WHITE PAPER

Artificial General Intelligence

Demystifying AI's Next Frontier

September 2024
By Akram Awad, Lars Littig, and Sophie Geraerdts

As Artificial Intelligence (AI) continues to achieve significant breakthroughs, some experts believe we are nearing the point where machines develop human-level cognitive abilities. Visionaries like Elon Musk, Sam Altman, and Ray Kurzweil have all suggested that AI could match or surpass human intelligence in coming years—with Musk predicting it could happen as soon as 2026.

But what exactly does this mean? Historically, AI has been divided into three levels. The first level is **Artificial Narrow Intelligence** *(ANI)*, which is also known as "Narrow AI" or "Weak AI." ANI refers to artificial intelligence that is trained in specific tasks within familiar environments. The second level is **Artificial General Intelligence** *(AGI)*, which is also called "General AI" or "Strong AI." AGI displays human-level intelligence and adaptability across a wide range of tasks. The third, and most advanced, level is **Artificial Superintelligence** *(ASI)*, which surpasses human intelligence in all areas and operates without human input or oversight (Exhibit 1).

Until recently, it has been widely accepted that we are firmly in the ANI era. But the rapid advances in **Generative AI**—an ANI subset that emulates human behaviors and capabilities—have fueled speculation that we might be entering the age of AGI.

In this article, we will demystify AGI by exploring what it entails and examine the key initiatives and investments being made by different stakeholders. We also will delve into the uncertainties surrounding AGI's timeline, the potential risks and ethical considerations it brings. Most importantly, we will also outline the critical steps required to assure that this transformative technology is used responsibly.

