

1

A Brief History of Artificial Intelligence

COPYRIGHTED MATERIAL

“Artificial intelligence is growing up fast, as are robots whose facial expressions can elicit empathy and make your mirror neurons quiver.”

– Diane Ackerman

“The science of today is the technology of tomorrow”

– Edward Teller

IN 1997, IBM’s Deep Blue computer famously defeated world chess champion Garry Kasparov in a six-game match. This event marked a major milestone in the development of AI, as it demonstrated that a machine could outthink a human in a complex game with countless possible moves. The jubilation felt on achieving such a feat was mixed with hand-wringing that the age of machines was about to eclipse the age of humankind. Kasparov himself could not believe a machine could have defeated him and insisted this was a modern version of the Mechanical Turk, a 19th century con where a small person hid inside a supposed automaton and played chess.^{1,2} Despite these expressions of disbelief, the match captured the world’s attention. Chess was, after all, an ancient game highly revered as an expression of human mental ability. This event sparked a new interest in the abilities of machines that could think and adapt and even outshine humans.

Nearly seven decades since the prefix “artificial” was attached to intelligence, we live on the cusp of one of the largest disruptions in human society. When the CEO of Google, Sundar Pichai, calls AI one of humanity’s most profound inventions,³ and other tech luminaries such as Bill Gates argue, “The development of AI is as fundamental as the creation of the microprocessor, the personal

computer, the Internet, and the mobile phone,”⁴ and Elon Musk goes so far as to deem it potentially more dangerous than nuclear weapons,⁵ it is hard to dismiss the furor around this new technology as hyperbole. We may indeed be living in a time of profound change.

Artificial intelligence’s rise and awesome potential have been a topic of discussion among tech insiders for quite some time. Now, with the emergence of ChatGPT, a much greater slice of humanity is witnessing firsthand the impact of this technology in their daily lives. If there are skeptics questioning the impact and abilities of artificial intelligence, their doubts are certainly being challenged.

AI manifests in our lives in the form of self-driving cars, virtual assistants such as Alexa and Siri, and unprecedented information via search engines. It is even more prevalent behind the scenes, powering medical assistants, farming, and disaster response. AI developments are quickly transforming the way we work, communicate, and even think. The invention of the automobile changed landscapes and economies, while radio and telephone transformed communications and society. AI is poised to join these ranks of major disruptors in the coming years. We are witnessing the birth of a transformative force that will change how we make decisions and perceive the world around us.

However, the implications of this technological transformation are not without their challenges. There are concerns over privacy, security, and job displacement. Evidence shows that AI reflects some of society’s worst habits, such as racial and societal bias. As AI continues to become more sophisticated and more integral to our lives, individuals and society must carefully consider its

ethical implications. With the proper safeguards in place, the undeniable benefits of AI could usher in a new era of progress and prosperity for all.

How Innovators Throughout History Paved the Way for Modern AI: From Babbage to Turing

Artificial intelligence was long the province of fiction, fantasy, folklore, and myth. Inanimate objects developing human-like intelligence and abilities beyond our own are common in the stories we share. From figures such as mystical golems in Jewish tales and enigmatic homunculi of the Middle Ages to the evil computer HAL in *2001: A Space Odyssey* and the iconic droids in *Star Wars*, these legends reflect our curiosity and desire to create intelligence in our image.

Next, we trace the broad outlines of AI's emergence, from early conceptualizations of universal calculating machines to the first manifestations of what we today call AI.

Charles Babbage

The first practical steps toward AI happened in the last 200 years. Charles Babbage (Figure 1.1) emerged as a seminal figure in the history of AI, revered by many as the progenitor of this field. Babbage, a brilliant mathematician and inventor, possessed an indomitable spirit, a penchant for spectacle, and an insatiable curiosity that led him to his brilliant achievements in computing.^{6,7} His fascination with automatons mimicking human intelligence was sparked at age eight when his mother whisked



Figure 1.1 Drawing of Charles Babbage

Credit: The Illustrated London News / Wikimedia Commons / Public Domain.

him away to a museum of scientific artifacts and wonders. There, he saw an artful creation—a dancer cradling a bird—so exquisitely crafted that it appeared lifelike. From that moment forward, Babbage’s destiny was irrevocably entwined with the pursuit of crafting machines capable of emulating human behavior.

In his late 20s in the early 1800s, Babbage designed the first mechanical computer, the Difference Engine. This groundbreaking machine could perform complex mathematical calculations, such as producing tables of logarithms.^{8,9} Indulging his showman tendencies, Babbage delighted in donning extravagant attire as he showcased his creation to the venerable Royal Society in London and other esteemed venues across England.

Tales of his eccentricities, ranging from chasing musicians away from his abode when they impinged on his concentration to his fastidious craftsmanship, where gears and tools personally ground by him remained in use long after his death, embellished the legend of this extraordinary man.

The Difference Engine was never completed during Babbage's lifetime. It wasn't until the 1990s that it was finally built according to Babbage's design. It is on display at the London Science Museum, and a second one remains in the possession of a private donor who financed its creation.

