CHAPTER

1

The Evolution of AI

Artificial intelligence (AI) is a complex field that has evolved continuously over several decades. Its origins trace back to the period just before World War II, when the discipline we now call computer science had yet to take shape.

The journey from those early days to today's AI has been full of twists and turns. At times, external events accelerated its development; at others, progress stalled in long periods of relative stagnation.

But in recent years, a seismic shift has taken place.

Driven by rapid advances in computation, storage, digital data, and global connectivity, the field has undergone a dramatic and sudden expansion. New applications and possibilities are emerging at an astonishing pace. It's now impossible to ignore the conclusion that AI will play a defining role in the future of human development.

Yet, amidst all this promise—and potential peril—it all traces back some 90 years to a conceptual and infinite roll of virtual tape.

The Birth of Electronic Computation

In 1936, screen legend Charlie Chaplin captivated cinema audiences with his latest creation, *Modern Times*. In this film, he reprised his iconic character, the *Little Tramp*, this time confronting the relentless march of industrialisation.

Chaplin's character personified the ordinary person struggling to cope with a rapidly changing world. Over the course of the film, he unwittingly becomes a symbol of resistance against the dehumanising forces of modernisation.

The theme of humans reduced to cogs in the machinery of automation had already been explored a decade earlier in Fritz Lang's cinematic masterpiece, *Metropolis*. There, humans were consumed by machines—pitiless instruments of social inequality.

While mechanical automation had been a source of concern since the Industrial Revolution, audiences in 1936 would not have been thinking about electronic automation. Concepts such as digital systems, programming, and computation had not yet entered mainstream thought—and were far from the minds of ordinary people.

Fortunately for us, in that same year, an extraordinary person was hard at work.

Alan Turing and Computability

In 1936, within the hallowed halls of King's College, Cambridge, a young mathematician was quietly reshaping the course of history. Alan Turing, though unaware of the full extent of his influence at the time, was on the cusp of introducing ideas that would help give birth to our modern world.

Turing's journey began with a profound mathematical question posed by German mathematician David Hilbert: Could one devise an algorithm to determine the truth or falsehood of any mathematical proposition? This question, known as the *Entscheidungsproblem*, intrigued the mathematical community. It was part of a broader goal by Hilbert to establish a solid foundation for all of mathematics using formal logic. But his ambitions would face a formidable challenge from a young, audacious maverick—Kurt Gödel.

Gödel, an unconventional thinker, developed an innovative technique known as Gödel numbering, which he used to prove a groundbreaking concept we now call Gödel's Incompleteness Theorem. In this theorem, he showed that in any formal mathematical system, there exist true statements that cannot be proved from *within* that system. Though initially met with scepticism, his discovery would later shake the very foundations of mathematics.

Inspired by the lingering mystery of the *Entscheidungsproblem* and the challenge posed by Gödel's discoveries, Turing began to explore the nature of computation itself. In the spring of that year, he submitted his findings to the London Mathematical Society in a paper titled *On Computable Numbers, with an Application to the Entscheidungsproblem*.

Within this landmark paper, Turing introduced the concept of a universal machine—later commonly referred to as a Turing machine. This was an abstract computing device capable of executing any mathematical

computation expressible as an algorithm. It was one of the first formal descriptions of a theoretical model of computation.

His Turing machine comprised an infinite roll of tape marked with symbols, a read-write head, and a set of rules for symbol manipulation. Turing demonstrated that this deceptively simple yet powerful machine could, in theory, carry out any computation described algorithmically. He also made a startling discovery: Some problems can never be solved by any machine, no matter how powerful. This result—known as the halting problem—defined fundamental limits of algorithmic computation.

Turing's work was instrumental in both introducing and formalising the idea of a universal computing machine—one that could simulate the behaviour of any other computing machine, given the right programme. His work forms the bedrock of the computational theories underpinning today's AI.

Turing would famously go on to play a key role in the team that broke the Nazis' Enigma code and helped to save countless lives. He will return later in our story—for, not content with introducing computation as we know it, he would later light the touchpaper for machine intelligence.

Claude Shannon and Information Theory

As Alan Turing was making waves with his seminal paper, another luminary in the world of computing was quietly beginning his graduate studies in electrical engineering at the Massachusetts Institute of Technology.

Claude Shannon was already excelling as both a mathematician and an electrical engineer. At MIT, he studied mathematics and electrical engineering while working under Vannevar Bush on the Differential Analyser—an early analogue computer. He also explored the logic of George Boole, whose work would later influence many aspects of computing.