Exhibit 1. The 3 Stages of Artificial Intelligence

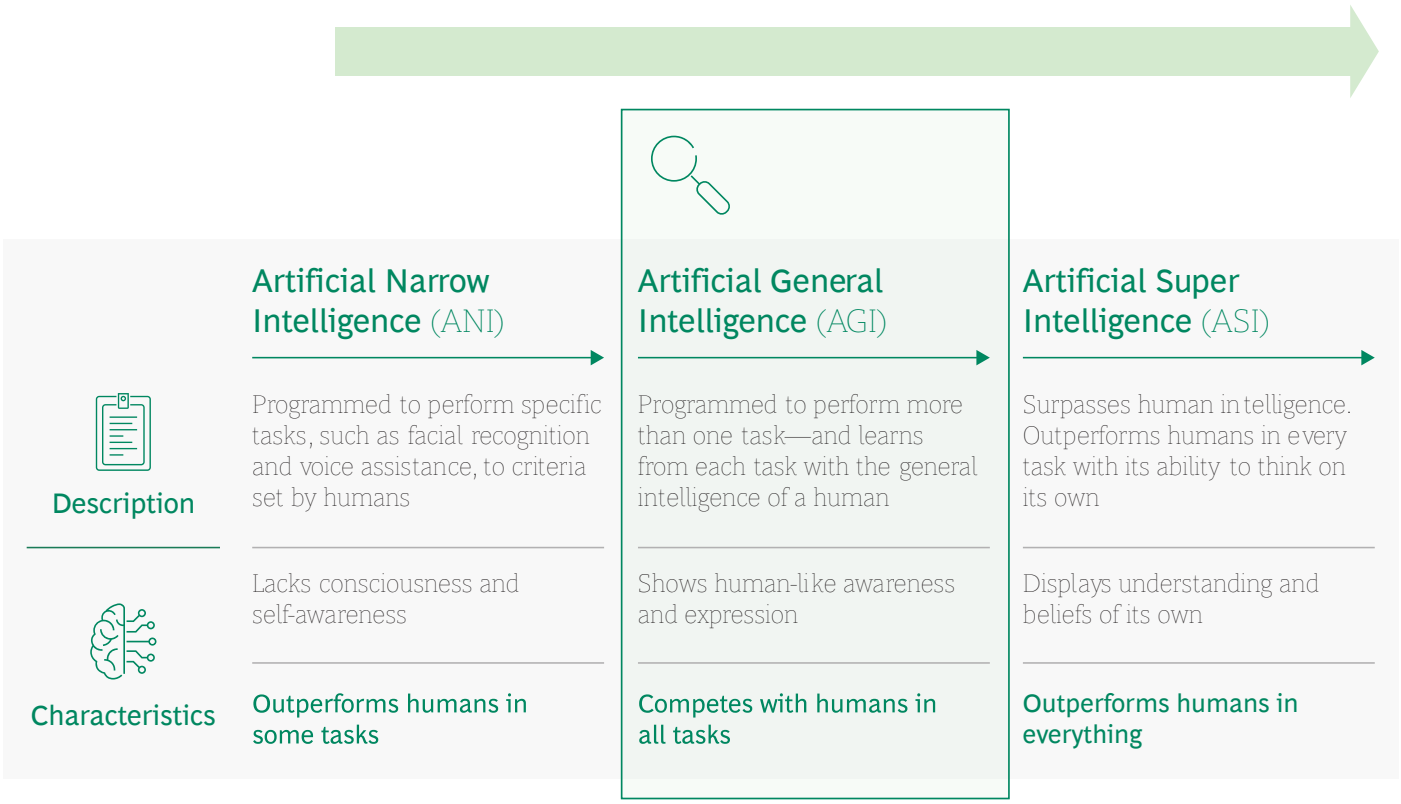

Sources: Expert interviews; BCG analysis.

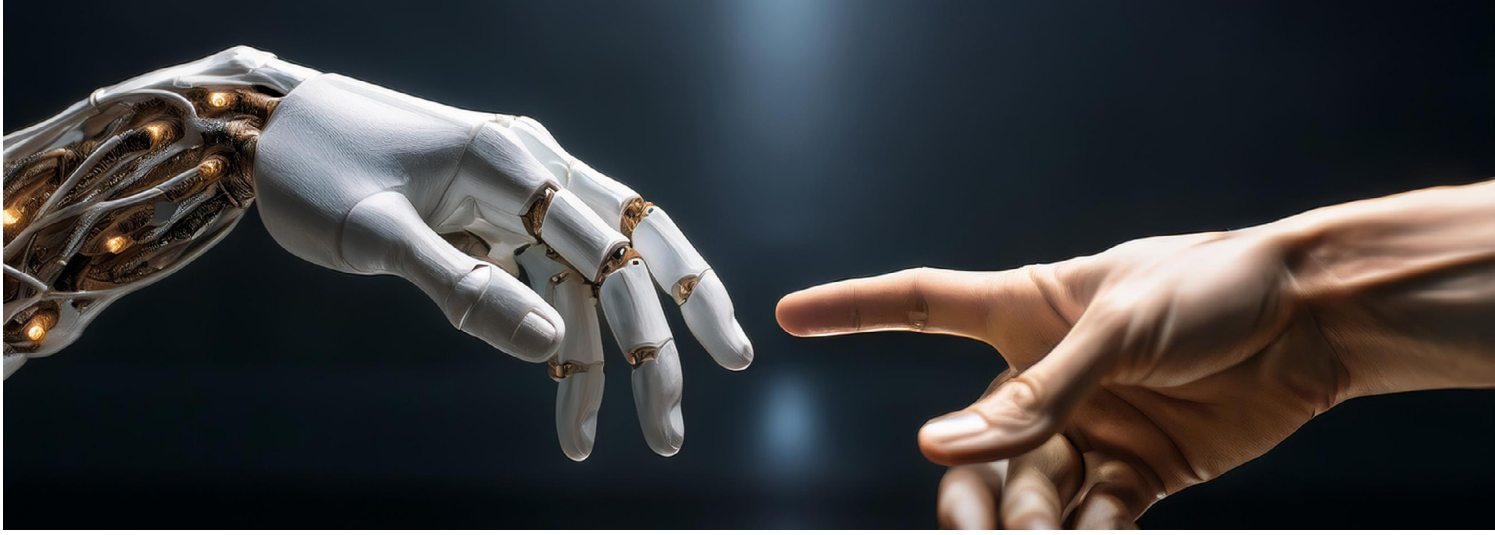

What is Artificial General Intelligence?

The concept of AGI has captured the imagination of technologists, academics, and policymakers since the 1950s. And yet, despite its growing popularity, there is no standard definition of AGI. Most agree that AGI will be as smart as humans (or even smarter), with the ability to reason, learn, and apply knowledge across a range of tasks. But the specifics of what this means—and how AGI systems will work—varies depending on who you ask.

For instance, neurobiologists focus on AGI's potential to have human-like motivations, emotions, and decision-making processes. Meanwhile, machine learning experts study whether AGI can truly think logically, draw casual inferences, and learn and adapt like the human brain does. For their part, philosophers debate whether AI systems can ever attain human-level intelligence or merely simulate it. This lack of consensus on the exact definition and requirements of AGI highlights the complexity and challenge of replicating the full breadth and depth of human intelligence through artificial means.

