



ARTIFICIAL INTELLIGENCE

What CEOs Need to Know About Artificial General Intelligence

By Greg Emerson and Ulrich Pidun

ARTICLE SEPTEMBER 10, 2026 12 MIN READ

As AI models continue to make dramatic leaps in sophistication, a question looms large for CEOs. Is the world on a trajectory toward artificial general intelligence (AGI)—the point at which the technology achieves virtuoso-level intelligence across the full breadth of cognitive human capabilities? Some AI lab leaders have predicted that AGI is a near-term possibility, which would signal the end of most “knowledge work” jobs. Other AI leaders are highly skeptical of this timeline and have suggested that today’s large language model (LLM) architecture will not lead to AGI. For company leaders, this divergence of opinion amplifies the confusion.

Adding to leaders’ challenge is the sheer magnitude of AGI’s implications for the global economy and society at large—implications that are impossible to fully anticipate. It is clear, however, that past technologies automated just a portion of tasks, leaving humans a comparative advantage in other areas. A system that could match or exceed the best human at everything cognitive would eliminate this advantage, along with jobs, wages, and commensurate spending.


In response, CEOs today should focus on a plausible and actionable scenario in which, as LLMs scale, AI’s “jagged frontier” becomes more pronounced. First identified in a 2023 study by researchers at the BCG Institute working alongside academics from leading institutions, the jagged frontier is an indication of AI’s uneven capabilities: it is a virtuoso at some skills, above average at others, and below average at others still. This jaggedness exists across many different jobs today and is precisely why so many reports have shown that AI is not driving mass unemployment. Many jobs involve assigning humans to oversee AI where it excels and spend the bulk of their time working on tasks where AI remains deficient. (See Exhibit 1.)

EXHIBIT 1


Task-Level Capabilities Reveal the Jagged Frontier of AI

A snapshot of task-level AI profiles, mid-2026

Solo/small-firm litigation attorney
Task allocation approximated within a measured billable envelope



Primary care physician—ambulatory and board-certified
Task allocation approximated with measured face-time and EHR blocks



Sources: B. G. Arndt et al., “Tethered to the EHR: Primary Care Physician Workload Assessment Using EHR Event Log Data and Time-Motion Observations.” *Annals of Family Medicine*, 2017; D. Charlotin, AI Hallucination Cases database, accessed 2026; Clio Legal Trends Report, 2023–2025; M. Dahl et al., “Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models.” *Journal of Legal Analysis*, 2024; M. R. Morris et al., “Levels of AGI for Operationalizing Progress on the Path to AGI.” *International Conference on Machine Learning*, 2023; J. Pereyra, “Extending Legal Agent Bench to M&A Due Diligence.” Harvey, 2026; N. Grupen, G. Pereyra, & J. Pereyra, “Initial Results on Legal Agent Benchmark (LAB).” Harvey, 2026; Microsoft “Time Warp” developer survey, 2025; C. Sinsky et al., “Allocation of Physician Time in Ambulatory Practice: A Time and Motion Study in 4 Specialties.” *Annals of Internal Medicine*, 2016; T. Tu et al., “Towards Conversational Diagnostic Artificial Intelligence.” *Nature*, 2025; V. Liévin et al., “Towards Conversational Artificial Intelligence for Disease Management.” *Nature*, 2026; BCG Institute analysis.

Note: Charts are illustrative. Model ability to complete tasks is AI generated with expert review. Actual performance will vary across specialization and tenure. APP = advanced practice provider; CME = continuing medical education; EHR = electronic health record; QI = quality improvement.

Debating the Future


Currently, no industry consensus exists on how to define AGI. In this article, however, we define it as technology that achieves virtuoso-level depth and competence across the full breadth of human capabilities, which is why it's characterized as *general intelligence*. Another term sometimes used is *superintelligence*. This technology does not exist today.