Although Babbage was not able to see his design take life, it inspired his later, more audacious creation, the Analytical Engine. This was a much more ambitious endeavor, surpassing the Difference Engine in its versatility. Babbage intended it to be a general-purpose computing machine that could be instructed to perform any type of calculation. He envisioned tables of mathematical values being formulated, and these tables of values would inform calculations of things like dates of eclipses. Crucially, the Analytical Engine encompassed the fundamental duality of modern computers: the ability to store and process vast troves of data.

Regrettably, quarrels with his engineers and the drying up of funding meant that the Analytical Engine, like the earlier Difference Engine, was never built. It nevertheless stands as a major milestone in the history of computing. It was the first machine designed to be truly programmable. And it also helped to popularize the idea of artificial intelligence.

Ada Lovelace

Now recognized as the world's first computer programmer, Ada Lovelace (Figure 1.2) collaborated with Charles Babbage on his prototypes. When recounting the history of science and technology, the contributions of women have often been overlooked or underrepresented. But Ada Lovelace, daughter of the romantic poet Lord Byron and Anne Isabelle Milbanke, left her mark as indelibly as any male pioneer. Despite being born in the 19th century, when women's opportunities were limited, Ada Lovelace defied societal norms and fervently pursued her passion for mathematics and science. Her mother was responsible in large part for Ada's education. Seeking to shelter Ada from her father's perceived and infamous instabilities, she ensured Ada got a firm grounding in logic and mathematics.¹⁰



Figure 1.2 Ada Lovelace, watercolor painting, possibly by Alfred Edward Chalon in 1840

Credit: Science Museum Group / Wikimedia Commons / Public Domain.

When she was 17, Ada Lovelace met Charles Babbage at the house of Mary Sommerfield, a Scottish scientist and mathematician. Sommerfield had recognized a keen scientific intelligence in Lovelace and consciously brought about this intellectual match. Lovelace and Babbage became collaborators.

Her insight into Babbage's Analytical Engine went beyond his own ideas. She envisioned its potential beyond mere calculation. She recognized that the Analytical Engine could be used for more than just crunching numbers; it could be a tool for creativity and generating complex outputs. Her notes included an algorithm for calculating Bernoulli numbers, which is widely regarded as the world's first computer program. This visionary insight earned her the title of the world's first computer programmer.

Unfortunately, like many bright intelligences, she succumbed to her body's infirmities at age 36. But her legacy in computer science guides researchers and engineers to this day.¹¹

John von Neumann

John von Neumann (Figure 1.3) is another of the most prominent people to lay the foundations of computer science. Hailing from Budapest, Hungary, von Neumann was a child prodigy, a versatile intellectual who hungered for mathematics and physics. His unconventional, multi-disciplinary approach to studying made many skeptical of his seriousness, and, like his predecessor Charles Babbage, he gained a reputation as a maverick.¹²



Figure 1.3 John von Neumann

Credit: Los Alamos National Laboratory / Wikimedia Commons / Public Domain.

Von Neumann's extraordinary intellect carried him to doctorates in chemical engineering at the University of Zurich and mathematics from the University of Prague. When he submitted his doctoral dissertation to the faculty at the University of Zurich, the professors found it so profound and complex that they couldn't fully understand it. They asked him to simplify it, but with characteristic conviction, he firmly declined. To his mind, if they failed to comprehend the magnitude of his ideas, they lacked the qualification to pass judgment upon them. As a result, his dissertation remained unfinished and was never formally submitted, yet it still significantly impacted the field of mathematics and was later published as a monograph.¹³

Von Neumann moved on to the University of Berlin, where he continued to baffle his peers and students. Many stories of his time there illustrate his brilliance. Once, a student in a statistics lecture asked him a challenging question about a complex mathematical calculation. Without skipping a beat, von Neumann proceeded to solve the problem mentally and provided the answer within seconds. His lectures were often marked by brilliant expositions, which the students would then spend hours deciphering amongst themselves.

In the 1930s, he landed a teaching appointment at Princeton University. There, his genius would shine most brilliantly, and his pioneering contributions would forever transform the field of computing. Today, we take for granted the CPU as the brain of a computer and memory where computer programs are stored. Von Neumann was the genius who formulated these concepts and helped make them a reality, like UNIVAC, one of the first computers ever built.¹⁴

Alan Turing

For decades, the Turing test was held up as the holy grail of computing and artificial intelligence. It was an answer to the question of how we would know when machines had become intelligent. Mathematician Alan Turing (Figure 1.4) proposed his eponymous test, though he called it the Imitation Game.^{15,16} The test consists of questions posed to the machine and humans. If the answers are indistinguishable, one cannot tell which answers came from the machine, then the machine has won the game and passed the test.



Figure 1.4 Alan Turing

Credit: Dunk/Flickr/Public domain.

Until the early 2000s, beating the test seemed like a very difficult, nearly impossible task. This seemingly insurmountable challenge for artificial intelligence researchers imbued the Turing test with an aura of mystery and intrigue. It became a symbol of the quest for artificial intelligence.