Despite the lack of a shared definition, certain intelligence and physical traits are widely recognized as indicative of AGI (Exhibit 2). For something to be considered AGI, it must meet all the criteria below, simultaneously. Each trait alone is not enough; it is the combination of all these abilities that defines AGI:

Intelligence Traits

Intelligence traits are related to AI's ability to learn and adapt. This is on display in how AGI might be used to optimize a company's supply chain. AGI must possess the ability to:

Reason under uncertainty. Reasoning amid uncertainty means making rational decisions when information is incomplete or ambiguous. It also means using different cognitive skills with little to no human intervention. This includes reasoning that is deductive (logical), inductive (probabilistic), and abductive (inferential). AGI's ability to reason is of value to a company looking to optimize its supply chain. Supply chain managers often have to deal with incomplete data about supplier reliability and customer demand. That means the AGI system must weigh the possible risks in using limited data—and make educated guesses about how best to optimize the supply chain.

Use "common sense." This trait enables the system to understand basic facts about the world and everyday life. In the case of the company looking to optimize its supply chain, the AGI system diversifies suppliers to reduce risks—assuming that's a common practice in supply chain management. Likewise, it implements a "just-in-time" management process to reduce inventory costs.

Make decisions autonomously. This attribute involves solving complex problems on its own. In the example above, the AGI system negotiates contracts and adjusts inventory levels without human intervention. It adds new suppliers that meet the company's guidelines for quality and cost—and adjusts inventory to match demand.

Communicate in natural language. This trait enables the AGI systems to interact, understand, and respond in a human-like manner. For example, the AGI system helps the supply chain manager draft emails and reports to update senior executives, using natural language to explain the new optimization strategy and the expected benefits.

Transfer learning and understand cause and effect. This characteristic enables an AI system to apply knowledge from one task to another. In the scenario above, the AGI system applies what it knows about best practices from other industries to fit this company's particular operations. The AGI anticipates the consequences of its decisions—including the risk that new suppliers could prove less reliable, and that its "just-in-time" inventory system could sometimes leave it out of stock.

Physical Traits

Physical traits describe how AGI is physically integrated into—and interacts with—its surroundings. To qualify as AGI, it must have all of the following abilities:

Sight and Hearing. AGI should be able to process environmental sensory data like humans can. This includes:

- **Visual Inspection:** An AGI-enhance supply chain system can use cameras to scan and identify products that need to be sorted, stored, or shipped.
- **Auditory Monitoring:** Having been programmed with sophisticated auditory capabilities, the AGI system listens to the sounds of running machines to detect issues or inefficiencies.

Performing complex movements: For the company above, AI systems with this physical can direct robots to sort and pack goods efficiently in the company's warehouse. The AGI system must also be able to adjust inventory levels across the company's distribution network in response to regional changes in customer demand.

Exhibit 2. Despite Lack of Definition, Some Traits are Accepted as AGI

Sources: Expert interviews; Russell & Norvig (2003); Luger and Stubblefield (2004); Poole, Mackworth, Goebel (1998); Nilsson (1998); Pfeifer, R., Bongard, J. C., and How The Body Shapes The Way We Think: a new view of intelligence (The MIT Press, 2007).

Barriers to Achieving AGI

While the concept of AGI signifies artificial intelligence that can rival—and even exceed—the breadth and flexibility of human cognition, the path to AGI remains complex and uncertain. Researchers and developers face a host of technical, ethical, and regulatory hurdles that will make the journey forward difficult, costly, and often controversial (Exhibit 3). These include:

- **Lack of Computing Power:** The current infrastructure is inadequate to manage the immense calculations required for AGI, both in terms of computational capability and energy consumption. Quantum computing is seen as a potential game-changer, offering exponential increases in computational power that could bring us closer to the capabilities needed for AGI. Meanwhile, some tech giants are exploring innovative energy solutions, such as building small nuclear reactors to provide the substantial energy demands that their AI data centers require (and exceed what local power grids can supply).

- **Algorithmic Complexity:** The algorithms that underpin current AI models remain relatively basic, particularly when compared to the intricate reasoning and decision-making processes that AGI will require. Significant research is needed to enhance how AI models learn, reason, and make decisions, especially in the "autoregressive loops" that are critical for tasks such as natural language generation. These loops allow AI to consider the context of previous information to accurately predict future outputs—a capability that is essential for AGI.

- **Ethical Considerations:** As AI evolves toward AGI, ethical concerns about its impact on society grow increasingly complex. AGI systems, with their human-level intelligence, will influence many facets of human life—from employment and healthcare to surveillance and privacy. For example, AGI might face ethical dilemmas when recommending medical treatments or inadvertently reinforce biases in decisions like hiring or firing. Ensuring transparency, fairness, and accountability in AGI systems is crucial for their responsible and ethical development.