Consider software engineering, for example. It sits at the most advanced point of applied AI, yet AI still can't perform all aspects of a software engineer's job. This was a key argument in a recent BCG Institute study on AI's impact on the labor market. Software engineering is not cleanly divisible: AI accelerates code generation and testing, but it can't yet own system-level judgment, architecture tradeoffs, management of stakeholder tradeoffs, and accountability for the outcome. Over the past two years, AI models have made extraordinary progress in code writing, and engineers have shifted their focus to areas of the job that AI can't perform. (See Exhibit 2.) AI capabilities have become more jagged, humans have become more valuable, and employment of software engineers continues to rise.

EXHIBIT 2

From 2024 to 2026, AI's Grasp of Software Engineer Skills Has Evolved Dramatically

Software engineer—midlevel, back-end
Capacity shift over two years; task allocation approximated and held constant




Sources: ARC-AGI-3 interactive reasoning benchmark, ARC Prize, 2026; ActiveState Developer Survey, 2019; C. E. Jimenez et al., "SWE-bench: Can Language Models Resolve Real-World GitHub Issues?" *International Conference on Learning Representations*, 2024; T. Kwa et al., "Measuring AI Ability to Complete Long Software Tasks," *Model Evaluation and Threat Research*, 2025; A. N. Meyer et al., "The Work Life of Developers: Activities, Switches and Perceived Productivity," *IEEE TSE*, 2017; M. R. Morris et al., "Levels of AGI for Operationalizing Progress on the Path to AGI," *International Conference on Machine Learning*, 2023; A. Sergeyuk et al., "Evolving with AI: A Longitudinal Analysis of Developer Logs," *International Conference on Software Engineering*, 2026; SWE-bench Pro (Scale AI), Scale Labs, accessed 2026; Tidelfit/The New Stack, 2019; T. Xie et al., "OSWorld: Benchmarking Multimodal Agents for Open-Ended Tasks in Real Computer Environments," *Neural Information Processing Systems*, 2024; B. Yu et al., "SWE-ABS: Adversarial Benchmark Strengthening Exposes Inflated Success Rates on Test-based Benchmark," *International Conference on Machine Learning*, 2026; L. Zhang et al., "SWE-bench Goes Live!", *Neural Information Processing Systems*, 2025; BCG Institute analysis.

Note: Charts are illustrative. Model ability to complete tasks is AI generated with expert review. Actual performance will vary across specialization and tenure. CI/CD = continuous integration and continuous delivery.

The key question is whether this jaggedness will persist. Scaling laws assert that additional computing power and training data reliably improve model performance. But will the next decade of LLM models close the jaggedness gap? Or is jaggedness a structural property of the way LLMs acquire capability? If the latter view is correct, achieving AGI will require a completely different AI architecture with an uncertain timeline. This debate over scaling versus architecture divides many leading AI thinkers. (See Exhibit 3.)

EXHIBIT 3
AI Leaders Don't Agree on Whether Scaling LLMs Will Reach AGI



Source: BCG Institute analysis.
Note: This exhibit is an illustrative summary of public positions and is not intended as a survey or consensus map. In the exhibit, we define AGI as a technology that has achieved virtuoso-level depth and competence across the full breadth of human capabilities.

Members of the AI research community skew even more conservatively than AI lab leaders do. In a 2025 survey by the Association for the Advancement of Artificial Intelligence (a group consisting primarily of academics), 76% of respondents agreed that scaling up current AI research approaches was “unlikely” or “very unlikely” to reach AGI.

Understanding the Technical Bottlenecks

Lacking a single, expert point of view to embrace, CEOs need to understand where matters stand today from a technical standpoint. A number of significant characteristics within LLM architecture

suggest that jaggedness may be a structural property of LLMs. Noteworthy bottlenecks include data exhaustion, the limitations of reinforcement learning with verifiable rewards as a training method, methodological constraints restricting on-the-job learning; and a practical ceiling on available computing power, energy, or hardware.

Data Exhaustion

Much of the earliest general progress in LLMs came from training on larger and larger data sets. AI researchers consistently observed that the more data they fed a model, the more effectively the model generalized. The challenge is that high-quality, readily accessible human-generated public text is finite. As OpenAI cofounder Ilya Sutskever observed, pretraining in its current form it will end, because there is only one internet.