The Turing test had profound philosophical implications. If a machine is equivalent to a human, then what does it say about human intelligence? What can it tell us about consciousness? It had practical implications as well, which we're now seeing firsthand. ChatGPT and DALL-E by OpenAI have taken the world by storm, and there's no doubt that ChatGPT can pass the Turing test.

The Turing test did, and still does, have its skeptics, who saw it as a limited indicator of machine intelligence.

They argued that relying on mimicking human speech patterns did not reflect on general intelligence. Now that we have reached an honest reckoning, it's unclear whether this holy grail is as significant as we thought it was. Chat-GPT is undoubtedly very human-like in its responses, but it is clearly still a non-conscious machine.

The eponym of this test, Alan Turing, was an Englishman who led the successful effort to break the German code during World War II, and then developed his theories of computing at the National Physical Laboratory. While Babbage's work was foundational for computing, and von Neumann influenced architected computer designs, Turing was a pioneer of theoretical computer science and artificial intelligence. His notion was of a Universal machine, known as a Turing machine, that could compute anything given a set of instructions. If this sounds like Babbage's Analytical Engine, it's because fundamentally they both had the same underlying idea of a flexible computing machine. Turing's mathematical concept, though, laid practical foundations for the development of computers.

The history of artificial intelligence is populated by thousands of mathematicians, engineers, psychologists, and scientists. However, among this vast sea of contributors, these four pioneers serve as human faces for the early development of artificial intelligence.

The Emergence of Modern AI

From the 1950s onward, the story of AI has taken on a certain canonical shape, which will be sketched here. Like its older, more venerable cousin of theoretical

physics, its coming to maturity is evolving into a story we tell ourselves and each other, a narrative that shapes our collective understanding. The story begins with Babbage and Lovelace and continues with Turing and von Neumann, and then comes one of the nodes, a turning point, in the 1950s.

The Dartmouth Conference: A Turning Point

In the summer of 1956, a group of researchers gathered on the campus of Dartmouth College to discuss a new field whose name had just been made up by one of the organizers. John McCarthy put “artificial intelligence” in the name of the conference and in the proposal for its funding.¹⁷ The Dartmouth Conference was a gathering of some of the leading researchers in computer science, mathematics, philosophy, and psychology, and would come to be seen as AI’s genesis moment.

The organizers were old friends. John McCarthy and Marvin Minsky had been roommates at Princeton University and had remained close friends ever since. Nathaniel Rochester and Claude Shannon were former colleagues from Bell Labs and collaborated on the development of computer languages and hardware. By most accounts, the gathering was somewhat chaotic, with a loose flow of ideas, brilliant minds each pursuing their own agendas, and people coming and going as they wished. Marvin Minsky brought his electric guitar and played late into the night, entertaining his colleagues with his musical skills.

It seemed nothing would come out of this gathering. Turing and von Neumann, who would have been

expected to be major figures at the conference, were dead (Turing) or ailing (von Neumann).^{18,19} But, in the years that followed the Dartmouth Conference, many of the participants went on to become leaders in the field of AI. John McCarthy, for instance, went on to develop the Lisp programming language, which became a vital tool in AI research. Marvin Minsky co-founded the MIT Artificial Intelligence Laboratory and became one of the most influential figures in the field. Nathaniel Rochester continued to work at IBM, overseeing the development of some of the earliest computer systems. Two of them, Nash and Simon, went on to win Nobel prizes for other endeavors.

Minsky continued the work he had begun during his doctoral research and shaped the research direction for the new field with his colleague Samuel Papert. Psychologists had been interested in how the brain worked and tried to model the behavior of individual cells. Frank Rosenblatt put together what would become the most famous neural network of all, the perceptron.²⁰ Incredibly simple compared to the monumental edifices that AI scientists now build, it was nevertheless an astounding demonstration of how cells could exhibit behavior. Minsky wrote his doctoral thesis on neural networks, and his book with Samuel Papert, *Perceptrons*, made the titular neural network famous.²¹ Rather than celebrating Rosenblatt's perceptron, the book argues that the network was too simple. Consisting of just a single layer of artificial neurons, they could be used to solve only very basic problems. They proposed a theoretical multi-perceptron, a neural network with multiple layers that could handle more sophisticated tasks. Ironically, the harsh critique

contained in the book about the perceptron, along with no practical way of implementing a multi-perceptron, was an early, unintended salvo that crashed the enthusiasm for AI.

From Optimism to Pessimism: The Story of the AI Winter

The 1960s were characterized by optimism and a focus on fundamental research. In comparison, the 1970s were a more challenging time for AI research, with a shift toward applied research and the development of expert systems. The enthusiasm of the 1950s and 1960s was exemplified by statements that promised human-level intelligence within a few years. *Life* magazine published this quote from an interview with Marvin Minsky, “[In] three to eight years, we will have a machine with the general intelligence of an average human being.”²² A few years later, in 1973, economist Herbert Simon, one of the creators of the world’s first artificial intelligence program, famously declared that “machines will be capable, within twenty years, of doing any work a man can do.”²³

These predictions turned out to be wildly optimistic. The power of computers at the time wasn’t enough to make their dreams a reality. Sure, computers had come a long way since the days of Turing machines that cracked the German codes in World War II, but the theory quickly outpaced the hardware. It was like trying to build a skyscraper with just hammers and nails.