During the war, Shannon joined Bell Labs to work—like Turing—on cryptography. While Turing was breaking codes, Shannon was laying the mathematical foundations of modern cryptography, helping prove the security of one-time pad encryption and defining principles for secure communication.

In 1948, he published a groundbreaking article in two parts in the *Bell System Technical Journal*, titled "A Mathematical Theory of Communication." This seminal work laid the foundation for what we now know as *information theory*.

At its core, this theory seeks to quantify and understand the transmission of information across communication channels. Shannon's genius lay in his ability to abstract complex problems into elegant mathematical

constructs. He formalised the concept of a *bit* as the fundamental unit of information—a language for measuring data that we still use today.

Shannon also introduced a concept he called *entropy* to measure uncertainty in information. Although inspired by thermodynamics, Shannon's entropy was a novel mathematical framework for quantifying how much information is contained in a message. This insight became pivotal to the design of error-correcting codes—an essential element of modern digital communication.

Despite the abstract nature of his work, Shannon's ideas permeate our daily lives. They underpin the efficient transmission of data over the internet, the design of reliable digital storage, and the compression algorithms that allow us to stream high-definition video with ease.

Claude Shannon is often hailed as the "father of information theory." His work revolutionised digital communication and laid the groundwork for many of the technologies that power the digital age—including AI.

John von Neumann and Stored Programmes

In the immediate aftermath of World War II, the landscape of computing was dominated by machines with fixed, hard-wired programmes. These early computers were built for specific purposes. Changing their function often required rewiring, structural modifications, or even complete redesigns. As demand grew for machines that could tackle a wider range of tasks, the limitations of this rigid architecture became increasingly clear.

Among the brilliant minds involved in the Manhattan Project—the effort to build the first atomic bomb—was a Hungarian-American mathematician, physicist, and polymath named John von Neumann. Faced with the enormous computational demands of the project, von Neumann quickly recognised the impracticality of *reprogramming* machines by hand. This challenge pushed him towards a groundbreaking innovation.

The idea of a stored-programme computer was already taking shape among pioneers like J. Presper Eckert and John Mauchly. But it was von Neumann who, in 1945, formally articulated the concept in a way that would shape the future of computing. His work brought together emerging ideas about memory, logic, and control into a single, coherent framework that would define the architecture of modern computers.

The key insight was that programme instructions could be stored in memory alongside data, allowing computers to be reprogrammed without any physical modification. This became the foundation of what is now known as the *von Neumann architecture*.

At the heart of this model was the idea of *memory* capable of holding both instructions and data. This shift introduced programmatic flexibility—enabling a single machine to perform many tasks simply by changing the software it ran.

Today, the distinction between programme files and data files is so embedded in everyday computing that it's hard to imagine it had a starting point. Yet it did. Von Neumann's architecture laid the foundation for modern computing—making it possible to build the complex software systems that underpin everything from operating systems to artificial intelligence.

A New Science

Emerging from the nightmare of global war, early computer pioneers were driven by the belief that they could create tools essential for rebuilding a shattered world in a fresh and modern way.

The trio of Turing, Shannon, and von Neumann stand as towering figures in this transformative era—but they were far from alone. The landscape of early computing was teeming with visionaries, each contributing to the nascent field in their own distinctive ways—too many to name here. These pioneers worked across university mathematics departments, military institutes like Bletchley Park and ARPA, and renowned research centres and organisations such as Bell Labs, IBM, and the Rand Corporation.

Their collective efforts laid the groundwork for what would later become known as *computer science*—a discipline that blends mathematics, engineering, and logic.

In these academic and industrial crucibles, ideas were exchanged, theories forged, and innovations born. From the fundamental concepts of algorithms and computation to the practical engineering of computing machines, this dynamic community laid the foundation for the digital revolution that was yet to come.

The Dartmouth Conference

In the summer of 1956, Soviet leader Nikita Khrushchev broke with tradition and began openly criticising the regime of his predecessor, Joseph Stalin. A US federal court ruled that racial segregation on Montgomery buses was unconstitutional. A young, upcoming singer named Elvis Presley had just released a new song called *Heartbreak Hotel*. At the cinema, audiences queued to watch the science fiction film *Forbidden Planet*, with its thrilling themes of intelligent robots and out-of-control technology.