Reinforcement Learning with Verifiable Rewards

Many recent breakthroughs in LLM progress—particularly in connection with autonomous agents—have come from a technique called reinforcement learning with verifiable rewards (RLVR). Essentially, RLVR involves having AI researchers ask AI to do a task repeatedly and then reward it if it performs the task correctly. RLVR has produced large gains in model performance in situations where it is possible to check outcomes cheaply and reliably, especially in mathematics and code.

But much of the knowledge work economy has no immediately verifiable answer. If the question is, “Is this the right product to launch?” the way to find out is to test it in the market. If the question is “Will this stock investment strategy yield five-year returns above the S&P?” the answer is unknowable until the five years expire. Extending RLVR techniques to ambiguous, long-horizon work is hard because rewards are sparse, delayed, and contestable. In short, although RLVR makes AI extremely good at some parts of the job, it may not be applicable to many skills.

On-the-Job Learning

Every leader understands the importance of on-the-job learning. Even the most brilliant and seasoned executive needs time to understand an organization’s history, culture, and strategic and operational priorities. Instead of coming from formal documentation, much of this learning is absorbed through interactions with coworkers and customers over time.

Although systems can access new organizational information through documented context, retrieval, tools, and memory, today's LLMs generally do not update while handling a task. In other words, they do not "learn on the job." The unresolved question is whether they can autonomously build durable, reliable tacit knowledge from experience, as employees can. The constraints on their capacity to learn limit these models' ability to perform many of the activities expected of employees, in situations where organizational context is critical to judgment and decision making.

A Ceiling on Computing Power, Energy, or Hardware

One core observed pattern in AI research is that throwing more computing power at a problem reliably leads to a better model. In theory, supplying sufficient computing power could lead to a model that achieves virtuoso-level expertise at every skill. But this prospect is subject to real physical limitations. It may simply not be possible to allocate enough power generation, grid capacity, and semiconductors to a training run that might yield AGI.

Regulatory Regimes

Even if LLMs can lead to AGI, and researchers clear the various technical hurdles and build AGI, that is no guarantee that it AGI will ever become widely available. One reason is that governments might choose to gate access on national-security grounds, capping the technology's economic, business, and societal impact regardless of its underlying capability. In this scenario, access to full-force AGI would be limited to a handful of government bodies and guardrailed commercial use cases, and a more jagged AI would remain dominant in most of the world.

The potential for regulatory impact is not merely hypothetical, as several recent examples illustrate:

- On June 12, 2026, the US government ordered Anthropic to suspend all access to Mythos 5 and Fable 5 for foreign nationals, citing the risk of cyberattacks. Although the government subsequently reversed this prohibition, the episode demonstrates that the government claims the authority to ban future models if it believes that they may be too dangerous to release.

- On June 25, 2026, news outlets reported that the White House had asked OpenAI to preemptively restrict release of GPT-5.6 (Sol/Terra/Luna) before its public launch. This marked the first time that the US government asked a company to restrict a model before releasing it. In response, OpenAI initially released its flagship (Sol) to only about 20 government-vetted partners.
- On July 14, 2026, Demis Hassabis, the co-founder and chair of Google Deep Mind, published a manifesto titled “A Framework for Frontier AI and the Dawning of a New Age,” proposing that the US create a standards body for frontier AI that would have the power to approve new models prior to release to US with American users.
- On July 28, 2026, a statement titled “Pacing the Frontier” (posted at pacingthefrontier.com) called on the US government to help build tools to deliberately slow frontier AI. More than 1,300 employees of OpenAI, Anthropic, Google DeepMind, and Meta have signed this letter.

Of course, there are many examples of regulators building distinct swim lanes for government and commercial availability. For example, Boeing maintains separate divisions for serving commercial airlines and the defense industry, and some of the company’s most cutting-edge innovation in aviation reserved solely for government access.

