1

Understanding Artificial Intelligence

In the early hours of 8 October 2024, a man lay sleeping in a cheap California motel room. He had taken the next day off from his job as an AI researcher and university professor to undergo an MRI at a nearby hospital.

At about 2:00 a.m., his phone rang, abruptly pulling him from a deep sleep. Half-conscious, he instinctively reached for his phone and answered in a groggy voice. On the other end of the line, a calm but firm voice spoke after confirming his identity: "Dr. Hinton, we are calling to inform you that you have been awarded the 2024 Nobel Prize in Physics."

For a brief moment, he wasn't sure if he was dreaming.

Perhaps it was some cruel prank or a misdialed number. Nobel Prizes aren't exactly announced in the middle of the night, and, at that moment, he certainly didn't feel like someone who had just joined the ranks of history's most celebrated scientists. After all, his research was in AI, not in physics. Still half asleep, he squinted at his phone's screen and noticed the caller ID indicated a Swedish number. He thought, "Wait a minute. . . . The Nobel Prize committee is based in Sweden. Could this actually be real?" The unmistakable Swedish accent of the caller was a detail that lent credibility to the surreal moment, but he hesitated further. He knew well that AI-generated voices could mimic foreign accents, as AI was precisely the kind of technology he had spent his entire career researching and developing. But as he listened

more closely, he realized multiple people—all with the same distinct Scandinavian intonation—were on the line. That's when it hit him: this was real. He—Geoffrey Hinton, professor at the University of Toronto, often hailed as the godfather of AI—had just won the 2024 Nobel Prize in Physics, which was jointly granted to John J. Hopfield of Princeton University. The award came as a surprise given that AI is not traditionally linked to physics. But as the reality sank in, it became clear that this was a historic recognition of Hinton's and Hopfield's pioneering work from the 1980s on neural networks, which are the foundation of modern AI and draw inspiration from the laws of physics. Once dismissed as abstract theory, their work now underpins some of the most powerful AI systems of the twenty-first century, driving everything from Large Language Models (LLMs)—AI systems trained on vast amounts of text to understand and generate humanlike language—to groundbreaking scientific discoveries.

In fact, neural networks led to yet another significant AI milestone at the 2024 Nobel Prizes—in Chemistry. The award recognized a trio of scientists who harnessed machine learning to predict protein structures, a breakthrough with far-reaching implications for medicine and biotechnology. Among the recipients were Demis Hassabis, co-founder of DeepMind (acquired by Google in 2014), and his colleague John Jumper. Their groundbreaking AI model, AlphaFold, solved a decades-old biological mystery, accurately predicting the intricate 3D structures of proteins from their amino acid sequences—a challenge that had long defied conventional scientific methods.

The excitement surrounding the awarding of Nobel Prizes has always been about more than just accolades; the prizes can be seen as a reflection of the collective *zeitgeist*, capturing the defining conversations of each era, and in 2024, that conversation was unmistakably about AI.

But how did we get here? After all, for decades, Artificial Intelligence was more of a theoretical dream than a reality. In the 1940s, Alan Turing laid the groundwork, envisioning machines that could simulate human reasoning, and the field officially took shape in 1956, when John McCarthy introduced the term *Artificial Intelligence* at the Dartmouth Conference, sparking waves of optimism. But progress

wasn't linear. AI went through cycles of hype and disillusionment—so-called "AI winters"—when early systems failed to meet expectations, leading to funding cuts and skepticism.

Today, however, as evidenced by the recipients of the 2024 Nobel Prizes in both Physics and Chemistry, optimism has replaced skepticism. AI has reached an unprecedented level of maturity, moving from theoretical promise to real-world impact.

What Exactly Is AI?

Although *AI* is a term that today dominates headlines, boardroom discussions, and even casual conversations, its true meaning remains widely misunderstood. At its core, *Artificial Intelligence* refers to the ability of computational systems to perform tasks traditionally associated with human intelligence—such as learning, reasoning, problem-solving, perception, decision-making, and even recognizing emotions.

AI functions by processing vast datasets through sophisticated algorithms, allowing it to identify patterns, draw insights, and improve their performance over time. But unlike traditional software, which follows rigid, pre-programmed instructions, AI systems can adapt dynamically, refining their outputs based on new information and past experiences. This ability to learn and evolve is what makes AI not just a tool, but a transformative force shaping industries, innovation, and the future of work.

Data, Hardware, and Algorithms

A convergence of key factors has enabled AI to evolve and integrate seamlessly into both daily life and the workplace. At the heart of this progress lies a powerful triad: hardware, data, and algorithms.

- *Hardware* provides the physical infrastructure, or computing power, for AI. Specialized processors like GPUs (Graphics Processing Units) and TPUs (Tensor Processing Units)—designed specifically to handle large-scale mathematical operations—have dramatically accelerated model training by enabling

massive parallel computations. The hardware allows today's complex models to be trained in days instead of years.

- *Data* is the fuel AI depends on to learn. Thanks to the Internet of Things (IoT)—a vast network of connected devices like sensors, wearables, and even home assistants—we now generate unprecedented amounts of data every second. Cloud computing, which refers to the delivery of computing services (like storage, processing, and software) over the internet, helps store and scale this data efficiently across distributed systems.
- *Algorithms* are the engines designed to recognize patterns, make predictions, and generate outputs from data. Among the most important are the neural networks—algorithms inspired by the structure of the human brain, composed of layers of interconnected nodes (or "neurons") that process information and learn from data through training. Neural networks form the foundation of deep learning, which is a subset of machine learning—the branch of AI focused on building systems that improve their performance over time by learning from data, without being explicitly programmed.

It's the convergence and maturation of these three elements—hardware, data, and algorithms—that has made AI not only viable, but widely adopted across industries, transforming how we work, make decisions, and live our daily lives.

Narrow AI and General AI

When thinking about the role of AI in the world today, there are two types of AI that we should consider: Narrow AI and General AI.

