McKinsey & Company

McKinsey Explainers

What is artificial general intelligence (AGI)?

Artificial general intelligence (AGI) is a theoretical AI system with capabilities that rival those of a human. Many researchers believe we are still decades, if not centuries, away from achieving AGI.

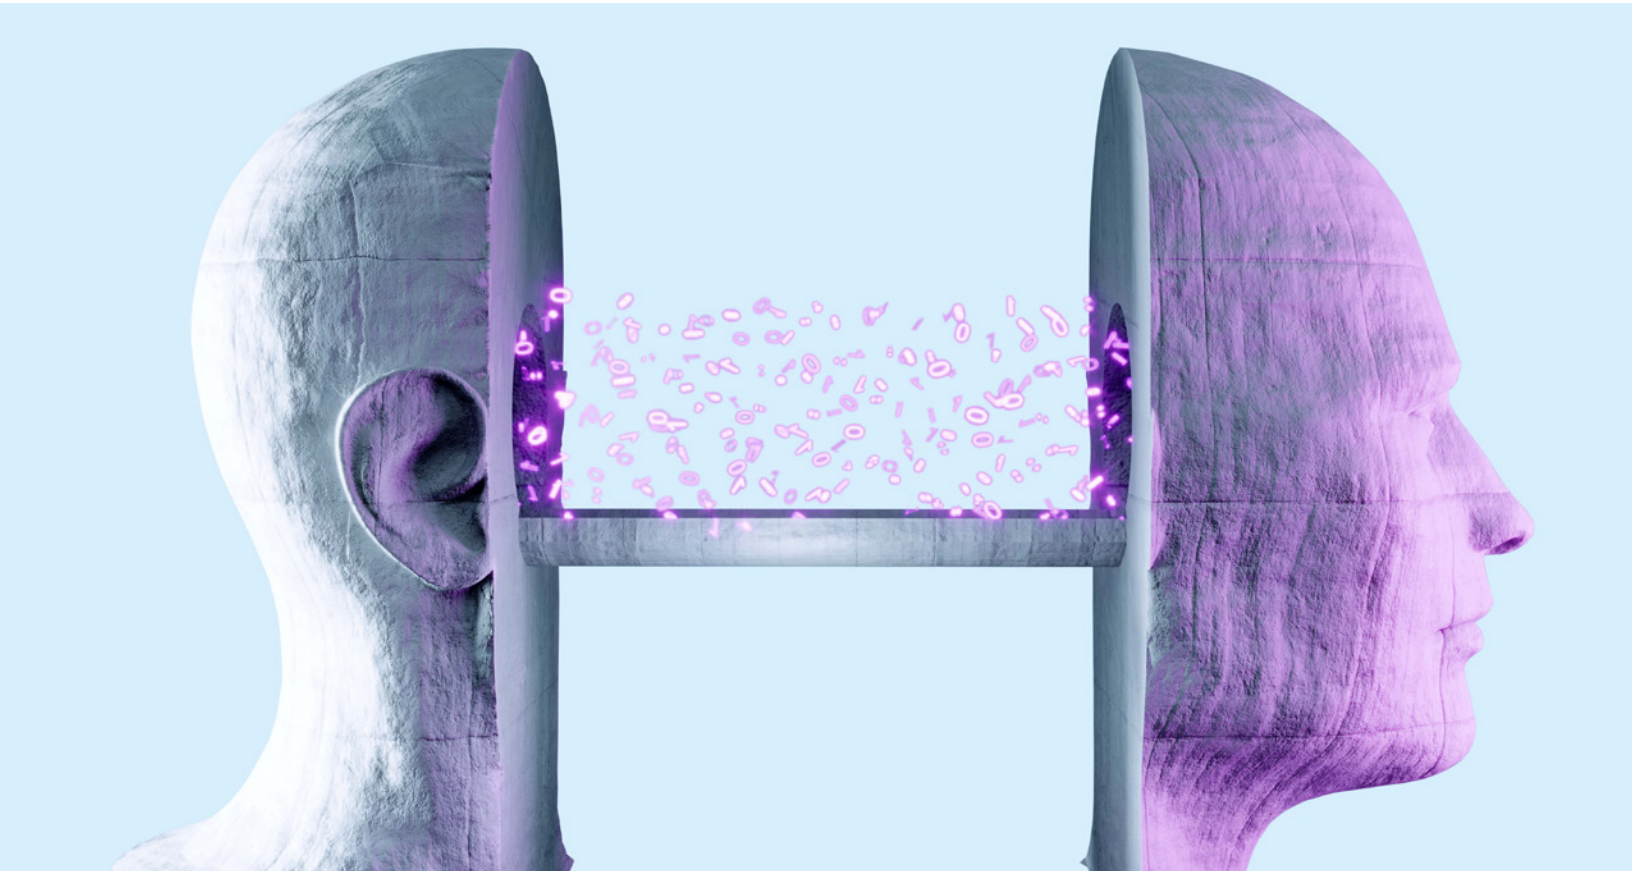

March 2024

You've read the think pieces. AI—in particular, the generative AI (gen AI) breakthroughs achieved in the past year or so—is poised to revolutionize not just the way we create content but the very makeup of our economies and societies as a whole. But although gen AI tools such as ChatGPT may seem like a great leap forward, in reality they are just a step in the direction of an even greater breakthrough: artificial general intelligence, or AGI.

AGI is AI with capabilities that rival those of a human. While purely theoretical at this stage, someday AGI may replicate human-like cognitive abilities including reasoning, problem solving, perception, learning, and language comprehension. When AI's abilities are indistinguishable from those of a human, it will have passed what is known as the Turing test, first proposed by 20th-century computer scientist Alan Turing.

But let's not get ahead of ourselves. AI has made significant strides in recent years, but no AI tool to date has passed the Turing test. We're still far from reaching a point where AI tools can understand, communicate, and act with the same nuance and sensitivity of a human—and, critically, understand the meaning behind it. Most researchers and academics believe we are decades away from realizing AGI; a few even predict we won't see AGI this century (or ever). Rodney Brooks, a roboticist at the Massachusetts Institute of Technology and cofounder of iRobot, believes AGI won't arrive until the year 2300.

If you're thinking that AI already seems pretty smart, that's understandable. We've seen gen AI do remarkable things in recent years, from writing code to composing sonnets in seconds. But there's a critical difference between AI and AGI. Although the latest gen AI technologies, including ChatGPT, DALL-E, and others, have been hogging headlines, they are essentially prediction machines—albeit