The pendulum swung toward pessimism. In 1975, mathematician James Lighthill published a report for the British government that criticized the state of AI

research at the time, arguing that progress had been “grossly exaggerated” and that the field was unlikely to deliver significant results in the near future.²⁴ He argued that the combinatorial explosion of choices that most decision processes would face would never be overcome. A few years earlier, philosopher Hubert Dreyfus had argued in his book “What Computers Can’t Do”²⁵ that AI was fundamentally flawed because it was based on a flawed understanding of human intelligence. Dreyfus’s book became a bestseller and helped popularize the view that AI was over-hyped and unlikely to succeed in the near future.

And so, the curtain fell on the first act of the artificial intelligence saga. Responding to the mood of the times, the flow of research funds from the US government, primarily through the Defense Advanced Research Projects Agency (DARPA), dried up. The reduction in funding fueled the perception that AI was over-hyped, creating a negative feedback loop.

The Rise of Expert Systems

Artificial intelligence is not a monolithic concept, though, and can mean other types of structures. Neural networks, particularly deep neural networks, are the most successful version yet. The 1970s through the 1990s were dominated by so-called expert systems as interest in neural networks waned.

These early AI models were designed to emulate the decision-making abilities of a human expert in a specific domain, such as medical diagnosis or financial planning. Expert systems were essentially decision trees, a series of

if...then statements. They sought to capture and encode an expert's decision-making process so it could live as a program on a computer. Expert systems enjoyed considerable popularity until the dawn of the new millennium.

One example of an expert system is INTERNIST-I, a system to aid physicians in diagnosing medical problems. A team at the University of Pittsburgh developed a set of rules and heuristics from medical textbooks and consulted with medical experts. It started with formalizing the decision-making process of exactly one person, John D. Myers, M.D.²⁶ However, it never captured the confidence of the physicians it was supposed to help and largely remained a research tool. It was difficult to keep up with new knowledge, and coding it in. It also could not take a broad view of problems; the way it was structured could not consider a patient's history.

The finance industry today still uses expert systems. FICO, the credit analysis and fraud detection giant, uses an ingenious algorithm called RETE III, a rule-matching algorithm originally developed by computer scientist Charles Forgy.²⁷ Systems built on RETE III can efficiently sift through countless financial transactions, identifying patterns and trends that would be difficult for humans to detect, let alone at the algorithm's scale and speed.

AI Revival: A Fitful Resurgence

Artificial intelligence, as most commonly used today, is in the form of deep neural networks. *Neural* means a simulation of the function of the neural cells in our brains. From breathing to planning our next meal, training our bodies to lift heavy objects, and recognizing sounds, everything

we do happens because of the way neurons in our brains are structured and communicate with each other. The emergence of these experiences from relatively simple electrical connections between neurons is still a matter of intense research. Still, researchers have accepted that it works and hope to achieve something similar by simulating that structure and function.

They build networks of artificial neurons, trying to create a simplified model of what they see in the brain. Admittedly, we're like a child watching construction workers build a house, mimicking their actions by attaching two pieces of wood to a plank and proudly declaring it a home.

The term "deep" means there are multiple layers within the network. Think of a shallow pool with just a few layers of water compared to a deep ocean teeming with life at various depths. In AI, a shallow network may have few layers (maybe even just one or two) while a deep network may boast many more. The perceptron we mentioned earlier is a shallow neural network containing just one layer.

The concept of artificial intelligence, embodied by neural networks as proposed in Minsky and Papert's book *Perceptrons*, experienced a revival with the 1986 publication of a groundbreaking paper in the journal *Nature*. Titled "Learning representations by back-propagating errors,"²⁸ this seminal work was authored by three visionary research psychologists: David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams.

They were visionary in the sense that they worked in the field of neural networks when it was dismissed and not seen as a fit subject for proper scientists. In fact,

Hinton, as a graduate student in cognitive psychology, was advised to explore neural networks only in his spare time. He persisted, and after defending his doctoral thesis on vision, he continued to work on neural networks.

Since those early days, Hinton has become a towering figure in the world of artificial intelligence. He is associated with many of the major breakthroughs in the field, earning him the moniker “Godfather of deep learning.” Hinton wears that mantle well, even as he has sounded warnings about the pace and direction of recent developments in AI.^{29,30} He comes from a family with an illustrious lineage in science. His great-great-grandfather was George Boole, who invented Boolean logic, the mathematics underpinning computer circuit construction. He is also related to the 19th century explorer Sir George Everest.