The AI tools we interact with daily—especially in the workplace—fall under the category of *Narrow AI*. These systems are designed to perform specific tasks with remarkable efficiency, whether it is recommending products, optimizing logistics, detecting fraud, or automating customer service. However, despite their advanced capabilities, they operate within strictly defined parameters and lack the broad adaptability, reasoning, and self-awareness of human intelligence. They can analyze, predict, and generate content, but they do so within the

confines of its training data. You can think of Narrow AI as a hyperspecialized expert.

But Narrow AI is just the beginning. Progress in AI development is now bringing us closer to the next frontier: *General AI*, or *Artificial General Intelligence* (AGI). Unlike Narrow AI, which is limited to specific tasks, AGI would possess the ability to reason, learn, and adapt across a broad range of fields and activities, much like a human. It would also think abstractly and make decisions with a level of flexibility and independence that no AI system has yet achieved.

Think of AGI as a true generalist thinker. Its arrival would mark a profound shift—moving AI from a set of powerful but specialized tools to machines capable of true cognitive flexibility, problem-solving across disciplines, and even self-directed learning. While AGI remains a theoretical goal, the rapid pace of AI development suggests that its arrival is a question of *when*, rather than *if*.

AI's Impact in the Real World

Throughout history, technology has continuously reshaped the way we worked: the steam engine mechanized labor, electricity powered factories, and the internet digitized communication and commerce. Like AI, these innovations are classified as *General Purpose Technologies* (GPTs)—foundational advancements that drive widespread economic and societal transformation and that share three key characteristics: (1) they have broad applications across all industries, (2) they continuously evolve and improve over time, and (3) they spur the development of other innovations.

These developments aren't just incremental—they are revolutionary. For the first time, we have a technology that doesn't just automate tasks, but amplifies human potential, unlocking productivity, creativity, and insights at a scale never before imagined. From streamlining workflows to generating new ideas, AI is quickly becoming a co-pilot for professionals across nearly every industry: it enables smarter decisions, faster execution, and more personalized experiences. In that sense, AI isn't just reshaping work—it's reimagining what humans are capable of achieving.

What was once confined to research labs and to the minds of academics such as Geoffrey Hinton and John Hopfield is now embedded in the very fabric of how businesses operate, make decisions, and compete in an increasingly AI-driven landscape. Consequently, AI is enabling companies of all sizes—whether long-established industry leaders or agile startups—to reinvent their business models, optimize processes, and enhance their tools with a clear end goal: to deliver greater value to customers and drive accelerated growth.

Several forward-thinking companies are already leveraging AI to stay ahead of the curve:

- Agricultural machinery leader John Deere is integrating AI-powered automation and computer vision into its farming equipment, enabling precision agriculture and maximizing crop yields.
- Consumer goods giant Unilever is harnessing AI-driven analytics to optimize its supply chain, predict consumer trends, and develop hyper-personalized marketing strategies.
- The world's most valuable carmaker Tesla continues to push the boundaries of AI with its self-driving technology, using machine learning to improve vehicle autonomy and manufacturing efficiency.
- Engineering powerhouse Siemens is embedding AI into its "smart factories," enhancing predictive maintenance, increasing operational efficiency, and reducing energy consumption.
- Music industry disruptor Spotify is revolutionizing music discovery by employing AI-driven recommendation engines to curate personalized playlists and enhance user engagement.

The Risks of AI

AI is no silver bullet, and it does not come without risks. While some tech enthusiasts are eager to embrace AI's potential, most people remain deeply concerned. Worries range from ethical dilemmas and AI turning against humanity to job displacement and the risks of AI hallucinations (when AI generates information that sounds plausible but is entirely false or fabricated)—and not without reason.

For decades, Hollywood has reflected these fears, often portraying AI as a malevolent force. Think of the Replicants in *Blade Runner*, humanoid robots that rebel against their creators, or Skynet in *The Terminator*, an autonomous superintelligence that triggers a nuclear apocalypse to eradicate humanity. These anxieties have even extended beyond fiction, prompting a group of AI experts and tech executives to issue an open letter calling for a pause in the development of AI systems more powerful than Open AI's GPT-4, one of the most advanced language models available, known for its ability to generate human-like text, write code, and perform complex reasoning tasks. Their plea, however, went unanswered.

But when we zoom into AI's specific impact on the human skill set, three key risks emerge: (1) the risk of substitution, (2) the risk of dependence, and (3) the risk of commoditization.

Risk of Substitution

The first risk is substitution—the fear that AI will replace so many of our skills that humans themselves become redundant. A report by investment bank Goldman Sachs[1] warned that AI could displace the equivalent of 300 million full-time jobs and automate a quarter of all work tasks in the United States and Europe over the next decade. These numbers scream "substitution," fueling concerns that AI's rapid advancements may one day render many human skills obsolete.

While many experts are rightly concerned about substitution, a smaller group views this shift from a different perspective: as an opportunity for new skill development. This perspective is best understood through the example of calculators.

When calculators were first introduced in the 1980s, they sparked widespread fear. After all, they allowed anyone—regardless of deep mathematical knowledge—to achieve the same results as someone with a PhD in arithmetic, simply by knowing how to use the tool effectively. This dramatically equalized outcomes, or at the very least, democratized access to solving complex mathematical equations.

Suddenly, performing advanced calculations was no longer limited to a select group of experts who had devoted their lives to studying mathematics; it became a skill accessible to most people.

With AI, the implications are even more profound. For the first time in history, we have a technology that mimics the functioning of the human brain—but on steroids. Unlike previous innovations, AI is not just automating tasks; it is expanding access to a vast range of skills that were once uniquely human—such as learning, decision-making, writing, and creativity. By making complex tasks more accessible and processing vast amounts of information effortlessly, AI is fundamentally reshaping what it means to acquire and apply knowledge.

Fast forward to today's world: we now know that calculators have not eliminated the need to study mathematics—but they have reshaped which skills matter most for mathematicians. The focus has shifted from simply solving problems to understanding and deconstructing them, selecting the right data inputs, and effectively leveraging the tools that calculators provide.