very good ones. In other words, they can predict, with a high degree of accuracy, the answer to a specific prompt because they've been trained on huge amounts of data. This is impressive, but it's not at a human level of performance in terms of creativity, logical reasoning, sensory perception, and other capabilities. By contrast, AGI tools could feature cognitive and emotional abilities (like empathy) indistinguishable from those of a human. Depending on your definition of AGI, they might even be capable of consciously grasping the meaning behind what they're doing.

The timing of AGI's emergence is uncertain. But when it does arrive—and it likely will at some point—it's going to be a very big deal for every aspect of our lives, businesses, and societies. Executives can begin working now to better understand the path to machines achieving human-level intelligence and making the transition to a more automated world.

Learn more about QuantumBlack, AI by McKinsey.

Most researchers and academics believe we are decades away from realizing AGI; a few even predict we won't see AGI this century (or ever).

What is needed for AI to become AGI?

Here are eight capabilities AI needs to master before achieving AGI.

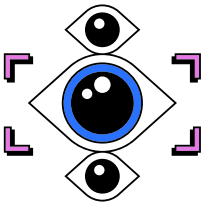

Visual perception

AI systems are a long way from achieving human-like sensory perception. Systems trained through deep learning, for example, still have poor color consistency. Some autonomous cars have been fooled by small pieces of black tape or stickers on a red stop sign, ultimately making the (incorrect and potentially dangerous) decision that the stop sign is something else.

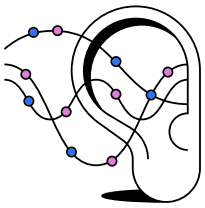

Audio perception

Humans use sound to determine the spatial characteristics of an environment with little to no effort. We can hear background noise and determine the location of the speaker, for instance, whether they're behind us or to our right. AI systems, on the other hand, have a more limited ability to extract and process sound, constrained by their hardware and software. Further, even when they have best-in-class speakers, microphones, and algorithms, AI systems struggle to interpret sound as well as a human.