- **Workforce Implications:** The rise of AGI poses significant challenges for the global workforce. As AI edges closer to human-level intelligence, the pressure to keep people employed in meaningful work will intensify. The International Monetary Fund has predicted that 60% or more of jobs in advanced economies may be impacted by AI. To adapt, companies will need to invest heavily in recruiting AI-savvy talent, upskilling their current workforce, and reengineering their processes to integrate AGI effectively. By devoting 10% of their effort to developing algorithms, 20% to building a robust data infrastructure, and 70% to business and people transformation, organizations reap the benefits of AI..

- **Governance and Regulation:** The development of AGI will necessitate new governance models and regulatory frameworks to protect workers and society at large. While new laws like the European Union's AI Act are vital to ensure ethical AI development, they risk stifling innovation if written too restrictively. Governments must find a balance by creating flexible regulatory approaches such as "regulatory sandboxes" that allow for experimentation and growth while maintaining safety and compliance.

Exhibit 3. A Blueprint to Overcome Current Challenges

Barriers to the development of AGI	Roadmap to overcome barriers
Lack of computing capabilities	Quantum computing could provide the processing power required to develop AGI
Increase in AI-limiting regulations	Government leaders and regulators provide more incentives and support to advance the development of AGI, including creation of regulatory sandboxes
Societal concerns over the ethics of AI	Industry and policymakers jointly identify guardrails to advance AGI responsibly, transparently and with accountability
Algorithms have limited capabilities	Companies continue to invest in R&D, including AI models' technical ability to reason (the autoregressive loop)
Challenges to the global workforce	Companies will need to invest heavily in recruiting AI-savvy talent, upskilling their current workforce, and reengineering their processes to integrate AGI effectively

Sources: Expert interviews; BCG analysis.

Who Is Focusing on AGI and What Is Happening in the Space?

While significant barriers stand in the way of achieving AGI, these challenges have not deterred a surge of investment and innovation in the AI space. Global private investment in generative AI, for example, skyrocketed from roughly $3 billion in 2022 to $25 billion in 2023, according to Stanford University. This surge is driven by players across various sectors trying to push AI systems to exhibit the human-like intelligence that characterizes AGI. These entities—ranging from big tech companies to academic institutions and industry leaders—are channeling resources into technologies that could make AGI a reality. Below is a closer look at the main contributors driving this momentum:

- **Big Tech Companies:** Companies like Google, Microsoft, and OpenAI are at the forefront of AGI research, investing heavily in the development of foundational technologies and infrastructure. These companies are exploring everything from advanced machine learning algorithms to sophisticated sensors that enable AI to adapt to its physical surroundings. For instance, Google is developing Pathways, a next-generation AI architecture designed to handle multiple tasks simultaneously and learn new tasks quickly. This approach allows a model trained on one task—such as predicting landscape elevation from aerial images—to learn another task much faster, like predicting the flow of flood waters through that same terrain.

- **Academic Institutions:** Universities and research centers are focusing on the theoretical, ethical, and practical implications of AGI. They are also working on developing the next generation of AI talent, which is crucial for driving innovation in this field. For example, researchers at Delft University of Technology developed a new method in 2024 that helps AI learn from very small amounts of data more effectively, an attribute that has long distinguished human learning. Moreover, several researchers study the concept of consciousness and how it relates to AI.

- **Industry Players:** Large companies in sectors like healthcare, finance, and manufacturing are exploring how AGI can transform their operations. They are investing in pilot projects and collaborations with AI startups to integrate AGI technologies into their workflows.

- **Institutional Investors:** Venture capital firms and sovereign wealth funds are pouring money into AI startups and research initiatives.

The role of governments in this ecosystem is equally critical, as they can shape the development and deployment of AGI through policy, regulation, and investment.

How Governments are Responding to the New AGI Era

As AGI inches ever closer to reality, governments around the world recognize AI as the next engine of economic growth and prosperity. Eager to improve their competitive standing, many countries are moving aggressively to become a global leader in AI research and development. The countries that pull ahead in this race will capitalize on their newfound competitive advantage. Their efforts include:

- **Seeding Early Development.** Many countries are investing heavily in AGI-linked research, often through their defense budgets or grants to national universities. For example, China's Beijing Institute for General Artificial Intelligence is recruiting top global talent in brain cognition and neuroscience to develop AGI.

- **Providing Access to Key Resources.** To build their AI capabilities, many countries are working to attract top talent and the necessary resources. In early 2024, the United States launched a "National AI Talent Surge" to hire experts in AI development and governance quickly. For its part, the United Arab Emirates bought thousands of NVIDIA computer chips to support its AI sector, which includes a project to build an open-source Large Language Model (LLM) called Falcon. Nations like the UK and the US are also establishing national quantum research labs—and even whole R&D clusters—to build the infrastructure for the next generation of AI, which will rely on quantum computing.

