



ARTIFICIAL INTELLIGENCE

AI Will Reshape More Jobs Than It Replaces

By [Greg Emerson](#), [Matthew Kropp](#), [Julie Bedard](#), [Lisa Krayer](#), [Viacheslav Romanov](#), [Megan Hsu](#), [Luis Sanchez Boedo](#), and [Diya Mohnot](#)

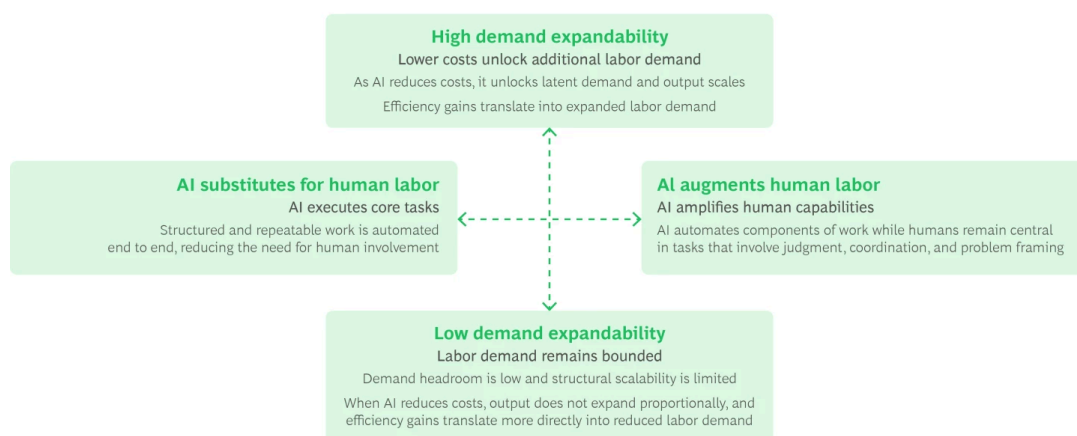
ARTICLE APRIL 03, 2026 15 MIN READ

Over the next two to three years, 50% to 55% of jobs in the US will be reshaped by AI. For many employees, this will mean that they retain the same or a similar role but face radically new expectations for how they work and what they produce. For company leaders, it will require a clear vision for how the transformation is managed, including a scaled, strategic approach to upskilling and reskilling and the restructuring of career ladders.

This shift is already happening—and will pick up speed as AI adoption spreads. Our analysis, based on microeconomic modeling, identified a significant swath of the labor force for which AI will meaningfully augment current roles. Moreover, when the productivity gains from AI use trigger increased end product demand and the potential for augmentation is high, we believe there will be a need for more and, in some cases, new human roles. (See Exhibit 1.) While job augmentation and new-job creation will happen rapidly, full substitution of jobs by AI will be slower. Five years from now—or perhaps further in the future—10% to 15% of jobs in the US could be eliminated.¹ This level of potential job loss is considerable and creates an important call to action for business leaders.

EXHIBIT 1

Two Critical Factors in AI’s Impact Are Labor Substitution and Demand Expandability



Source: BCG Henderson Institute analysis.

Critically, our analysis is not intended to be an unemployment forecast. It does not account for macroeconomic factors such as geopolitics or inflation, nor does it contemplate the impact of new AI breakthroughs beyond the capabilities of today’s frontier models. Moreover, our model can’t solve for powerful and influential unknowns, such as the future impact of AI on the accessibility of jobs and the speed with which the technology is adopted. (See “Methodology: Translating Automation Potential into Labor Market Outcomes” for more on our approach.)

— Methodology: Translating Automation Potential into Labor Market Outcomes

This analysis applied a structured microeconomic framework to assess how AI reshapes employment outcomes. Rather than equating automatable tasks with job loss, we evaluated three distinct forces: task-level automation potential, substitution versus augmentation dynamics, and demand expandability.