Risk of Dependence

Democratization comes with a trade-off: as AI expands access to skills, it also raises the risk of overdependence—where reliance on technology weakens our ability to think critically and function independently. When technology takes over certain tasks, the very skills it replaces can begin to atrophy. People who become too dependent on technological tools—both in life and at work—may lose the opportunity to develop a deep understanding of the principles behind them, and as a result, their ability to function independently diminishes when access to technology is limited.

Think of what happens all too often at the end of a meal with a large group of friends. When the bill arrives, instead of quickly doing the math in our heads, we instinctively reach for our phones to calculate the split. The convenience of technology has, in many ways, reduced the need for certain types of mental agility, making us worse at performing even simple calculations on our own—something we realize only when our phone battery dies or there's no internet connection.

A recent study by Macnamara and colleagues[2] warns that frequent use of AI assistance can lead to unintended skill degradation—with users often unaware that their abilities are declining. The study highlights how AI-driven decision-making tools, particularly in high-stakes

fields like medicine and aviation, can lead professionals to unknowingly lose their expertise over time.

This phenomenon aligns with what AI expert and author Pascal Bornet calls "AI obesity"[3]—a term describing the excessive reliance on AI at the expense of human capabilities. This overdependence is particularly risky because AI is far from perfect: with its biases, ethical concerns, and occasional hallucinations, blindly trusting AI can degrade the quality of our work rather than enhance it.

Much like relying on a calculator without understanding its logic, overreliance on AI can leave us unable to recognize even obvious mistakes. Imagine sending out a critical weekly report to your boss without noticing a glaring error—such as an extra zero in an Excel formula—simply because you didn't double-check the AI-generated output; or not verifying the source of an AI-generated meeting note-taker and unintentionally spreading biased or incorrect information. Without the ability to critically assess AI's results, we risk amplifying its flaws rather than leveraging its strengths.

Research from Carnegie Mellon and Microsoft,[4] conducted with 319 knowledge workers, reinforces this concern. The study found that the more people use AI, the less critical thinking they engage in. A key irony of automation is that by mechanizing routine tasks and leaving exception-handling to the human user, it deprives them of regular opportunities to practice judgment and strengthen cognitive skills, and over time, this leads to mental atrophy, leaving users unprepared to handle complex situations when they inevitably arise.

In our survey, 44% of respondents expressed concern that using AI tools could make them overly dependent on them. This fear echoes the "calculator effect" described above, where reliance on a tool—like calculators in education—raises questions about whether we are losing core skills in the process and highlights the delicate balance between leveraging technology for efficiency and preserving essential human capabilities.

Risk of Commoditization

The challenges associated with the use of AI go beyond the risks of substitution and dependence. Another major concern is

commoditization. When AI grants widespread access to the same knowledge, problem-solving abilities, tools, and even ideas, what will set one professional apart from their colleagues? Likewise, if businesses all have access to similar data, algorithms, and AI systems, what will differentiate one company from its competition?

These questions illustrate why standing out becomes increasingly difficult in the age of AI. When everyone has equal access to the same AI tools, traditional skills alone are no longer enough to create a competitive edge. If nearly everyone is adopting AI, simply using these tools without evolving one's skillset adds little value. True differentiation comes not from adoption alone, but from developing new capabilities that complement and enhance AI's potential.

Human Intelligence Versus AI Intelligence

The 2024 Nobel Prizes raise a profound question: Who was the true winner eventually—Human Intelligence or Artificial Intelligence itself? Nobel Prizes are traditionally awarded for human achievement, but in 2024, for the first time, they recognized work that might not have been possible without AI. Were these awards a recognition of human ingenuity in advancing AI, or of AI overcoming human ingenuity?

The answer is complex.

Let's start by saying that yes, the Nobel Prizes were awarded to Human Intelligence—to the ingenuity behind developing AI and applying it to some of the most complex scientific challenges. The distinction, subtle yet crucial, is that the recognition still belongs to the human minds that created and harnessed AI, not to the technology itself, and that it remains a celebration of human brilliance, innovation, and discovery.

But now let's consider the second part of the question. If today, Demis Hassabis and his team at DeepMind were awarded a Nobel Prize in Chemistry for an AI model that accurately predicts protein structures—outperforming human researchers in a task that had stumped scientists for decades—what does that tell us about the role of AI in the future? If AI can solve scientific challenges more efficiently than even the most brilliant human minds, at what point does it stop

being a tool and start becoming a true independent contributor to discovery and innovation? Eventually, if AI can improve itself by writing new code, will it be able to generate new AI systems better than human beings do?

These questions are no longer theoretical. As AI takes on an increasingly central role in our world, far beyond the realm of science, the line between human achievement and machine contribution grows ever more blurred. With each new AI breakthrough, our very definitions of intelligence, creativity, and innovation are being reshaped—challenging long-held assumptions about what it means to be good at what we do.

An Exercise in Ancient Greek Translation

To illustrate the path forward—one that allows us to harness AI's potential without being replaced by AI—I want to take a step back in time. Specifically, to my teenage years in Italy, when I was studying Ancient Greek in school.

Picture this: On a chilly winter morning in Italy—the country that gave birth to the Renaissance and proudly upholds a deep-rooted tradition in the humanities—I stepped into a classroom carrying a heavy Ancient Greek dictionary. The room was orderly, with rows of neatly arranged desks facing a blackboard, ready for a group of 14-year-olds embarking on their first year at *Liceo Classico*.

Here, we studied philosophy, literature, Latin, Ancient Greek, and art history, among other disciplines, in preparation for college. It felt like a world apart—yet it was just the year 2000. At the time, the world was undergoing a digital revolution, fueled by the rise of the Internet. Meanwhile, the dot-com bubble was at its peak, sparking a question on everyone's mind: Would digital technologies truly change the world, or was it all just a passing fad?