Hinton credited the invention of backpropagation to David Rumelhart (also a psychologist) after the publication of the paper in *Nature*. Backpropagation is a way of teaching neural networks. It was simple enough to apply to large and complex neural networks, and it remains the foundation of neural network training today. To conceptualize how it works, consider an analogy of a group of musicians tuning their instruments in an orchestra. When a musician plays a note that is in tune, the conductor nods in approval. If the note is out of tune, the conductor provides guidance on whether the note is too high or too low, helping the musician make the necessary adjustments. This process continues until all the instruments are in harmony, much like how backpropagation refines the neural network’s performance.

Often the case in science, Hinton, Rumelhart, and Williams did not develop with these ideas in isolation. The ideas had been formulated by other researchers,²⁴ but their paper in *Nature* inspired and catalyzed the field. Backpropagation reignited widespread interest in neural networks because it was easy to understand, simple to implement, and computationally efficient. Despite these advantages, training large neural networks remained a prohibitive task in terms of time and money. This hints as to why neural networks still did not come to the forefront of AI. Once again, despite enormous advances in CPUs, the techniques proposed by the three mathematical psychologists were outpacing the day's technology.

Following this success, the nineties were marked by more remarkable developments in neural network research. New computational techniques were developed. One of the pioneers in this period was one of Hinton's students, Yann LeCun. Working against the tide in computer science research, where conventional wisdom still regarded artificial intelligence as a topic leading nowhere, he was responsible for some major breakthroughs. He built some of the first AI models called Convolutional Neural Networks, or CNNs, and taught them to recognize handwritten digits, a technological feat soon adopted by banks to read checks.³¹ Eventually, LeCun and his colleagues Corinna Cortes from Google and Christopher Burges from Microsoft created one of the most storied datasets in modern AI, the MNIST database of handwritten digits.³² It is almost a rite of passage for AI researchers learning the basics of artificial intelligence to creating a neural network model, replicating LeCun's work, that

can recognize handwritten digits with high accuracy. The sense of mastery from this task propels them on to greater discoveries.

No account of the history of artificial intelligence would be complete without acknowledging the third member of the triumvirate that has indelibly shaped the field. Along with Hinton and LeCun, Yoshua Bengio, a Canadian originally from France, born of Moroccan Jewish heritage, jump-started AI and continues to guide the field even today. In 2000, his groundbreaking paper³³ demonstrated how neural networks could understand language, even with all the multifarious ways of expressing meanings. Bengio has pushed the boundaries of what was possible in AI and is responsible for many of the breakthroughs that have become tools of the trade for AI researchers today. In 2018, in recognition of their groundbreaking contributions to artificial intelligence, Yoshua Bengio, Geoffrey Hinton, and Yann LeCun were jointly honored with what is informally known as the “Nobel Prize of Computing,” the prestigious Turing Award.




Figure 1.5 From left to right: Yann LeCun, Geoffrey Hinton, Yoshua Bengio. As imagined by Midjourney, prompted by the author.

The Birth of Modern AI

Then came a time when processors became fast enough, and memory became cheap enough that complex, deep learning neural networks could be trained. This was in the first and second decades of the new century. In 2012, the field experienced a sea change when Big Tech started making huge bets on AI. Once again, the Godfather of deep learning, George Hinton, was a central figure in this development.

Greater developments in neural network research marked the 1990s and the 2000s. New computational techniques were developed. The term big data became vogueish, followed soon by machine learning. Techniques were developed that could glean insights from reams of data using advanced statistical techniques. Some of these techniques were applied to the problem of computation-hungry neural networks.

George Hinton and two of his students at the University of Toronto, Alex Krizhevsky and Ilya Sutskever, built a model called AlexNet that could identify the content in images. It was a powerful demonstration of what computing power could finally achieve. It was also an anticipated progression. The trio were responding to an annual challenge called the ImageNet Large Scale Visual Recognition Challenge, organized by Professor Fei-Fei Li and her team at Stanford University.

Fei-Fei Li is another person who stands out as a mover in the field. Growing up in China during the Cultural Revolution, her father imbued in her a love for photography, and she has described her passion as a way to connect with the natural world and understand

its complexities. At Princeton, she understood that the field was ripe for neural network models, and large-scale datasets were needed to help researchers move forward. Therefore, she developed ImageNet, a database of digital images which now holds 14 million images across 1,000 categories. The profound influence of ImageNet is evident as legions of AI researchers have honed their skills, drawing invaluable insights from this repository. Like LeCun's MNIST database, ImageNet has become one of the most consequential objects in the field of AI.^{34,35}

The year 2012, the year AlexNet took the world by storm, was the first year neural network models were used to classify the images in ImageNet. Hinton and his students founded a company called DNNresearch. The large tech companies, sharklike instincts immediately recognizing the practical power of their algorithm, descended on Lake Tahoe in December 2012. There, in a casino hotel room, they engaged in a tense auction for the fledgling three-person company. This was a high-stakes secret auction, and the company representatives soon found themselves bidding outside their authorized limits. Phone calls and emails to headquarters to sign off on ever-increasing bids resulted throughout the day, and finally, one company walked off with the prize. The amount of the winning bid was never officially disclosed. Still, there were reports in the media that Hinton shut down the auction at \$44 million³⁶ despite every indication that the amount could go higher. Hinton later stated, tongue in cheek, "I signed contracts saying I would never reveal who we talked to. I signed one with

Microsoft and one with Baidu and one with Google.” Hinton and his students joined Google soon after.