AI Task Automation Potential

We began by estimating automation potential at the task level for each role. Using granular occupational data (Revelio Labs' 1,500-role taxonomy, O*NET task decomposition, and Revelio headcount data), we evaluated the share of individual work activities within each role that are automatable given current AI capabilities.

Each task was assessed using a structured rubric and classified as automatable if it met the following criteria:

- The task does not require a significant physical human presence or manual interaction in the real world.
- The task can be executed without substantial emotional intelligence, negotiation, or complex interpersonal judgment.
- The task is sufficiently structured to be performed without excessive ambiguity or open-ended reasoning.
- The necessary data inputs are observable or available to an agentic AI system.
- The task's outcome is governed by rule-based logic grounded in documentation, precedents, or established procedures.

For each role, we applied the share of time spent on each automatable task to estimate the overall task automation potential. Roles with automation potential below the average of 40% were classified as lower automation, as they are less likely to be prioritized for material disruption in the next four to five years. We segmented lower-automation roles into enabled and limited-exposure roles using a 25% automation potential cutoff. Roles exceeding the 40% threshold were evaluated in subsequent stages of the analysis.

Substitution Versus Augmentation

For roles with meaningful automation potential, we assessed whether AI is more likely to substitute for labor or augment it. This depends on the degree of human value embedded in the role, which we evaluated on two dimensions:

- **Human Interaction and Judgment.** We determined whether reliable performance requires significant emotional intelligence, negotiation, or

complex interpersonal judgment. Roles scoring low involve routine, transactional interactions in which persuasion needs are minimal and outcomes hinge on objective criteria requiring limited discretion; these roles are more likely to be substituted. Roles scoring high require nuanced interpretation of emotional and social cues, and outcomes depend on trust, persuasion, and contextual judgment; these are more likely to be augmented.

- **Process Structure and Repeatability.** We evaluated how structured and codifiable the workflow is. Roles scoring low involve structured, repeatable processes with well-defined inputs and outputs; these roles are more susceptible to substitution. Roles scoring high involve open-ended problem solving and frequent exceptions requiring expert judgment; these roles are more likely to be augmented.

Demand Expandability

Even when AI substitutes for humans in executing tasks, labor outcomes depend on whether productivity gains expand total demand. We evaluated demand expandability through two lenses:

- **Demand Headroom.** We evaluated whether AI-driven reductions in unit cost or cycle time unlock additional output. We assessed this on the basis of two empirical signals. First, we estimated industry-level price elasticity by relating changes in industry price indices to changes in real industry output, normalized for overall real GDP growth.* Industries where output is especially responsive to price movements were classified as more elastic, implying that AI-driven cost reductions are likely to translate into higher demand for output.

Second, as a proxy for unmet demand, we used labor market indicators, particularly job openings from the US Bureau of Labor Statistics' Job Openings and Labor Turnover Survey. Openings were normalized for baseline turnover to avoid overstating demand in high-churn occupations. A persistently elevated job opening rate suggests that employers face supply constraints, indicating latent or underserved demand.

Together, these measures help distinguish industries where demand is expandable from those that are saturated.

- **Structural Scalability.** We used this lens to evaluate whether supply can scale if demand increases. Even when demand headroom exists, expansion may be constrained by capital intensity (such as physical capacity, capex

requirements, or long lead times) or by credentialing barriers (including education level, licensing, certification, regulatory requirements, and extended training pipelines). In roles with high structural scalability, demand expansion is more likely to translate into employment growth.

Defining Roles Versus Jobs

We applied our model at the level of roles. We evaluated each role and translated the resulting insights into impacts on the total number of jobs. A role refers to a type of work defined by a specific set of responsibilities, while jobs refer to the positions within a role. In our dataset, approximately 165 million US jobs (the employment level as of January 2026 reported by the US Bureau of Labor Statistics) are distributed across roughly 1,500 distinct roles defined by Revelio Labs. (Revelio aggregates publicly available professional profile data, including LinkedIn profiles, and applies weighting adjustments to account for workers not represented in the data.) Our analysis focuses on the US labor market, where Revelio's data is more comprehensive than elsewhere.