My teacher, though, was not part of everyone. She was not concerned about that at all. That morning, she had us return to the roots of nearly every modern language, guiding us through the fundamentals of translating an Ancient Greek text. Back then, there was no Google Translate, let alone ChatGPT. I remember flipping through the pages of heavy dictionaries, searching for the Greek words needed to piece

together a translation. The real challenge, however, wasn't just finding the right words. Our final grade depended on how natural and coherent our translation sounded—how well it captured the essence of the original text.

I was lucky. In my grandparents' time, they weren't even allowed a dictionary—they had to memorize most words by heart so that their grades depended not only on how well their translation read but also on their ability to recall the words entirely from memory. At least I had the dictionary to assist me. But so did all my classmates.

While my grandparents' final grade reflected both memorization and translation skills, mine was based primarily on my ability to translate. But that raised a new challenge: If everyone had access to the same tool—the dictionary—how could I stand out?

First, it came down to how effectively I used the dictionary. Not everyone approached it the same way. Some were faster; some were slower. Some relied on it mechanically, while others navigated it more strategically. But beyond that, differentiation came from how well the final translation flowed: creativity, contextual interpretation, and the subtle human touch all played a role in setting one translation apart from another.

As I lugged those heavy Latin and Greek dictionaries around, I often asked myself, "Why? Why study ancient languages that are no longer in use? Why memorize so many translation rules by heart?" At the time, AI was mostly confined to academic discussions and was far from mainstream, but even the bulky mainframe computer we had at home made me suspicious that things would change. It seemed inevitable that one day, the way through which we would study these languages might become obsolete. So, what would *studying* Ancient Greek look like in the age of AI?

With the rise of generative AI tools, particularly after the launch of ChatGPT in 2022, one of the first things that came to mind was to test its capabilities—specifically, to see if it could translate the same Ancient Greek text from my high school final exam in 2004, the *Maturità*. I remember having four hours to complete the exam, and I barely made it. I spent nearly three hours and fifty minutes meticulously translating the text, leaving myself only the last 10 minutes to review it. Many of my classmates didn't even finish in time.

The process was grueling, requiring intense focus, patience, and a deep understanding of the language.

For context, the text I had to translate for my *Maturità* exam was an excerpt from Plato titled "The Long Journey of Humanity Toward Civil Coexistence." The passage comes from Protagoras's speech in Plato's *Protagoras*, where he recounts the myth of Prometheus and the origins of human civilization.

In its original Ancient Greek, it looked like this:

> Επειδη δε ο ανθρωπος θειας μετεσχε μοιρας, πρωτον μεν δια την του θεου συγγενειαν ζωων μονον θεους ενομισεν, και επεχειρει βωμους τε ιδρυεσθαι και αγαλματα θεων: επειτα φωνην και ονοματα ταχυ διηρθρωσατο τη τεχνη, και οικησεις και εσθητας και υπο δεσεις και στρωμνας και τας εκ γης τροφας ηυρετο. Ουτω δη παρεσκευασμενοι κατ' αρχας ανθρωποι ωκουν σποραδην, πολεις δε ουκ ησαν: απωλλυντο ουν υπο των θηριων δια το πανταχη αυτων ασθενεστεροι ειναι, και η δημιουργικη τεχνη αυτοις προς μεν τροφην ικανη βοηθος ην, προς δε τον των θηριων πολεμον ενδεης - πολιτικην γαρ τεχνην ουπω ειχον, ης μερος πολεμικη - εζητουν δη αθροιζεσθαι και σωζεσθαι κτιζοντες πολεις: οτ' ουν αθροισθειεν, ηδικουν αλληλους ατε ουκ εχοντες την πολιτικην τεχνην, ωστε παλιν σκεδαννυμενοι διεφθειροντο.

Almost 20 years later, I decided to put AI to the test. I asked ChatGPT:

> If you were an 18-year-old Italian student taking the prova di maturità and wanted to get the highest grade, how would you translate the following text from Ancient Greek to English?

The result was mind-blowing:

> Since man partook in a divine nature, he was, first of all, the only living being to conceive the idea of gods due to his kinship with the divine. Thus, he began to build altars and statues in their honor. Then, with the aid of his craft, he swiftly articulated speech and created names, and he invented dwellings, clothing,

> footwear, bedding, and means of sustenance derived from the earth.
>
> At first, however, human beings lived scattered, without cities; as a result, they were being annihilated by wild beasts, for they were weaker than them in every way. Their technical skills provided sufficient aid for survival but were inadequate for waging war against the beasts, since they had not yet acquired the art of politics, of which warfare is a part. Consequently, they sought to gather together and ensure their safety by building cities. However, once assembled, they began harming one another, as they lacked the knowledge of political organization; thus, they dispersed again and continued to perish.

The first particularly interesting thing I noted is how this excerpt—perhaps a premonition—serves as a powerful analogy between fire and AI. Just as Prometheus gave fire to humans, granting them the ability to transform their environment and propel civilization forward, we could say that AI is our modern-day fire. It is a tool of immense power, capable of elevating humanity to new heights, yet one that also carries profound existential challenges. In the passage, early humans lacked political wisdom, leading to chaos and destruction when they gathered. The same pattern is emerging with AI today: we have incredibly powerful tools—Large Language Models, automation, predictive analytics—but without clear regulations, ethical frameworks, and governance, these technologies risk causing more harm than good.

The lesson from this passage is not just about survival but about the crucial role of wisdom in handling powerful tools. AI, much like Prometheus's fire, is neither inherently good nor bad; its impact depends entirely on how we use it. If we fail to establish a guiding framework, much like the *art of politics* in the myth, we risk plunging into disorder rather than progress.

The second thing that truly impressed me was the speed at which AI translated the passage—roughly 10 seconds. And you know what? It sounded pretty good. I'd give it a solid 9 out of 10. "Why wasn't ChatGPT around back then?" I couldn't help but think as I reviewed the AI-generated translation. With nearly four hours left on the clock, I could have spent all that time refining the text, ensuring every

sentence flowed naturally, and aiming for a perfect score. Because yes, some parts of the AI translation sounded slightly off—but the same was true for my original translation. The difference was that I didn't have the luxury of time to refine mine. Although I ended up getting a very good grade, I could have done much better with the use of AI.