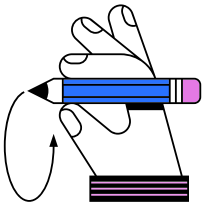

Fine motor skills

AI-powered robots have yet to achieve the kind of fine motor skills that would inspire us to trust them to braid our hair or independently perform surgery on our loved ones. But they're getting closer. In 2019, a single OpenAI robot hand solved a Rubik's Cube in under four minutes. What's more, it was able to continue working even with disruptions it didn't encounter during training, such as being prodded with a stuffed giraffe.

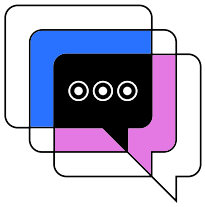

Natural language processing

To rival human-level cognition, AGI would need to consume human sources of information—books, articles, videos, etcetera—with full comprehension. AGI would also need to operate from a place of human-level general knowledge and common sense. When humans communicate, a vast amount of information is assumed and unsaid. AGI would need to fill these gaps. Recent gen AI tools have demonstrated improved natural language processing, but they still lack true understanding and context comprehension. These models rely on statistical patterns and correlations in large data sets to generate text, rather than truly comprehending the meaning and nuances of what's being communicated.

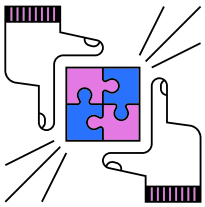

Problem-solving

An AGI system needs to be able to diagnose and address problems—for instance, recognizing that a light bulb is out and replacing it. To do this successfully, an AGI system would need a degree of common sense or the ability to run simulations to determine possibilities, plausibility, and probabilities. AGI systems should also be able to learn from their environment and experiences and adapt to new situations without explicit programming from humans.

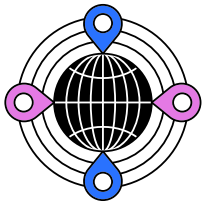

Navigation

GPS, combined with capabilities like simultaneous location and mapping (or SLAM, currently used in self-driving cars and robot vacuums), has made good progress. But years of work are still required to create robot systems that can navigate autonomously with no human priming.

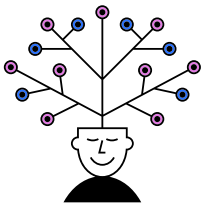

Creativity

Sci-fi fantasies posit that AI will not just reach human-level intelligence but exceed it. For this to happen, AI systems will need to rewrite their own code. This will require them to understand the vast amounts of code humans put together to build them and identify novel ways to improve that code. ChatGPT may be able to write a sonnet, but it's not yet ready to rival human-level creativity.

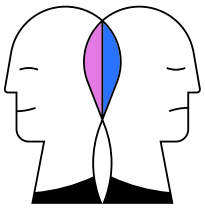

Social and emotional engagement

For robots and AI systems to be truly successful, humans should *want* to interact with them. The robot would need to be able to interpret facial expressions and changes in tone that reveal underlying emotions. Some AI systems can do this already, to a limited extent. Some systems in call centers, for instance, can tell a human is in distress when they raise their voice. But humans themselves struggle to correctly identify emotions; AI that is capable of empathy is still a distant prospect.

How will people access AGI tools?

Today, most people engage with AI in the same ways they've accessed digital power for years: via 2D screens such as laptops, smartphones, and TVs. The future will probably look a lot different. Some of the brightest minds (and biggest budgets) in tech are devoting themselves to figuring out how we'll access AI (and possibly AGI) in the future. One example you're likely familiar with is augmented reality and virtual reality headsets, through which users experience an immersive virtual world. Another example would be humans accessing the AI world through implanted neurons in the brain. This might sound like something out of a sci-fi novel, but it's not. In January 2024, Neuralink implanted a chip in a human brain, with the goal of allowing the human to control a phone or computer purely by thought.

A final mode of interaction with AI seems ripped from sci-fi as well: robots. These can take the form of mechanized limbs connected to humans or machine bases or even programmed humanoid robots.

What is a robot and what types of robots are there?

The simplest definition of a robot is a machine that can perform tasks on its own or with minimal assistance from humans. The most sophisticated robots can also interact with their surroundings.