Quantifying Potential Job Impact

To estimate the share of jobs potentially vulnerable due to AI, the model considers:

- **Substituted Roles:** the number of jobs where AI substitutes for human workers multiplied by the role's automation potential
- **Divergent Roles:** the number of jobs where AI substitutes for human workers multiplied by the role's automation potential multiplied by a demand-expansion adjustment (0.5–1.0) reflecting how much employment persists as output grows
- **Aggregation Across Roles:** adding together these adjusted impacts across roles produces a point-in-time estimate of 10% to 15% of US jobs that are vulnerable over the next four to five years

Quantifying Potential Share of Reshaped Jobs

We defined reshaped jobs as those in which AI materially changes how work is done, even if the job itself remains. This includes amplified and rebalanced roles, where AI augments work; remaining roles in substituted and divergent categories, where workers will need significant upskilling; and enabled roles with more than 25% automation potential, where task exposure is high enough to alter day-to-day

responsibilities. Aggregated across roles, this suggests that 50% to 55% of US jobs could be reshaped over the next two to three years.

Scope and Limitations

The analysis reflects current AI capabilities, with a focus on large language models and autonomous driving technologies because of the demonstrated commercial progress in both domains. Other forms of physical AI and robotics were excluded, as large-scale deployment across most industries remains limited. While advances are accelerating, we do not expect diffusion to materially alter labor market outcomes in the near term. The results therefore reflect potential exposure under current AI capabilities; realized outcomes will depend on how quickly organizations deploy these technologies.

If AI systems were to consistently perform open-ended, judgment-intensive tasks autonomously at human-level proficiency, the distribution of roles across the substitution and augmentation categories would shift materially, and the framework would need to be revisited.

This is a microeconomic assessment of labor impact, not a forecast of aggregate unemployment or near-term layoffs. Employment outcomes are also influenced by macroeconomic forces—including inflation, interest rates, trade policy, recession cycles, and geopolitical dynamics—which are outside the scope of this analysis.


That said, the analysis provides clear guidance for CEOs on how to act in the face of this workforce shift. They are making decisions today that affect their people and the viability of their business, and they need a nuanced view of AI's impact. Those who cut their workforce beyond AI's ability to replace it will see productivity drop, institutional knowledge disappear, and critical talent walk away. Those who fail to dramatically rethink work will see their competitors grow faster and more profitably.

Task Automation Doesn't Have to Mean Job Loss

To build our model, we calculated the number of jobs involving tasks that are at least 40% automatable. This represents the average level of automation for all US-based occupations and the threshold after which role and organization redesign becomes a stronger business case. The 43% of jobs that exceed this threshold formed the focus of our analysis. The other 57% are jobs that depend heavily on the physical presence of human workers, on hands-on work, or on sustained human interaction, all of which limit the potential for automation and thus are less likely to be disrupted by AI. (See Exhibit 2.)

EXHIBIT 2

Agentic AI May Drive High Levels of Task Automation in 43% of Jobs



Sources: Revelio Labs; O*NET; US Bureau of Labor Statistics; BCG Henderson Institute analysis.

Substitution Versus Augmentation. To illustrate the difference between substitution and augmentation, consider call center representatives and software engineers, two roles that are already deploying agentic AI at scale.

A call center representative is typically responsible for resolving a defined set of customer inquiries within established workflows. Much of the work involves structured interactions such as account lookups, policy explanations, and scripted troubleshooting. When AI systems can reliably handle these repeatable inquiries end to end, fewer representatives are required. In this setting, the workflow can often be clearly bifurcated, with AI handling first-line interactions and humans focusing on escalations and exceptions.