The key learning from this story is that the best outcomes from using tools that not only outperform human skills but also operate at greater speed come from the combination of two factors. First, whenever we have a tool—whether it's something as simple as a dictionary or as sophisticated as AI—what ultimately sets us apart isn't the tool itself, but how well we use it. If everyone has access to the same technology, the real differentiator lies in how effectively we leverage it.

Second, the final outcome of any tool, no matter how advanced, is shaped by the human contribution we add to it. A tool can assist, enhance, or even take over parts of a task, but its true impact depends on the creativity, judgment, and refinement we bring to the output—often using the time saved by the tool to elevate the final result.

A New Skills Paradigm

By automating and augmenting many of our tasks, AI not only saves us time but also reduces our reliance on technical skills. This shift presents an opportunity—or perhaps even a necessity—to develop and refine human skills that were previously overlooked. These skills were often deprioritized because they were difficult to measure, time-consuming to develop, and perceived as inefficient in traditional work environments.

AI is ushering in a new skills paradigm—one we call the *hybrid skillset*. To thrive in this evolving landscape, professionals and leaders must urgently adapt by developing two complementary sets of skills: on the one hand, *AI literacy skills* that make us effectively leverage AI tools, navigate their complexities, and maximize their potential; on the other hand, *human literacy skills*—the uniquely human capabilities that AI's efficiency allows us to focus on, such as critical thinking, creativity, empathy, and decision-making.

This shift is not just about technology—it's about us, humans. The professionals, leaders, and individuals who choose to master both AI

literacy and human literacy will shape the future. The rest? They risk becoming relics of the past—outpaced not by machines, but by those who understand how to work alongside them.

To understand and develop these skills, we must first map out a fundamental question: What does AI now do *better* than humans, and what will remain uniquely human?

The End of Human Uniqueness?

Between February and July 2024, researchers conducted a simulation[5] involving 344 participants facing real-world decisions typically made by industry CEOs. These decisions were based on historical pricing, market shifts, sales trends, and disruptions like COVID-19.

Participants were up against a powerful opponent: OpenAI's GPT-4o, the latest version of its LLM. The goal of the game was simple: while maximizing market capitalization, survive as long as possible without being fired by a virtual board.

The results were striking: GPT-4o consistently outperformed human participants, designing products with cost efficiency, responding swiftly to market signals, and achieving greater profitability three rounds ahead of the top human performer.

These findings are consistent with what Scott Aaronson, professor of Computer Science at the University of Texas, explained in his keynote address at TEDx Palo Alto.[6] He highlighted what he sees as "the problem with human specialness in the age of AI"—namely, that for any task with a reasonably objective measure of success or failure, such as games, competitions, or structured challenges, AI will inevitably match or surpass human performance.

In the OpenAI GPT-4o simulation, AI did struggle with outlier events like COVID-19—highlighting the fact that AI's success depends on the availability of past data patterns.

But according to Aaronson, "Given any task with a reasonably objective metric of success or failure, and on which an AI can be given suitably many examples of success and failure, it's only a matter of years before not only AI, but AI on our current paradigm, will match or beat human performance." In other words, if a task can be quantified and AI has access to sufficient training data, its dominance over humans is

just a matter of time. Aaronson refers to this inevitability as the "Game Over Theory."

As humans, it has been years since we had any hope of outperforming AI in games like chess and Go—milestones that are now well behind us—and it's only a matter of time before AI surpasses human performance in other competitive domains, such as the Mathematics Olympiad. GPT-4, for instance, can pass the Uniform Bar Examination with ease, scoring in the top 10% of test takers. It can also answer 90% of questions correctly on the U.S. Medical Licensing Examination. By the time you're reading this book, AI will likely have achieved even more advanced feats.

AI in Cancer Diagnosis[7]

Scientists at Harvard Medical School have developed a versatile, ChatGPT-like AI model capable of performing a wide range of diagnostic tasks across multiple forms of cancer. Named CHIEF, this AI system goes beyond typical models that are trained for specific tasks, such as detecting the presence of cancer or predicting a tumor's genetic profile. Using 15 million unlabeled images, CHIEF was tested on 19 cancer types and outperformed existing models across 11 datasets.

With an accuracy rate nearing 94%, CHIEF exemplifies how general-purpose AI is expanding beyond narrow tasks—tackling complex problems that once required highly specialized human expertise.

Differences in Operating Strengths

But how do we determine if AI is better than we are in fields like writing, the humanities, or creativity, where there are no reasonably objective metrics of success? Measuring performance in these areas isn't straightforward—there's a high degree of subjectivity involved. For example, would you say your writing is better than ChatGPT's? What would your friends, readers, or family say? (Though let's be honest,

your mom doesn't count—she's always going to say your writing is better.)

To some extent, rather than asking which is *better*, it may be more useful to explore the differences in *how* each operates. If AI is transforming the way we work and think, how do its capabilities compare to those of human intelligence?

The table below maps out the differences between human and artificial intelligence across key cognitive and functional domains, illustrating how humans and AI process information, make decisions, and interact with the world. Understanding these distinctions can help us better navigate the evolving relationship between humans and AI.

Domain	Human Intelligence	Artificial Intelligence
Decision-Making	Humans make decisions using reasoning (deductive, inductive, and abductive), intuition, and experience. Decision-making is influenced by emotions, biases (e.g., confirmation bias, availability heuristic), and contextual understanding.	AI makes decisions through logical, data-driven models such as decision trees and neural networks. It follows consistent patterns based on data but may inherit biases from training datasets. It lacks human-like intuition and contextual adaptability.
Creativity	Human creativity emerges from emotional experience, cultural influences, and the ability to connect unrelated ideas. It is shaped by curiosity, intrinsic motivation, and personal expression.	AI generates creative outputs (e.g., music, art, and text) using algorithms trained on large datasets. While it can replicate patterns and styles, it lacks intrinsic motivation, originality, and emotional depth.