A Seminal Conference

Also that same month, at Dartmouth College in New Hampshire, a young and eager computer scientist was preparing rooms for the arrival of colleagues from around the world. They were gathering to discuss a niche and abstract topic: *machine intelligence*.

The college students had left for the summer break, and John McCarthy had seized the opportunity to organise what would later be recognised as a landmark event in the history of technology.

But that wasn't clear to McCarthy in those summer months in New England. At just 28, he had already built a strong reputation as a mathematician, having earned a PhD from Princeton. He had recently become one of the youngest professors on campus and was pressing ahead with cutting-edge research in the emerging field of computer science.

The Infancy of Computing

Computing was still in its infancy and, by modern standards, only a handful of computers existed at all. Those that did were used by research institutions, government agencies, universities, and large corporations. Their use was limited by extremely high costs, vast physical size, and complex, specialised operation.

If you had a spare million dollars in 1956, you could have purchased an IBM 704 with a few tens of kilobytes of memory—capable of processing 40,000 instructions per second. That's many billions of times less capable than a one-dollar computer chip available today.

AI Visionaries

So one might ask—given the state of the art at the time—why was McCarthy organising a conference on such an advanced topic? The answer is that he, like the other attendees, was a visionary.

In those early days of computing, the field generated enormous excitement among researchers. The focus quickly honed in on human-level capabilities—both physical (robotics) and cognitive (machine intelligence). At the time, these areas were seen as intertwined, not separate disciplines. Today, robotics and artificial intelligence are distinct fields, but in the 1950s, they were all part of a thrilling, unified vision of the future.

Earlier that decade, Alan Turing had written a paper titled *Computing Machinery and Intelligence*, where he set the stage for the idea of "thinking machines." Turing explored the possibility of machines exhibiting intelligent

behaviour indistinguishable from that of a human. His paper sparked intense debate within the nascent computer science community and lit the fuse on a topic that still resonates today.

That fuse-wire burnt its way right through that sleepy college in New England in the summer of '56.

The Gathering

Turing had already—sadly and tragically—died before the conference began. But many of those who attended would, like Turing, go on to become the rock stars of computer science. In 1956, however, they were more scruffy academics than revered experts.

It's almost tempting not to name individuals, as so many great minds took part. But let's at least acknowledge the principal organisers of the event. Alongside McCarthy, there were: Marvin Minsky (who later founded the MIT AI Lab), Nathaniel Rochester (a key figure in early IBM computer development), Claude Shannon (the father of digital circuit design theory and information theory), and Oliver Selfridge (a pioneer in perception, pattern recognition, and machine learning).

The conference brought together experts from diverse fields, fostering discussions that laid the groundwork for artificial intelligence and helped shape both its early development and long-term direction. It was not a short affair, stretching over several weeks. Roaming the empty maths building, the participants would meet in ad hoc groups to probe various aspects of the problem domain.

Legacy

In hindsight, a great deal of mythology has developed around the conference—much of it deserved, but some rooted in hero worship.

The Dartmouth Conference is often miscredited as the birthplace of AI. In reality, its participants were building on ideas that had been circulating for over a decade. Some attendees were active in the cybernetics movement, which had already laid important groundwork for machine intelligence research.

Two key technical themes emerged from the proceedings: a focus on symbolic methodologies and a contrast between deductive and inductive systems (both of which we explore in later chapters). But perhaps the most enduring outcome of the workshop was the popularisation of the term *artificial intelligence*. This simple, memorable, and powerful phrase has echoed through the decades to the present day.

In the conference proceedings, we find a striking sentence:

> Every aspect of learning, or any other feature of intelligence, can in principle be so precisely described that a machine can be made to simulate it.

Right from the outset, computer scientists planted a marker pointing directly towards a future of intelligent machines. Was this prophetic insight—or hubristic boldness? One thing is clear: This was the beginning of a long tradition of overly optimistic and embellished claims—a pattern that has shadowed the field for decades and still persists today.

The Timeline of AI

The path of artificial intelligence to date has been winding and erratic—marked by moments of great promise, frustrating dead ends, and long periods where progress seemed to stall indefinitely.

Yet with each setback, the AI community persisted, building on the thoughts and achievements of those who came before them. It was an era of great ambiguity, rudimentary technology, and uncertain outcomes. Countless individuals toiled away, driven by a shared vision of machines that could one day think, act, and behave like us.

