprecedented levels of accuracy in tasks such as image classification, speech recognition, and natural language processing.

What sets deep learning apart from earlier AI technologies is the advancement in computational power, the availability of massive data sets, and the use of intricate algorithms that optimize neural networks with more than just a few layers. This multi-layered architecture enables the system to model complex relationships in the data, leading to remarkably precise results.

Deep learning-enabled systems

- » **Stand as the engine powering an extensive range of AI applications in use today.** Deep learning revolutionizes automation by enabling systems to perform complex analytical and predictive tasks entirely autonomously, without any human intervention. Whether you use digital voice assistants such as Siri or Alexa, voice-activated TV remotes, or advanced driver-assistance systems in modern automobiles, deep learning acts as the key technology underpinning many of these innovations.
- » **Promise to offer even more cross-domain intelligence in their next generation.** These future systems will likely require less data for effective learning, operate more efficiently on increasingly sophisticated processors, and employ even more advanced algorithms. People developing AI technologies want to bring artificial intelligence closer to mimicking the complexities and capabilities of the human brain.



REMEMBER

Although scientists and programmers may still be decades away from achieving *artificial general intelligence* — a state where AI possesses reasoning, learning, and common sense akin to human cognition — deep learning undeniably serves as a significant step toward that lofty goal.

Demonstrating intuition in the age of AI

The Turing test raised the seminal question, “Can machines think?” People began to ponder whether humans could distinguish between a machine and a human during a text-based interaction. (See the section “Alan Turing and Machine

Intelligence,” earlier in this chapter, for info about the Turing test.) This question appeared to find a definitive answer in the groundbreaking 2016 victory of AlphaGo over Lee Sedol in a game of Go.

AlphaGo was the brainchild of DeepMind, a British AI company that Google later bought. Unlike conventional AI programs, AlphaGo was purpose-built to master the game of Go, an ancient board game that boasts a complexity far surpassing that of chess. Although the game has simple rules, the sheer number of possible moves adds astronomical complexity. Top Go players — such as Lee Sedol, a leading figure in the world of Go — are revered for their intuition, creativity, and analytical skills.

In preparation for its 2016 face-off with Lee Sedol, AlphaGo underwent rigorous training, using a combination of machine learning methodologies, including deep learning, along with other algorithms such as the probability-based Monte Carlo tree search. The program analyzed thousands of historical Go matches and, perhaps more impressively, honed its skills by playing countless matches against itself. This self-play allowed AlphaGo to simulate various strategies and tactics, thereby enhancing its own game-playing capabilities.

When AlphaGo beat Lee Sedol in a five-game series, the global AI community sat up and took notice of two startling realizations:

- » **The unexpected display of AI ingenuity:** AlphaGo’s ability to make apparently creative and intuitive strategic choices — qualities that many assumed were the exclusive domain of human cognition. Sergey Brin of Google — whose company acquired DeepMind — was in Seoul for the third game and said, “When you watch really great Go players play, it is like a thing of beauty. So I am very excited that we have been able to instill that kind of beauty in our computers.”
- » **The profound capabilities and future potential of AI:** AlphaGo’s win provided more than just a technological milestone; it created a paradigm shift that raised the awareness of leaders across various sectors — from scientists and politicians to business leaders and the general public.



REMEMBER

This historical event where AlphaGo beat a consummate human Go player served as an irrefutable testament to the advancements in deep learning, indicating that AI can indeed perform tasks that many people previously thought only human intelligence could do.

Creating content with generative AI

Advancements in AI after 2010 saw dramatic innovation, particularly in the development of *generative models* (which can generate new synthetic data such as text or images). And by the 2020s, generative models found applications in a variety of fields ranging from art and entertainment to scientific research and drug discovery. Two specific developments provided the necessary foundation for generative models' advancement:

- » **Generative adversarial networks (GANs):** Introduced by computer research scientist Ian Goodfellow and his colleagues in 2014, GANs had the capability to generate incredibly realistic images, text, and other types of data. This significant leap forward provided a robust framework for generating intricate, high-quality digital assets. Subsequent advancements in GANs led to models such as StyleGAN, which can generate high-resolution, highly realistic images.
- » **The Transformer architecture:** Initially designed for natural language processing, it was adapted for generative tasks later in the same decade. The adaptation to generative tasks culminated in models such as OpenAI's GPT series that can generate human-like text.

I cover generative AI models extensively in Chapter 8.