Domain	Human Intelligence	Artificial Intelligence
Learning	Humans learn through experience, observation, trial and error, and social interactions. Learning is adaptable, requires minimal data, and generalizes across different contexts.	AI learns by processing large datasets through supervised, unsupervised, or reinforcement learning. It requires extensive data, struggles to generalize outside its training set, and lacks experiential learning.
IQ-Related Tasks	Humans demonstrate intelligence through logical reasoning, verbal comprehension, problem-solving, and abstract thinking. Performance on IQ tests is influenced by education, culture, and prior knowledge.	AI can outperform humans in IQ-related tasks such as pattern recognition, logical reasoning, and mathematical problem-solving. However, its intelligence is narrow, task-specific, and lacks generalized understanding.
Task Execution	Human task performance is influenced by cognitive load, physical limitations, time constraints, and external distractions.	AI executes tasks with near-limitless efficiency, constrained only by computational power. It can process vast amounts of data rapidly but lacks adaptability beyond predefined parameters.

(*continued*)

(*continued*)

Domain	Human Intelligence	Artificial Intelligence
Content Production	Humans create content based on personal experiences, emotions, and cultural context. Writing, storytelling, and artistic expression are deeply influenced by intention and creativity.	AI generates text, images, and media based on training data and structured prompts. While it can mimic human language and style, it lacks true intent, emotional nuance, and original thought.
Emotional Intelligence	Humans possess self-awareness, empathy, emotional regulation, and social adaptability. These skills are critical for leadership, communication, and interpersonal relationships.	AI can analyze sentiment and detect emotions using affective computing but does not experience emotions. Its responses are simulated and lack genuine empathy or self-awareness.
Social Cognition	Humans understand social dynamics, read nonverbal cues, and infer the mental states of others. Social cognition is essential for building relationships and navigating complex interactions.	AI can simulate social interactions (e.g., chatbots, virtual assistants) but lacks genuine understanding of human emotions, intentions, and social complexities. Its interactions are scripted and rule-based.

Domain	Human Intelligence	Artificial Intelligence
Trust-Building	Humans build trust through relationships, consistency, emotional connection, and social experiences. Trust is dynamic and influenced by past interactions and context.	AI does not inherently build trust but earns it through transparency, reliability, and accuracy. Trust in AI depends on explainability and the absence of bias rather than emotional or relational factors.
Perception	Humans integrate sensory information seamlessly, influenced by experience, emotions, and cognitive biases. Perception is adaptable but prone to illusions and misinterpretations.	AI excels in processing images, recognizing patterns, and analyzing audio through advanced algorithms. However, it struggles with ambiguous, incomplete, or conflicting sensory information and lacks holistic multisensory perception.

While this comparison highlights key differences between Human Intelligence and Artificial Intelligence, determining which is better at a given task is not always straightforward. Many skills exist on a spectrum—some have objective measures of success, while others are deeply subjective.

Skills with inherently subjective success criteria are not only less susceptible to automation but are also becoming more valuable in the age of AI. In fields like mathematics, games, and many aspects of business, performance is measured by clear-cut metrics—either you solve the problem or you don't; either you win or you lose. However, in areas such as business strategy, writing, art, and music—where creativity and personal expression take center stage—the concept of *better* is far more fluid. It is shaped by individual preferences, cultural context, and the inherent subjectivity of human taste.

Hard and Soft Skills in the Age of AI

It's no coincidence that the distinction between objective and subjective AI outcomes aligns with the divide between two major skill categories: hard skills and soft skills. And it is no coincidence that in the age of AI, the balance between these two is shifting dramatically.

In 1968, the upper ranks of the U.S. Army faced a pressing challenge: certain factors were directly influencing the success of military units, yet these elements were neither measured nor accounted for in training programs, development plans, or performance evaluations.

In other words, something intangible within battalion structures was shaping soldier behavior—at times fostering dedication and cohesion, at other times fueling internal conflicts or disrupting the flow of critical information between units. This unseen force was deeply affecting both the performance of officers and the outcomes of missions, yet it remained unquantified, unstructured, and unaddressed.

Lacking a precise definition, the Army's top minds launched a rigorous academic effort to categorize these complex, intangible abilities. They labeled them *soft skills*—a stark contrast to the easily measurable *hard skills* like driving a vehicle, operating a machine gun, or preparing a logistics report.

Despite acknowledging the undeniable value of soft skills—such as the ability to inspire, build trust, or think flexibly—Army scholars struggled to develop concrete methods for identifying, measuring, and teaching them. The following traits have traditionally made soft skills more challenging to prioritize than hard skills:

1. **Broad (vs. hard skills, which are specific).** Unlike hard skills, which are task-oriented and tailored to specific functions, soft skills are universally applicable across various business and life situations. Their versatility makes them indispensable for navigating complex environments, adapting to different contexts, and fostering effective human collaboration.
2. **Difficult to transfer (vs. hard skills, which are easier to teach and transfer).** While technical skills can be systematically taught and passed down, allowing one person to directly help another acquire the same ability, soft skills are far more personal and internal. Their development relies less on formal

instruction and more on individual growth, lived experiences, and self-awareness, making them harder to replicate or standardize.

3. **Difficult to measure (vs. hard skills, which are quantifiable).** Assessing a software developer's coding proficiency or a candidate's fluency in English is relatively straightforward since hard skills come with clear metrics, tests, and benchmarks. In contrast, evaluating someone's ability to collaborate, demonstrate empathy, or think critically is far more complex. Soft skills lack standardized measurement tools, making it challenging for organizations to assess and develop them effectively.

Paradoxically, despite the undeniable importance of soft skills, we often gravitate toward objective, quantifiable targets simply because they are clearer and leave little room for ambiguity. Unlike behavioral goals—where success can feel vague or subjective—numeric targets provide a sense of certainty, making them easier to track and evaluate, but this preference for clarity often comes at the cost of overlooking the skills that truly drive long-term success.