Their collective efforts—often against the odds—have brought us to where we stand today, reaping the fruits of their labour. Let's now explore the milestones, breakthroughs, and relentless spirit that have shaped the remarkable evolution of artificial intelligence.

It's worth noting that, while we present a timeline of AI here, *there is no single, universally agreed way to capture its genesis and history*.

Early Foundations (1940–1956)

The roots of artificial intelligence can be traced back to the 1940s—a time when the world was engulfed in the terrible events of World War II. Amid the chaos, a group of brilliant minds quietly laid the groundwork for a technological revolution that would one day shape the future of humanity.

At the heart of this foundational period was Alan Turing, whose work in the 1930s and 1940s was instrumental in defining the theoretical underpinnings of computation. His groundbreaking concept of the *Turing machine*—a theoretical model of computation—laid the foundation for modern computing.

At the same time, other trailblazers were making significant contributions in areas such as stored programmes, information theory, and programming languages.

Early artificial intelligence was closely intertwined with *cybernetics*—a field that gained prominence during this period. Coined by Norbert Wiener in 1948, cybernetics explored the relationships between control, communication, and intelligence—both in machines and living organisms.

Many early pioneers of AI were influenced by the emerging ideas of cybernetics. Alan Turing interacted with figures in the British cybernetics community, and Oliver Selfridge—a key figure at the 1956 Dartmouth Summer Research Project on Artificial Intelligence—made notable early contributions to pattern recognition that resonated strongly with cybernetic thinking.

Cybernetics also captured the public imagination. During this period, it featured in newspaper articles, radio programmes, and even popular films. The idea that machines could mimic intelligence and behaviour was already being explored—both in technical research and in cultural discourse.

This was a period marked by optimism, as pioneers set the stage for AI's emergence as a distinct field of study. It was the dawn of a new era—when the quest to replicate human intelligence and create thinking machines truly began.

The Birth of AI (1956–1969)

The 1950s ushered in a remarkable era of exploration as AI began to take shape. Emerging from a period of war and conflict, a mood of optimism fuelled intense research, leading to groundbreaking discoveries.

During this period, the dominant approach was *symbolic AI*—later referred to as "good old-fashioned AI" (GOFAI). Researchers sought to build intelligent systems using symbolic representations and rule-based reasoning. Early achievements included basic decision systems, game-playing programmes, and rudimentary natural language processing—all pioneering efforts to emulate human cognitive abilities.

However, not all researchers agreed on the best path forward. While symbolic AI dominated the era, alternative approaches—such as *connectionism*, which later inspired neural networks, and probabilistic methods rooted in statistical reasoning—were the subject of early research and theoretical debate.

The birth of AI was marked by great enthusiasm and a shared belief that intelligent machines were not only possible but also within reach. Researchers tackled challenges such as problem-solving, reasoning, and knowledge representation, convinced that these systems could soon match human intelligence.

Although early AI systems showed promise, they also revealed fundamental limitations—foreshadowing the challenges that would define AI's

development for decades to come. Nevertheless, the stage was set for AI to evolve, advance, and shape the future of human–machine interactions.

The First AI Winter (1970–1980)

The early promise and optimism of AI gave way to a challenging period known as the *first AI winter*, during the 1970s and 1980s. This phase was marked by scepticism, reduced funding, and a decline in AI research activity.

As the 1960s transitioned into the 1970s, the initial excitement surrounding AI began to fade. A major reason for this shift was the growing gap between lofty expectations and the actual progress being made. Early AI systems struggled with complex, real-world problems, exposing the limitations of existing approaches.

These difficulties were compounded by structural constraints—limited computational power, scarce data, and overhyped claims about AI's capabilities. As disillusionment spread, many governments and organisations began to question AI's practical feasibility, and research funding had begun to dry up.

Yet while mainstream enthusiasm faded, important progress continued behind the scenes. Basic rule-based tools for decision-making were beginning to emerge in specific industries. Meanwhile, alternative approaches based on networks of simple, interconnected units (the ancestors of modern neural networks) were quietly researched by a small but dedicated group of researchers.

Although AI had entered a winter, these quiet developments ensured the field would not remain dormant forever.

AI Research Resurgence (1980–1987)

Following the challenges of the first AI winter, the field experienced a revival in the 1980s and early 1990s. This period was marked by renewed enthusiasm, technological progress, and a broader approach to AI research.