In some cases, representatives can transition into higher-value roles focused on relationship building and proactive risk mitigation. But overall employment in the call center rep role will decline as the most structured tasks are absorbed by the system.

A software engineer, by contrast, produces a very different output. While coding includes routine elements, the core value of the role lies in system design, architectural judgment, tradeoffs between performance and cost, and the translation of business needs into technical solutions. AI can dramatically accelerate code generation and testing, but given today's capabilities, it cannot replace the system-level judgment required to own the outcome end to end—meaning the work

cannot be cleanly divided between system and engineer. Instead, software development becomes an ongoing interaction in which engineers define objectives, refine outputs, validate results, and integrate components into broader systems.

AI supports and accelerates these steps but does not replace the need for human judgment and accountability. As a result, AI helps engineers do their jobs more effectively rather than replacing them, making these roles less likely to experience direct displacement and shifting engineers' work toward system-level thinking, orchestration, and product and design tasks, rather than repetitive coding tasks.

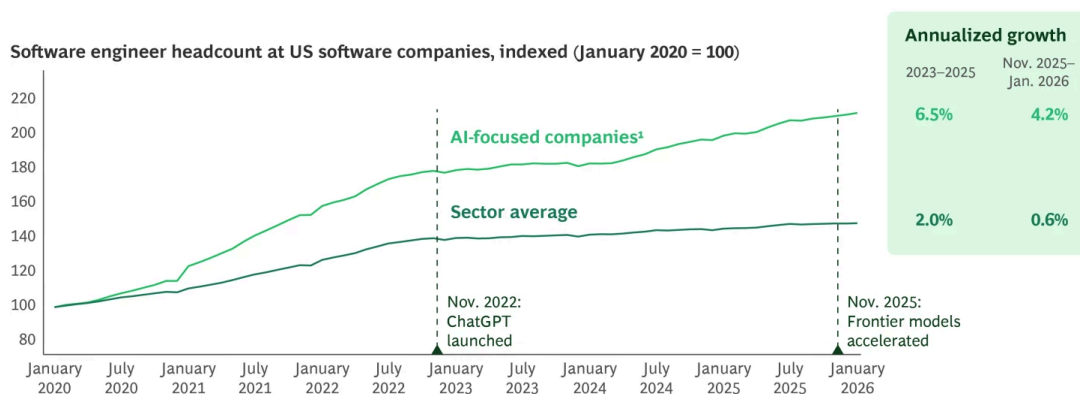
Low Versus High Demand Expandability. When AI lowers the cost of delivering a business outcome or end product, the question is whether demand for the output expands or remains bounded. If lower costs unlock unmet demand, increase accessibility, or accelerate consumption, total output may rise and employment may remain stable or even grow despite significant automation at the task level. If demand for the outcome is fixed, productivity gains are more likely to translate into fewer required workers.

This dynamic is not new. Economists have long observed that efficiency improvements can increase total consumption rather than reduce it, a phenomenon often referred to as Jevons Paradox. When the cost of a resource falls, usage can rise. The same logic applies to labor: whether productivity reduces employment depends on how demand for the output responds.

Software engineering illustrates expandable demand. Organizations across industries continue to face persistent unmet need for digital products, automation, and new features. As AI reduces the cost and time required to build software, organizations often build more. Output expands and overall job volume may remain stable or grow, even as the productivity of individual engineers rises, because humans continue to play a meaningful role. The continued growth of software engineering headcount in the years following the introduction of ChatGPT in 2022 illustrates this phenomenon. (See Exhibit 3.)

EXHIBIT 3

Software Engineering Headcount Has Risen Steadily in the Past Three Years



Sources: Revelio Labs; BCG Henderson Institute analysis.
¹Approximately 500 US-based software companies tagged as AI in PitchBook.

Call center representatives illustrate bounded demand. The volume of inbound interactions is largely determined by the size of the customer base and the frequency of service needs. When AI reduces the cost of handling routine inquiries, the number of interactions does not expand proportionally. In this context, productivity gains are more likely to reduce the number of representatives required.