Soft Skills in Practice: A Normal Day at the Office

Your manager calls you in for an end-of-year feedback session to set goals for the upcoming year. They start with the easy part—clear, quantifiable targets like "achieve 15% sales growth in Q1." These are concrete, measurable, and straightforward to track.

But then comes a different kind of objective: "Your goal this year is to be more collaborative." You nod politely, thank her, and walk out of the room—only to feel a wave of uncertainty.

What does being more collaborative actually mean? How do you show it? How do you measure where you stand today, or by how much you need to improve? And most important, how do you develop that skill in the first place?

Because they are so broad, difficult to teach, and hard to measure, soft skills have traditionally been deprioritized in the workplace—an inheritance of the Industrial Revolution mindset that prioritizes efficiency, growth, and maximized outcomes. However, in the age of AI, the very fact that soft skills are difficult to quantify doesn't make them less important—it makes them *more* important, as these skills are precisely what make humans complementary to AI rather than replaceable by it. As AI automates technical tasks and accelerates decision-making, the ability to think critically, collaborate, and bring emotional intelligence to problem-solving will be what truly sets us apart.

Data from our survey reinforce this point: 82% of respondents said they would prefer to hire a candidate who lacks sufficient hard skills for the role but has strong soft skills, rather than the opposite. When asked why, two reasons stood out in their comments. First, unsurprisingly, respondents noted that hard skills are much easier to teach and develop than soft skills. But the second reason is more surprising: as companies increasingly adopt AI tools, many hard skills are being automated and replaced by machines, while soft skills remain uniquely human and irreplaceable.

Why IQ Isn't Enough Anymore

This shift in focus—from hard skills to soft—was not always obvious to me. I grew up in a world that taught me the exact opposite—that my focus should be on developing strong hard skills, often measured by Intellectual Quotient (IQ) tests, while soft skills were rarely mentioned, let alone prioritized.

That lasted up until I read Daniel Goleman's *Emotional Intelligence*,[8] a book that completely shifted my perspective. It introduced me to the idea that while IQ may provide an early advantage, its influence diminishes over time: once inside the workplace, success depends far more on emotional intelligence (EI).

The transition from academia to the professional world makes this clear. In school, success is determined by structured tests, clear-cut problems, and individual performance. But the workplace is far more complex—an unpredictable environment where ambiguity is the

norm. Projects shift, stakeholders have conflicting interests, crises emerge unexpectedly, and success is determined not just by what you know, but by how well you navigate relationships, communicate effectively, and manage your emotions under pressure.

This is where IQ's limitations become apparent. High cognitive ability might help process information quickly or solve complex analytical problems, but it does little to address the human side of work. In a world where collaboration, negotiation, and adaptability are essential, interpersonal skills often outweigh raw intellectual horsepower.

Enter emotional intelligence: unlike IQ, which remains relatively stable, EI is a skill that can be developed and refined. Research consistently shows EI as a key predictor of workplace success, particularly in roles that require teamwork, leadership, and high levels of interaction.

Two studies highlight this point. David McClelland's 1998 study, "Identifying Competencies with Behavioral-Event Interviews,"[9] found that while cognitive and technical skills are necessary to enter many professions, emotional competencies are what truly distinguish top performers, especially in leadership. Similarly, Sigal Barsade and Olivia O'Neill's 2014 study, "What's Love Got to Do with It?"[10] revealed that teams with high collective EI exhibited greater cooperation, higher creativity, and overall better performance—demonstrating that EI is not just essential for individual success but also for building strong organizational cultures.

These insights heavily influenced Goleman's work and are even more relevant in today's AI-driven world. If IQ serves as the foundation for technical skills, EI is the fertile ground for socio-emotional skills—where, as we have established, humans still hold a decisive advantage over AI. When these human literacy skills are paired with AI literacy skills, they form the core of a new paradigm for thriving in the age of AI.

Thinking, Acting, and Feeling: The Core of What Makes Us Human

To define the new skills required to thrive in the age of Artificial Intelligence—what we have called the *hybrid skillset*—we must begin by examining the core activities that have always defined human

existence. Not specific tasks like cooking, writing poetry, or filling out an Excel spreadsheet—these vary from person to person. Instead, we need to focus on the fundamental, universal activities that all humans engage in. No exceptions. At our core, we do three main things: we think, we act, and we feel.

First, *thinking* encompasses our cognitive processes—reasoning, analyzing, imagining, problem-solving, and decision-making. It's how we process information, plan, and make sense of the world around us. It's also how we generate creative ideas and find solutions to problems. We engage in thinking constantly. A study[11] by psychology researchers at Queen's University in Canada, which analyzed brain scans, found that the average person has more than 6,000 thoughts per day, which is why it is so hard to get into that meditative state of "no thoughts"—our brains are wired to constantly generate new thoughts.

Second, *acting* involves the execution of plans, initiatives, and behaviors that produce tangible outcomes. It's about transforming thoughts into reality—creating tools, building systems, interacting with the environment, and bringing ideas to life. However, this is where a common challenge arises: the *insight–action gap*. While generating thoughts is effortless, acting on them is far more complex. Action requires effort, carries risks, and is often constrained by external factors. This gap explains why many ideas never materialize—because turning intention into execution is a far greater challenge than simply thinking about it.

Third, *feeling* encompasses emotions, intuition, and subjective experiences. It's how we connect with ourselves and others; process emotions like joy, fear, love, and sadness; and navigate interpersonal relationships. Feeling also plays a crucial role in self-awareness, shaping our perceptions, decisions, and interactions with the world around us.

Obviously, these three macro activities rarely function in isolation—they are deeply interconnected. Thoughts influence feelings, feelings drive actions, and actions, in turn, generate new thoughts and emotions. Likewise, decisions made by our brains often (though not always) translate into actions.