A significant development was the growing success of earlier forms of AI such as expert systems, which aimed to encapsulate domain expertise within a computational structure. These began gaining traction in industries like healthcare, finance, and engineering. Although these systems used narrow reasoning based on lists of rules and domain-specific knowledge, they nonetheless delivered some practical value and demonstrated a tangible return on investment.

Another key factor driving this resurgence was a gradual shift in research focus. Symbolic AI, which had dominated earlier decades, began to face

growing scrutiny as researchers recognised its limitations. This encouraged exploration of alternative paradigms and more practical approaches suited to real-world constraints.

Further momentum came from the growing availability of computational power and digital data. Advances in hardware, algorithms, and data storage enabled researchers to tackle more complex problems and process larger datasets than ever before.

AI research also became more interdisciplinary, drawing on fields such as cognitive psychology, neuroscience, and linguistics. Researchers aimed to develop systems capable of mimicking human cognitive functions—including natural language understanding, reasoning, and learning.

This resurgence laid the groundwork for AI's diversification into multiple subfields, including machine learning, computer vision, natural language processing, and robotics. It was a pivotal era that restored momentum to AI research and set the stage for the breakthroughs that would define the decades to come.

The Second AI Winter (1987–1997)

The 1990s marked a period of renewed challenges and scepticism, often referred to as the *second AI winter*. Following the resurgence of AI in the 1980s, enthusiasm began to cool and funding declined. Once again, AI was perceived as having overpromised and underdelivered.

Several factors contributed to this downturn. A key issue was the widening gap between AI's theoretical potential and the practical limitations of existing technology. While researchers had made meaningful progress, AI systems still fell short of the ambitious expectations set by funding bodies and industry leaders.

The limitations of these systems in handling uncertainty, context, and real-world complexity became increasingly apparent. Early AI struggled with tasks requiring common-sense reasoning or the interpretation of vague or ambiguous information. These shortcomings reinforced doubts about AI's viability beyond highly specialised applications.

Despite the setbacks, important work continued. Researchers refined existing methods and advanced newer approaches that had begun gaining momentum in the late 1980s, such as probabilistic reasoning and reinforcement learning. These developments strengthened the foundations of modern machine learning and helped prepare the ground for later breakthroughs in neural networks.

In hindsight, the second AI winter served as a critical learning period for the field. It underscored the importance of setting realistic expectations, fostering interdisciplinary collaboration, and committing to sustained

innovation—lessons that would ultimately drive AI's transformation in the years to come.

Rise of Machine Learning (1998–2015)

The mid-2000s marked a transformative era in AI research, driven by the resurgence of *machine learning*. This period saw major advancements in algorithms, computing power, and data availability—leading to breakthroughs that would redefine artificial intelligence.

A key factor behind this resurgence was the renewed focus on machine learning, a subfield of AI concerned with building models that learn from data. While machine learning had been part of AI since its earliest days, it gained new prominence as researchers moved beyond handcrafted, rule-based systems and embraced more flexible, data-driven approaches.

This era also saw the expansion of neural networks. Loosely inspired by the structure of the human brain, neural networks began to demonstrate value in tasks such as image recognition, speech processing, and natural language understanding. Their success paved the way for AI systems that could process and interpret information in increasingly human-like ways.

It is also important to clarify the relationship between neural networks and machine learning, as these terms are often confused within the media, but in reality they are closely intertwined. Within the chapters of this book, we treat them distinctly to help non-technical readers understand the concepts more clearly.

However, in practice, the two areas have become deeply integrated. Neural networks are not an alternative to machine learning, but a powerful family within it—one that has become the dominant class of machine-learning techniques. Modern neural nets use machine-learning methods to adjust their internal parameters automatically as they learn from data, enabling them to recognise patterns, make predictions, and generate content with remarkable accuracy.

Progress in machine learning was enabled by two critical developments. First, a substantial increase in computational power—driven by advances in hardware—allowed researchers to use more advanced algorithms, train more complex models, and process larger datasets.

Second, the rise of the internet, social media, and sensor-rich environments produced vast amounts of digital data. This gave researchers access to an unprecedented volume of information with which to train AI systems, enabling models to recognise patterns and make accurate predictions.

The rise of machine learning reinvigorated the AI community and sparked a wave of innovation and investment. Startups and tech giants alike began applying machine learning to a wide range of real-world problems.

During this period, the terms *machine learning* and *artificial intelligence* began to be used interchangeably—though not quite accurately. It was by the end of this era that AI-powered applications started to become a part of daily life, marking the beginning of AI's widespread societal impact.