What CEOs Can Do Today

Given the level of uncertainty about whether AGI will be achieved and what disruption it would create, company leaders cannot really prepare for true AGI—a form of AI that achieves virtuoso levels across all cognitive skills and rapidly absorbs all context and all tacit knowledge within most organizations. Such technology would effectively render knowledge work obsolete, eliminating competitive advantage and economic value in vast sectors of the economy. If these capabilities were further translated into robotics and physical AI (with LLMs acting as the world’s greatest robotics designer), the level of disruption would compound even further.

The idea that it will be possible to find meaningful work for humans to do rests on the assumption that AI’s capability frontier, like that of all past technologies, will remain jagged. But if, instead, AI generalizes, there will be no more cognitive work for humans to do. Even businesses that possess strong moats around cognitive work—in physical goods or human-centric services, for example—will struggle to adapt to an economic shock that causes a meaningful portion of their customers or their customers’ customers to lose their jobs, with AI readily replicating proprietary ways of working. No level of strategic planning or technology investment can help a CEO prepare for that scenario.

Meanwhile, in a world of sustained and increasingly pronounced jaggedness, competitive advantage comes from figuring out how AI can best augment the humans in a company’s teams. The BCG Institute describes this as “token-based competition.” Ultimately, change management, skills development, incentive setting, and culture are what will distinguish companies when all of them have access to the same powerful AI. If jaggedness persists, AI can be a powerful engine for growth and human prosperity without driving mass unemployment and dislocation.

But what is a CEO to do when faced with such vast uncertainty regarding these distinct outcomes? A French mathematician named Blaise Pascal proposed a solution to this question nearly 400 years ago. His analysis, known as *Pascal’s wager*, concludes that in a scenario where a person is uncertain of the outcome, where the possible outcomes are vastly different, and where making a choice is unavoidable, the logical way to proceed is to identify and embrace the value-maximizing choice. Applied to the AGI question, the wager can be expressed in the form of a simple table. (See Exhibit 4.)

EXHIBIT 4
Reinventing Pascal’s Wager for the 21st-Century CEO

CEO options regarding persistent jaggedness	The future situation	
	Jaggedness persists	AGI is achieved
Prepare	Competitive advantage is accelerated	Competitive advantage is severely reduced or eliminated
Don’t prepare (reasoning that AGI is coming, and you can’t prepare for it)	Competitive advantage is reduced (as rival firms are relatively prepared)	Competitive advantage is severely reduced or eliminated

Source: BCG Institute analysis.

According to this line of thinking, the correct choice is clear. To the extent possible, leaders should prepare for a future in which AI continues to grow in sophistication but has capability limits. And, critically, they should ensure that their company’s operating model is aligned to today’s jagged frontier and prepared to adapt to tomorrow’s.

The authors would like to acknowledge Dwarkesh Patel, Ethan Mollick, and Erik Brynjolfsson, whose prior work contributed meaningfully to the ideas in this article.

BCG Institute is Boston Consulting Group's real-world think tank. We research the future of thinking about the future. Built on BCG's decades-long intellectual legacy and a global network of academics, scientists, and institutional partners, we bring together expertise across disciplines to advance ideas before they become conventional wisdom. Tracking the forces reshaping business, technology, economics, and geopolitics, we turn emerging patterns into clarity, helping leaders around the world turn early insight into measurable outcomes.

Authors



Greg Emerson

Managing Director & Senior
Partner
San Francisco - Bay Area



Ulrich Pidun

Partner and Director, Corporate
Strategy
Frankfurt



ABOUT BOSTON CONSULTING GROUP

Boston Consulting Group bridges the gap between ambition and outcomes for the world's leading companies and organizations. We are built for this era of unprecedented change — bringing strategic clarity rooted in over 60 years of deep domain knowledge, combined with applied AI shaped by our practitioners. BCG works shoulder-to-shoulder with CEOs across industries and geographies to deliver transformative impact at scale: stronger returns, transferred capabilities, and change that sticks. For more information, visit bcg.com.