AlexNet was one of many successful image recognition models. A couple of years earlier, DanNet, developed by Dan Ciseran, had won other competitions as well. It was a more specialized algorithm focusing on medical images. What AlexNet, DanNet, and other algorithms made clear is that the age of deep learning had arrived, and an arms race had begun among the top tech companies.

AI Today

This inter-company arms race would get much more heated 10 years later when the OpenAI consortium unleashed ChatGPT on the world. This technology, which seems to be paving the way for the incredible spoken computer interface in the TV series *Star Trek*, has delighted the public with its oracle-like abilities, similar to the delight that Google search occasioned in the naughts. But the amount of research it enables is balanced by the amount of misinformation it blithely generates.

AI’s ability to understand natural language had already up-ended the translation industry. In yet another science-fiction foreshadowing, the babel fish from *The Hitchhiker’s Guide to the Galaxy* no longer seems like a ridiculous, incredible device. In the book, the babel fish, when inserted into your ear, would immediately translate any language to English or your preferred language. For tourists around the world, Google’s Translate app performs that very function in a more rudimentary way.

For several years, research in speech and language understanding continued, but at a different rate than computer vision. GPT-3.5 was truly an eruption in this landscape, an achievement that caught even the AI cognoscenti by surprise. It is part of a new trend in AI research, the creation of foundational models that can be applied to many situations. Other foundational models are ones that generate images, such as DALL-E 2 by OpenAI, and Midjourney. Foundational models are being built for scientists, such as Segment Anything by Meta³⁷ that enable AI programmers to quickly find objects in any image.

What this means is that the ability to harness AI is being opened up to more people. These models are generally available to anyone right now. But no doubt they will be part of the AI-powered economy of the 21st century.

Driver of the 21st Century Economy

AI, specifically connectionist AI, has evolved from a nascent concept into a juggernaut within just 20 years. Trillion-dollar companies like Microsoft and Google have recognized the financial potential of AI and have pivoted their businesses to integrate AI into every aspect of their operations.

Microsoft's Azure, for instance, has grown its cloud computing platform into an AI-driven powerhouse, enabling businesses to harness the power of AI and machine learning for their applications. Google has also invested heavily in AI research and development, with projects like TensorFlow and Google Brain propelling the company to AI dominance.

Meta, formerly known as Facebook, also made a bold move by shifting its focus to an AI-powered metaverse. This ambitious endeavor has seen the company take on losses of more than \$10 billion in order to develop its new business model. The metaverse, a virtual reality space where users can interact with a computer-generated environment and other users, was expected to be a game-changer for businesses. But disappointing results have led them to change their business: a de-emphasis on the metaverse and doubling down on AI.

The widespread adoption of AI across industries has also profoundly impacted the job market. Job descriptions now commonly mandate a requisite understanding of AI. Employers appreciate the worth of employees who possess the ability to navigate new technologies, even if such employers may not comprehend the technology themselves. Entirely new job roles have arisen, such as AI ethicists, who are responsible for ensuring that AI systems are designed in a manner that respects human values and ethical principles. Universities and online platforms have reacted to the surge in demand for AI-related skills by offering specialized courses to equip professionals with the necessary expertise.

It is no secret that big tech companies have been on a talent acquisition spree, snapping up the brightest minds in AI research and development. Google, for instance, has made significant strides in consolidating its position in the AI space, reportedly absorbing 30 AI startups at a cost of around \$4 billion.³⁸ This aggressive expansion led to concerns that the available AI talent pool is being monopolized by the tech giants, potentially stifling innovation and leaving little room for newcomers to thrive.

There has been a surprising surge in global funding for AI startups. The exponential increase from \$4.5 billion in 2014 to a staggering \$38 billion³⁹ in just the first half of 2021 suggests continued international investor interest and capacity. Yet it also left many wondering if this is an investment bubble driven by the scarcity of AI knowledge as an essential commodity and fear of missing out on the next big breakthrough in AI or if it is a genuine testament to the financial potential of AI. However, the Great Layoff of 2023 released a considerable depth of AI talent into the market, which should fuel another surge of AI startups and funding.

The global nature of AI funding highlights the fact that AI innovation is not limited to Silicon Valley or other established tech hubs. Emerging markets in Asia, Europe, and beyond are increasingly becoming hotbeds for AI research and development. In some countries, such as China, AI is rapidly becoming part of social government, while in Europe, politicians are looking to curb and limit intrusions AI might make on privacy. The impacts of these different approaches will be explored in Chapter 6.

This amount of migration to AI, of both people and money, must be happening because investors hope to make large profits. Sizing market potential is a necessary step for businesses, but this task is far from a straightforward scientific endeavor. In fact, it is more of an art, a combination of intuition, experience, and data analysis. Companies such as McKinsey have made billions of dollars by mastering that art or at least convincing the world that they have.