- **Building a Conducive AI Ecosystem.** Many countries are working to build a robust AI ecosystem. This involves not only adopting early AI use cases such as smart-city initiatives, but also fostering an environment where innovation can thrive. Governments are also providing resources to help startups develop and commercialize the AI technologies that lay the groundwork for AGI. They are also focusing on education: Schools are adding AI lessons to their curriculum, and governments are funding media campaigns to increase AI literacy, address public concerns, and prepare citizens for a future where AI plays a greater role.

These initiatives are crucial in creating a society that is both accepting of, and prepared for, the profound changes AGI could bring in the future. For example, Singapore's AI Trailblazers 2.0, a joint initiative involving government agencies and tech giants like Google, aims to equip organizations with the tools to develop AI solutions that tackle real-world challenges. By promoting AI literacy and innovation, countries are not only advancing their current AI capabilities; they are paving the way for the responsible development of AGI.

The Timeline for AGI

While AGI development continues apace, the AI community remains divided over how long it will take before AGI becomes a reality and achieves what is known as the "technological singularity". This is the point where ordinary human intelligence is enhanced—or overtaken—by artificial intelligence in ways that fundamentally transform society. The futurist Ray Kurzweil has been the most optimistic, predicting in his seminal 2005 book The Singularity Is Near that AGI could be achieved by 2029. His 2024 book, The Singularity Is Nearer, reaffirmed his belief in the rapid pace of technological progress. Other experts like Musk and Altman share this optimism, believing that AGI could emerge in the next few years.

But skeptics like Turing Award winners Yann LeCun and Judea Pearl argue that AGI may struggle to achieve the exponential progress needed. In the middle are the many experts who believe AGI is possible—but might not occur until the second half of this century. For instance, a 2018 survey from Oxford University's Future of Humanity Institute found that AI researchers believe there is just a 50% chance that AI could outperform humans in all tasks by around 2063. This ongoing debate reflects the challenges in predicting when AGI will become a reality.

What Could the Future Look Like with AGI?

The development of true AGI could bring about a transformative era for humanity, generating potentially game-changing breakthroughs in many areas. In healthcare, AGI's unique ability to analyze complex biological data could lead to innovative treatments to treat diseases. AGI systems could also help us detect, prepare for, and respond to natural disasters before they hit.

AGI could also give rise to a new era of cognitive cities that "think" and then act on behalf of their citizens. These cities could optimize complex supply chains, infrastructure, and self-driving transportation networks in real time—making life easier and more sustainable for residents and visitors. Beyond our planet, AGI could also lead a new age of space exploration, directing missions into uncharted areas and expanding humanity's presence in the cosmos.

Moreover, AGI could transform how we live and work. AI-powered personal assistants will intuitively understand and cater to the needs of individuals. At the office, AGI could act as a digital partner, helping us work more creatively, productively, and efficiently than we ever imagined. Outside of work, the development of immersive virtual realms could usher in a new era of entertainment and leisure—blurring the lines between the real and the virtual.

Navigating the Journey to AGI

The pursuit of AGI is not just about overcoming formidable technological challenges; it's also a journey that will force us to establish ethical frameworks, find the right balance between humans and machines, and in some instances, reconsider our core beliefs. While AGI offers limitless potential, its unpredictable nature creates the risk of unintended consequences. This journey will require us to address fundamental questions about consciousness, the nature of intelligence, and our evolving role in an AI-driven world.

Ultimately, the value of this journey extends far beyond simply reaching the goal of AGI. It provides an opportunity to unlock countless innovations and use cases that can benefit humanity along the way. As our curiosity helps us push the boundaries of what is possible, the philosophical and moral choices we make will determine how mankind interacts with machines. By approaching these challenges thoughtfully and collaboratively, we have the potential to create a future where AGI not just mimics—but broadens human capabilities.

About the Authors

Dr. Akram Awad is a Partner at the firm's Middle East offices. He is a Global Lead for Smart Cities and core member of the Center for Digital Government. You may contact him by email at Awad.Akram@bcg.com.

Dr. Lars Littig is a Managing Director & Partner at the firm's Middle East offices. He is the EMESA Leader of BCG's Center for Digital Government. You may contact him by email at Littig.Lars@bcg.com.

Sophie Geraerdts is a Project Leader at the firm's Middle East offices. She is a core member of BCG's Technology and Digital Advantage practice.

BCG