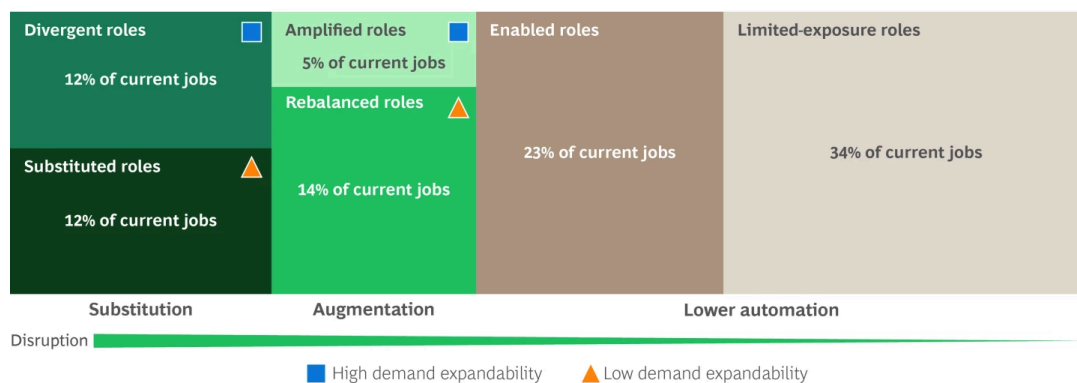
The Majority of Current Jobs Will Stay but Evolve

By analyzing roles against the factors above, we can identify where they belong in BCG Henderson Institute’s proprietary AI Labor Disruption Segments, which comprise six categories. (See Exhibit 4.)

EXHIBIT 4

In the US, Most Jobs Will Remain in Some Form in the Coming Years

AI Labor Disruption Segments (US workforce)



Sources: Revelio Labs; O*NET; US Bureau of Labor Statistics; BCG Henderson Institute analysis.
 Note: Demand expandability was not assessed for lower-automation roles, because the expected disruption is less significant.

Amplified Roles. When AI augments human capabilities and demand expands, employment may remain stable or grow. Humans remain central to value creation, and wage inflation may develop as higher productivity increases competition for skilled talent. This type of role represents 5% of current jobs.

Software engineering falls under this category. (For a “what if” scenario, see “What Would It Take for Software Engineering to Move from an Amplified to a Divergent Role?”) So do many lawyers, particularly those in advisory and judgment-intensive areas of law. Investment in legal tech startups, like Harvey AI, reached record-high levels in 2025, spurring significant discussion of how

AI will reshape this field. As AI accelerates research, drafting, and case preparation, legal services could become more accessible. In domains where unmet demand exists, such as compliance, regulatory review, and contract management, lower costs may increase the volume of legal work. Senior legal judgment would remain central, with AI amplifying professional capabilities.

— What Would It Take for Software Engineering to Move from an Amplified to a Divergent Role?

The future of software engineering is an incredibly controversial topic. Some AI leaders have publicly claimed that there's never been a better time to be a software engineer, while others have forecast the end of the profession. The expandability of demand around the output of software engineers is without question: the number of IT roadmaps is nearly infinite and AI implementation itself will be a massive driver of more engineering output. The question for the medium and long term is whether to categorize the software engineering role as amplified (suggesting that AI will augment the work of software engineers and most jobs will remain) or divergent (suggesting that AI will substitute for that work and a substantial share of jobs may potentially be eliminated).

At present, although AI has largely mastered the domain of coding, humans are still required in many parts of the software engineering process, including system design, architectural tradeoffs, quality checking for security and efficiency, and integration across complex systems. There is an open question—one that cannot be answered today—as to whether newer AI models will master these processes and come to serve as complete substitutes for software engineers.