In 2018, McKinsey⁴⁰ released a report on the value of AI, estimating that it could range between \$3.5 trillion and \$9.8 trillion by 2030. This staggering range is a testament to the complexity and uncertainty that characterizes market sizing. PricewaterhouseCoopers,⁴¹ one of McKinsey's competitors, put the figure at an even more astounding \$15.7 trillion.

These impressive numbers may turn out to be wildly inaccurate, but they strive to capture the potential of AI to revolutionize virtually every aspect of our lives. The technology has already shown remarkable promise in areas such as natural language processing, image recognition, and predictive analytics, to name just a few. And our successes in these areas have opened a window into what's possible in the future.

For instance, as we grapple with the repercussions of climate change, we will find AI invaluable in finding solutions for improved forecasting and disaster management. In the area of health, AI-powered medicine could significantly improve the accuracy and speed of diagnoses and prescribe personalized and more effective treatments for patients. In the finance field, AI could enable more sophisticated fraud detection and nuanced risk analysis. Drones, self-driving trucks, and cars could transform the way we move goods and people, improving safety and environmental health.

But the darker side of AI is waiting. If we flip the AI coin, we will see AI powering devices and machines that are not beneficial for society. We cannot wish away concerns by George Hinton that the technology he was such an instrumental part of may end up being detrimental

to humanity. AI can be used to spread misinformation, which can cost lives. It can amplify the worst of human society, like racism and sexism, wealth inequalities, and health inequity, because it learns from human society. What about the loss of human judgment and critical thinking? If AI is writing the news and creating music and art, what does it mean to be human? We will explore some of these issues in later chapters.

Scientists and politicians recognize these risks and are creating guidelines, standards, and laws that will aim to mitigate the negative consequences of this new technology. The more people are aware of AI's abilities and limitations, what it can and cannot do, the better our collective use of AI will be. AI literacy will be as important as traditional literacy and numeracy in making informed choices about how we will live with AI.

AI is not just another consumer product, a gadget on the shelf. Nor is it a mere extension of cloud services or another way to compute like Azure or Amazon Web Services (AWS). At its core, AI is a collection of algorithms and computational models designed to mimic the cognitive abilities of the human brain. By processing vast amounts of data, these algorithms can identify patterns and make correlations, thereby allowing machines to “learn” and adapt to new information. This ability to process and analyze data at lightning speed is what sets AI apart from traditional computing. The true essence of AI lies in its ability to enable humans to make efficient use of resources, a capacity that we commonly refer to as intelligence. In this sense, it is an intangible force that empowers us to coexist with the world and manage its resources in once unimaginable ways.

Final Thoughts

Historians caution that the tides of human progress are caused by factors much larger than individuals. There are broad societal and cultural forces that guide our energies and shape the direction of our progress and occasional regress. However, in narrating the past, certain personalities stand out as bright representatives of their times, shining light on their eras. We should also remember that countless other souls fueled their breakthroughs in thought and understanding.

The ebbs and flows of AI, its near emergence and temporary declines, yet a consistent upward trend that is clear in hindsight, mirrors that of other technologies. The parallel leads us to expect challenges we can barely imagine now. Let us hope that humanity chooses to use this technology for the good of our world.

References

1. Kasparov, G. K. & Greengard, M. *Deep Thinking: Where Machine Intelligence Ends and Human Creativity Begins*. (PublicAffairs, an imprint of Perseus Books, LLC, 2017).
2. Hsu, F. *Behind Deep Blue: Building the Computer that Defeated the World Chess Champion*. (Princeton University Press, 2004).
3. Google CEO: AI impact to be more profound than discovery of fire, electricity — CBS News. at <<https://www.cbsnews.com/video/google-ceo-ai-impact-to-be-more-profound-than-discovery-of-fire-electricity/>>
4. Gates, B. The Age of AI has begun. *gatesnotes.com* at <<https://www.gatesnotes.com/The-Age-of-AI-Has-Begun>>