Programmable robots have been operational since the 1950s. McKinsey estimates that 3.5 million robots are currently in use, with 550,000 more deployed every year. But while programmable robots are more commonplace than ever in the workforce, they have a long way to go before they outnumber their human counterparts. The Republic of Korea, home to the world's highest density of robots, still employs 100 times as many humans as robots.

But as hardware and software limitations become increasingly surmountable, companies that manufacture robots are beginning to program units with new AI tools and techniques. These dramatically improve robots' ability to perform tasks typically handled by humans, including walking, sensing, communicating, and manipulating objects. In May 2023, Sanctuary AI, for example, launched Phoenix, a bipedal humanoid robot that stands 5' 7" tall, lifts objects weighing as much as 55 pounds, and travels three miles per hour—not to mention it also folds clothes, stocks shelves, and works a register.

As we edge closer to AGI, we can expect increasingly sophisticated AI tools and techniques to be programmed into robots of all kinds. Here are a few categories of robots that are currently operational:

— *Stand-alone autonomous industrial robots:* Equipped with sensors and computer systems to navigate their surroundings and interact with other machines, these robots are critical

components of the modern automated manufacturing industry.

— *Collaborative robots:* Also known as cobots, these robots are specifically engineered to operate in collaboration with humans in a shared environment. Their primary purpose is to alleviate repetitive or hazardous tasks. These types of robots are already being used in environments such as restaurant kitchens and more.

— *Mobile robots:* Utilizing wheels as their primary means of movement, mobile robots are commonly used for materials handling in warehouses and factories. The military also uses these machines for various purposes, such as reconnaissance and bomb disposal.

— *Human–hybrid robots:* These robots have both human and robotic features. This could include a robot with an appearance, movement capabilities, or cognition that resemble those of a human, or a human with a robotic limb or even a brain implant.

— *Humanoids or androids:* These robots are designed to emulate the appearance, movement, communicative abilities, and emotions of humans while continuously enhancing their cognitive capabilities via deep learning models. In other words, humanoid robots will think like a human, move like a human, and look like a human.

Learn more about QuantumBlack, AI by McKinsey.

What advances could speed up the development of AGI?

Advances in algorithms, computing, and data have brought about the recent acceleration of AI. We can get a sense of what the future may hold by looking at these three capabilities:

— *Algorithmic advances and new robotics approaches.* We may need entirely new approaches to algorithms and robots to achieve AGI. One way researchers are thinking about this is by exploring the concept of embodied cognition. The idea is that robots will need to learn very quickly from their environments through a multitude of senses, just like humans do when they're very young. Similarly, to develop cognition in the same way humans do, robots will need to experience the physical world like we do (because we've designed our spaces based on how our bodies and minds work).

The latest AI-based robot systems are using gen AI technologies including large language models (LLMs) and large behavior models (LBMs). LLMs give robots advanced natural-language-processing capabilities like what we've seen with generative AI models and other LLM-enabled tools. LBMs allow robots to emulate human actions and movements. These models are created by training AI on large data sets of observed human actions and movements. Ultimately, these models could allow robots to perform a wide range of activities with limited task-specific training.

A real advance would be to develop new AI systems that start out with a certain level of built-in knowledge, just like a baby fawn knows how to stand and feed without being taught. It's possible that the recent success of deep-learning-based AI systems may have drawn research attention away from the more fundamental cognitive work required to make progress toward AGI.

— *Computing advancements.* Graphics processing units (GPUs) have made the major AI advances of the past few years possible. Here's why. For one, GPUs are designed to handle multiple tasks related to visual data simultaneously, including rendering images, videos, and graphics-related computations. Their efficiency at handling massive amounts of visual data makes them useful in training complex neural networks. They also have a high memory bandwidth, meaning faster data transfer. Before AGI can be achieved, similar significant advancements will need to be made in computing infrastructure. Quantum computing is touted as one way of achieving this. However, today's quantum computers, while

powerful, aren't yet ready for everyday applications. But once they are, they could play a role in the achievement of AGI.