Renaissance and Commercialisation (2015–Present)

This was an era of unprecedented growth and innovation, marking the renaissance and commercialisation of AI—a shift that continues to shape today's technological landscape. AI has moved from research labs into everyday life, becoming an integral part of both the digital and physical worlds.

This transformation has been driven by the convergence of several key factors. First, advances in machine learning and deep learning have enabled AI systems to achieve remarkable results in pattern recognition, natural language understanding, and decision-making. Deep learning, in particular, has propelled neural networks to the forefront of modern AI, making them the backbone of today's most powerful applications.

Another critical driver has been the availability of massive datasets and powerful computing infrastructure. The rise of the internet, mobile technologies, and data-driven industries has provided researchers with vast amounts of training data. Meanwhile, cloud platforms and specialised AI hardware (such as GPUs and TPUs) have made high-performance computing more accessible, allowing organisations to scale AI initiatives rapidly.

The late 2010s also saw significant investment in AI research and development from tech giants, startups, and governments alike. These investments pushed the boundaries in areas such as computer vision, reinforcement learning, and generative AI.

AI applications have become pervasive across industries—transforming healthcare, finance, manufacturing, and entertainment. AI-powered systems now drive personalised content recommendations, enhance medical diagnostics, optimise supply chains, and support the development of autonomous vehicles.

As AI's influence on society deepens, ethical considerations have taken centre stage. Issues such as bias, transparency, and accountability have led to the creation of ethical AI frameworks, guidelines, and legislation aimed at ensuring these systems are used fairly and responsibly.

The commercialisation of AI has become a powerful force behind its expansion. Businesses increasingly recognise AI's potential to improve operations, boost productivity, and unlock new revenue streams. Today, the AI industry is a thriving global ecosystem of startups, technology providers, and service companies supporting organisations in every sector.

Generative AI and Machine Creativity (2022–Present)

AI has entered a transformative phase, marked by the rise of *generative AI* and machine creativity—pushing the boundaries of what technology can achieve. This period represents a remarkable fusion of art, technology, and innovation, where AI systems are beginning to display creative capabilities once considered uniquely human.

At the heart of this evolution are generative AI models, which have shown astonishing abilities to produce content across a range of domains—including images, music, prose, and interactive media. *Machine creativity* is now a prominent theme, as AI systems are no longer just automating tasks but generating original artistic and intellectual work.

Advances in *natural language processing* have also transformed how machines understand and generate human language. Large-scale language models can now engage in complex conversations, generate human-like text, and utilise context with unprecedented depth. This has fuelled the rise of *conversational AI*, with chatbots, virtual assistants, and AI-powered customer service agents becoming increasingly human-like in their interactions.

Alongside the excitement surrounding generative AI, ethical considerations have taken centre stage. Questions of authorship, originality, bias, misinformation, and the societal implications of AI-generated content are now widely debated. The balance of power between human and machine creativity—once a science fiction scenario—has become a real-world policy issue, shaping conversations on regulation, intellectual property, and the future role of AI in the creative industries.

Summary

What lies ahead for AI is uncertain, but recurring patterns and promising directions are emerging—we will examine these more closely in the closing Afterword chapter.

The journey AI has taken so far stands as a testament to human ingenuity and perseverance. Throughout its history—briefly chronicled in this chapter—AI's evolution has been a tapestry woven with moments of triumph and uncertainty. From the post-war era to the present day, its trajectory has been marked by cycles of hope, disillusionment, and resurgence.

At the heart of this narrative lies a legacy of individuals who embodied innovation, resilience, and an unwavering belief in AI's potential. Armed with rudimentary tools and constrained by the limits of their time, they achieved remarkable feats in environments that would seem inconceivably primitive to today's AI developers.

Through scarce hardware, data limitations, algorithmic uncertainty, unclear use cases, and funding gaps, these pioneers persisted. Their unshakable faith in machine intelligence propelled them forward, laying the foundations for today's breakthroughs.

Now, we stand at the dawn of a new AI era—one in which AI has broken through into many facets of society. It is a time when its capabilities increasingly rival those of its human creators, reshaping how we live, work, and interact with the world.

As you engage with AI—whether seeking travel recommendations, refining an email, enhancing an image, or exploring creative tools—take a moment to reflect on the visionary minds who made this technology possible.

Let us thank them for leaving an indelible mark on our modern world—a legacy that will continue to shape our future for decades to come.