5. Clifford, C. Elon Musk: “Mark my words — A.I. is far more dangerous than nukes.” CNBC (2018). at <<https://www.cnn.com/2018/03/13/elon-musk-at-sxsw-a-i-is-more-dangerous-than-nuclear-weapons.html>>
6. Swade, D. *The Cogwheel Brain: Charles Babbage and the Quest to Build the First Computer*. (Little, Brown, 2000).
7. Hyman, A. *Charles Babbage: Pioneer of the Computer*. (Princeton University Press, 1985).
8. Campbell-Kelly, M. Charles Babbage’s Table of Logarithms (1827). *Annals of the History of Computing* 10, 159–169 (1988).
9. Babbage, C. *Passages from the Life of a Philosopher*. (2018). at <<https://www.gutenberg.org/ebooks/57532>>
10. Lethbridge, L. *Ada Lovelace: Computer Wizard of Victorian England*. (Short Books, Limited, 2004).
11. Essinger, J. *Ada’s Algorithm. How Lord Byron’s Daughter Ada Lovelace Launched the Digital Age*. (Melville House Books, 2013).
12. Macrae, N. *John von Neumann: The Scientific Genius Who Pioneered the Modern Computer, Game Theory, Nuclear Deterrence, and Much More*. (Plunkett Lake Press, 2019).
13. Bhattacharya, A. *The Man from the Future: The Visionary Ideas of John von Neumann*. (W. W. Norton, 2022).
14. von Neumann, J. & Kurzweil, R. *The Computer and the Brain*. (Yale University Press, 2012).
15. Turing, A. M. *Computing Machinery and Intelligence*. (Springer, 2009).
16. Turing, A. & Copeland, B. J. On Computable Numbers, With An Application To The Entscheidungsproblem. in *The essential Turing: seminal writings in computing, logic, philosophy, artificial intelligence, and artificial life, plus the secrets of Enigma* (Clarendon Press; Oxford University Press, 2004).
17. McCarthy, J., Minsky, M. L., Rochester, N. & Shannon, C. E. A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence, August 31, 1955. *AI Magazine* 27, 12 (2006).

18. The Dartmouth Workshop—as planned and as it happened. at <<http://www-formal.stanford.edu/jmc/slides/dartmouth/dartmouth/node1.html>>
19. Dartmouth AI Archives. at <<https://raysolomonoff.com/dartmouth/>>
20. Rosenblatt, F. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review* **65**, 386–408 (1958).
21. Minsky, M. & Papert, S. A. *Perceptrons: An Introduction to Computational Geometry*. (The MIT Press, 2017). doi: 10.7551/mitpress/11301.001.0001
22. Darrach, B. Meet Shaky, the first electronic person. *Life* **69**, 58B – 68 (1970).
23. Simon, H. A. *The New Science of Management Decision*. (Prentice-Hall, 1977).
24. Russell, S. J., Norvig, P. & Davis, E. *Artificial Intelligence: A Modern Approach*. (Pearson India Education Services Pvt. Limited, 2015).
25. Dreyfus, H. L. *What Computers Can't Do: A Critique of Artificial Reason*. (Harper & Row, 1972).
26. Miller, R. A., Pople, H. E. & Myers, J. D. in *Computer-Assisted Medical Decision Making* (eds. Reggia, J. A. & Tuhim, S.) 139–158 (Springer, 1985). doi:10.1007/978-1-4612-5108-8_8
27. Forgy, C. L. Rete: A fast algorithm for the man pattern/many object pattern match problem. *Artificial Intelligence* **19**, 17–37 (1982).
28. Rumelhart, D. E., Hinton, G. E. & Williams, R. J. Learning representations by back-propagating errors. *Nature* **323**, 533–536 (1986).
29. Metz, C. “The Godfather of A.I.” Leaves Google and warns of danger ahead. *The New York Times* (2023). at <<https://www.nytimes.com/2023/05/01/technology/ai-google-chatbot-engineer-quits-hinton.html>>

30. Tavernise, S. The Godfather of A.I. has some regrets. At <<https://www.nytimes.com/2023/05/30/podcasts/the-daily/chatgpt-hinton-ai.html>>
31. LeCun, Y., Bottou, L. & Bengio, Y. Reading checks with multilayer graph transformer networks. In *1997 IEEE International Conference on Acoustics, Speech, and Signal Processing* **1**, 151–154 (IEEE, 1997).
32. LeCun, Y., Cortes, C. & Burges, C. *MNIST handwritten digit database*. (Florham Park, NJ, 2010).
33. Bengio, Y., Ducharme, R. & Vincent, P. A neural probabilistic language model. *Advances in Neural Information Processing Systems* **13** (2000).
34. Li, F. F. *The Worlds I See: Curiosity, Exploration, and Discovery at the Dawn of AI*. (Flatiron Books, 2023).
35. Kaplan, J. *The Genius of Women: From Overlooked to Changing the World*. (Penguin Publishing Group, 2021).
36. The Secret Auction That Set Off the Race for AI Supremacy | WIRED. at <<https://www.wired.com/story/secret-auction-race-ai-supremacy-google-microsoft-baidu/>>
37. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A. C., Lo, W.-Y., Dollár, P. & Girshick, R. Segment Anything. Preprint at <https://doi.org/10.48550/arXiv.2304.02643> (2023)
38. Hurst, A. Google revealed to have acquired the most AI startups since 2009. *Information Age* (2020). At <<https://www.information-age.com/google-revealed-acquired-most-ai-startups-since-2009-15415/>>
39. CB Insights. Total funding of AI startups worldwide 2014-2021. *Statista* at <<https://www.statista.com/statistics/621468/worldwide-artificial-intelligence-startup-company-funding-by-year/>>

40. Sizing the potential value of AI and advanced analytics | McKinsey. At <<https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier-applications-and-value-of-deep-learning>>
41. PricewaterhouseCoopers. PwC's artificial intelligence services. *PwC* (2023). at <<https://www.pwc.com/us/en/services/consulting/cloud-digital/data-analytics/artificial-intelligence.html>>