At any given moment, as long as we are alive, we are engaged in one or a combination of these three dimensions. Unlike external circumstances—such as wealth, social status, or age—these fundamental activities remain constant. Whether rich or poor, young or old, part of one culture or another, we all think, act, and feel. Even if we renounce all material possessions and embrace a secluded monastic life, we are still bound to these three forces. They represent, in essence, what it means to be human.

It is no different in the workplace. Regardless of our role, industry, company size, or seniority level, these three human dimensions—thinking, acting, and feeling—are constantly at play. We think creatively during brainstorming sessions to generate original ideas, analyze data to determine the best strategy, and predict emerging business trends that could shape our industry.

We act by executing tasks, implementing strategies, and delivering tangible results. Whether coding software, assembling products, or presenting a business case, this dimension is where ideas and plans take form. Action is the backbone of progress—it's what transforms strategy into reality.

And although many might argue that there's little room for feeling in the workplace, they couldn't be more wrong. We feel when we make intuitive decisions based on a gut instinct, when we seek to understand customer behavior through empathetic conversations, and when we navigate team dynamics and workplace culture. Feeling is embedded in trust, engagement, and motivation—the driving forces behind both individual and collective success. It shapes how we lead, collaborate, and resolve conflicts, making emotional intelligence an essential component of any thriving professional environment.

The Hybrid Skillset Framework

As we stand at the intersection of Human Intelligence and Artificial Intelligence, it becomes clear that success in this new era will not come from competing with AI in areas where AI excels, but from complementing AI use with skills that remain uniquely human. The way forward is not about choosing between AI and human expertise—it is

about integrating both into a new hybrid set of skills that leverages the best of each.

This is where the Hybrid Skillset comes into play. Rather than being a single set of competencies, the hybrid skillset is best understood as a dynamic interplay of cognitive, behavioral, and emotional abilities—skills that enable professionals to work effectively alongside AI. To make this framework more actionable, we can break it down into three core pillars:

- **Cognitive Hybrid Skills:** AI can better process data, but humans must define the right problems, contextualize information, and generate original ideas. This requires expertise in skills like Prompting, Data Sensemaking, and Reperception (see Section I)—skills that ensure AI's outputs are meaningful and relevant.
- **Behavioral Hybrid Skills:** AI is efficient in executing tasks but lacks adaptability and true innovation. Humans must embrace skills like Augmentation, Adaptability, and Antifragility (see Section II)—allowing AI to handle routine tasks while focusing on experimentation and learning from uncertainty.
- **Emotional Hybrid Skills:** AI can recognize patterns in emotions but does not feel, empathize, or build deep trust. The most enduring human advantage will lie in skills like Empathy, Trust, and Agency (see Section III), as businesses and leaders recognize that emotional intelligence is the true differentiator in an AI-driven world.

The table below maps out key human activities across cognitive, behavioral, and emotional dimensions, illustrating the degree to which AI can be a substitute for these tasks, what AI does best, and where human strengths remain indispensable. It also highlights the critical skills professionals must develop to effectively collaborate with AI—ensuring that we don't just adopt new technology but actively shape its role in our work and decision-making.

Pillar	Human Activity	What AI Does Best	What Humans Do Better	Key Skill	Definition	Associated Behaviors
Cognitivez	Learning	Processes vast data, identifies patterns, retrieves knowledge	Prioritizes key info, defines questions, applies context	**Prompting**	Interacting effectively with AI through precise, structured inputs	Asking strategic, well-crafted questions to get optimal outputs
	Decision-Making	Analyzes data, finds correlations, predicts outcomes	Applies intuition, ethics, judgment	**Data Sensemaking**	Critically assessing AI output to extract meaningful, actionable insights	Evaluating recommendations through strategic and ethical lenses
	Idea Generation	Synthesizes knowledge, optimizes patterns	Integrates diverse experiences, has cognitive flexibility	**Reperception**	Seeing problems from new angles, challenging assumptions, finding new connections	Reframing challenges, generating novel perspectives

(*continued*)

(*continued*)

Pillar	Human Activity	What AI Does Best	What Humans Do Better	Key Skill	Definition	Associated Behaviors
Behavioral	Performing Familiar Tasks	Automates high-volume, familiar and repetitive tasks	Enhances strategy, brings quality and nuance	**Augmentation**	Leveraging AI to improve—rather than just replace—human work	Using AI for efficiency while adding critical oversight
	Performing Unfamiliar Tasks	Predicts change using past data	Adapts flexibly, improvises, handles complexity	**Adaptability**	Adjusting rapidly to new conditions and staying effective under pressure	Leading through uncertainty, making decisions in gray areas
	Innovation	Improves iteratively within defined frameworks	Thrives in uncertainty, learns from failure	**Antifragility**	Minimize cost of mistakes and maximize the learning from them	Embracing failure, experimenting quickly, and learning from mistakes

Pillar	Human Activity	What AI Does Best	What Humans Do Better	Key Skill	Definition	Associated Behaviors
Emotional	Connection	Detects emotional patterns, simulates responses	Feels, empathizes, and expresses emotion authentically	**Empathy**	Understanding and sharing others' emotional states	Active listening, emotional attunement, compassionate feedback
	Collaboration	Coordinates communication, schedules workflows	Builds trust, navigates team dynamics	**Trust**	Creating psychological safety and relational credibility	Demonstrating integrity, reliability, emotional presence
	Ethics	Goal maximization but without self-awareness or conscience	Ethical concerns limit goal maximization based on self-awareness and conscience	**Agency**	Taking accountability for its own actions, even if we cannot control them	Applying ethical standards to the use of AI

The reality is that AI will continue to advance, and the line between what it does best and what we do best will keep shifting. But the key to staying ahead is not resisting AI—it is mastering how to work with it. Those who develop a Hybrid Skillset will not only future-proof their careers but redefine the workplace in ways we are only beginning to imagine.

This book is a roadmap to mastering this skillset transformation. In the next chapters, we will explore how organizations and individuals can actively cultivate these skills—and what it takes to thrive in a world where AI is no longer just a threat, but a superpower to make us thrive in the workplace.