— *Growth in data volume and new sources of data.* Some experts believe 5G mobile infrastructure could bring about a significant increase in data. That's because the technology could power a surge in connected devices, or the Internet of Things. But, for a variety of reasons, we think most of the benefits of 5G have already appeared. For AGI to be achieved, there will need to be another catalyst for a huge increase in data volume.

  New robotics approaches could yield new sources of training data. Placing human-like robots among us could allow companies to mine large sets of data that mimic our own senses to help the robots train themselves. Advanced self-driving cars are one example: data is being collected from cars that are already on the roads, so these vehicles are acting as a training set for future self-driving cars.

What can executives do about AGI?

AGI is still decades away, at the very least. But AI is here to stay—and it is advancing extremely quickly. Smart leaders can think about how to respond to the real progress that's happening, as well as how to prepare for the automated future. Here are a few things to consider:

— *Stay informed about developments in AI and AGI.* Connect with start-ups and develop a framework for tracking progress in AGI that is relevant to your business. Also, start to think about the right governance, conditions, and boundaries for success within your business and communities.

— *Invest in AI now.* "The cost of doing nothing," says McKinsey senior partner Nicolai Müller, "is just too high because everybody has this at the top of their agenda. I think it's the one topic that every management board has looked into, that every CEO has explored across all regions and industries." The organizations that get it right now will be poised to win in the coming era.

— *Continue to place humans at the center.* Invest in human–machine interfaces, or "human in the loop" technologies that augment human intelligence. People at all levels of an organization need training and support to thrive in an increasingly automated world. AI is just the latest tool to help individuals and companies alike boost their efficiency.

— *Consider the ethical and security implications.* This should include addressing cybersecurity, data privacy, and algorithm bias.

— *Build a strong foundation of data, talent, and capabilities.* AI runs on data; having a strong foundation of high-quality data is critical to its success.

— *Organize your workers for new economies of scale and skill.* Yesterday's rigid organizational structures and operating models aren't suited to the reality of rapidly advancing AI. One way to address this is by instituting flow-to-the-work models, where people can move seamlessly between initiatives and groups.

— *Place small bets to preserve strategic options in areas of your business that are exposed to AI developments.* For example, consider investing in technology firms that are pursuing ambitious AI research and development projects in your industry. Not all these bets will necessarily pay off, but they could help hedge some of the existential risk your business may face in the future.

Learn more about QuantumBlack, AI by McKinsey. And check out AI-related job opportunities if you're interested in working at McKinsey.

Articles referenced:

— "Generative AI in operations: Capturing the value," January 3, 2024, Marie El Hoyek and Nicolai Müller

— "The economic potential of generative AI: The next productivity frontier," June 14, 2023, Michael Chui, Eric Hazan, Roger Roberts, Alex Singla, Kate Smaje, Alex Sukharevsky, Lareina Yee, and Rodney Zemmel

— "What every CEO should know about generative AI," May 12, 2023, Michael Chui, Roger Roberts, Tanya Rodchenko, Alex Singla, Alex Sukharevsky, Lareina Yee, and Delphine Zurkiya

— "An executive primer on artificial general intelligence," April 29, 2020, Federico Berruti, Pieter Nel, and Rob Whiteman

— "Notes from the AI frontier: Applications and value of deep learning," April 17, 2018, Michael Chui, James Manyika, Mehdi Miremadi, Nicolaus Henke, Rita Chung, Pieter Nel, and Sankalp Malhotra

— "Augmented and virtual reality: The promise and peril of immersive technologies," October 3, 2017, Stefan Hall and Ryo Takahashi

Get to know and directly engage with senior McKinsey experts on AGI

Aamer Baig is a senior partner in McKinsey's Chicago office; **Federico Berruti** is a partner in the Toronto office; **Ben Ellencweig** is a senior partner in the Stamford, Connecticut, office; **Damian Lewandowski** is a consultant in the Miami office; **Roger Roberts** is a partner in the Bay Area office, where **Lareina Yee** is a senior partner; **Alex Singla** is a senior partner in the Chicago office and the global leader of QuantumBlack, AI by McKinsey; **Kate Smaje** and **Alex Sukharevsky** are senior partners in the London office; **Jonathan Tilley** is a partner in the Southern California office; and **Rodney Zemmel** is a senior partner in the New York office.

Designed by McKinsey Global Publishing