Although software engineering is currently situated in the amplified section of our AI Labor Disruption Segments, we expect considerable churn given the extent to which the frontier labs are focused on this domain. Skills such as the writing and maintenance of code will be deprioritized, whereas higher-order systems thinking and proficiency with AI tools will grow in importance. Were the role to move to the divergent segment owing to a leap in the capabilities of AI agents, we would expect to see considerably more output from a smaller population of senior, deeply knowledgeable engineering leaders.

Leaders should monitor how these technological shifts interact with workforce demand and adjust their talent strategies accordingly. Rather than treating workforce planning as a static forecast, companies will increasingly need to update hiring, upskilling, reskilling, and role design dynamically as AI evolves.

Rebalanced Roles. When AI augments work but demand is bounded, headcount may remain steady while roles are redesigned. Routine tasks automate while more complex responsibilities expand. As work shifts toward higher-value activities, skill requirements rise and upskilling becomes essential. We estimate that 14% of current jobs fall under this category.

One example is content marketing, where demand is bounded by marketing budgets and strategic priorities. At the same time, with the fragmentation of audiences and with influencers and large language models reshaping customer journeys, brands need more, not less, content delivered in a more targeted way to the right audience at the right time. Marketer roles—instead of being broken up by channel, as they are now—will expand, with marketers and creative directors becoming omnichannel specialists who can think through end-to-end campaigns. Academic research is a similar case. Demand is constrained by institutional funding, but with increased productivity driven by agentic AI, the need increases for hypothesis formation, data interpretation, and scholarly judgment—which remain the province of humans.

Divergent Roles. Where AI substitutes for human tasks but demand remains expandable, the effect on employment becomes uneven; our analysis suggests that this applies to 12% of current jobs. In these roles, entry-level and junior positions are more exposed to automation in the short term. A sizable portion of the structured tasks traditionally performed at these levels are among the first to be automated, which means there will be some roles lost in this category. At the same time, responsibilities at senior levels persist and may grow to include higher volumes of work as demand expands. This creates a structural tension. The roles that remain require contextual judgment, oversight, and coordination skills that are typically built through on-the-job experience. This is a clear example of a scenario where the numbers alone can't tell the full story.

Insurance sales agents fall into this category. AI automates routine activities such as lead qualification, quote generation, and policy comparisons—tasks often handled by entry-level employees or sales assistants, especially for standardized products such as auto or basic life insurance. At the same time, significant protection gaps remain, particularly in broader life insurance coverage and among small businesses that remain underinsured. By lowering distribution and servicing costs, AI allows insurers to reach previously underserved customers, expanding overall market participation and unlocking latent demand. As a result, some routine roles decline, while others shift toward higher-value activities such as policy advisory for more complex products (commercial insurance or annuities, for example) and long-term client relationship management.

IT support technicians provide a similar example. AI can resolve routine tickets and automate diagnostics. As digital infrastructure expands and systems grow more complex, the need for advanced troubleshooting and systems oversight may rise. Entry-level resolution work may decline unless workers can be quickly upskilled, while higher-skilled technical coordination roles persist or grow.

Substituted Roles. Only when demand is capped and AI directly substitutes for human workers in core tasks do roles fall into this category—an outcome that our analysis estimates affects about 12% of current jobs. Efficiency gains convert into net job losses, and downward wage pressure develops for the positions that remain.

Certain financial-analyst roles fall under this category (as do call center representatives). The volume of financial analysis is largely tied to existing reporting cycles, investment mandates, and internal decision processes. When AI automates routine modeling, data aggregation, and initial interpretation, the output does not expand proportionally. In this case, productivity gains are more likely to reduce the number of analysts required than to drive additional hiring.

It is important to note that substitution does not imply permanent exclusion from the workforce. As augmented, less automatable, and emerging roles expand elsewhere in the economy, workers in substitution-exposed positions may transition to such roles with targeted reskilling and mobility support. Starting to plan now for these redeployments—even though they will not happen immediately or all at once—will be critical, as transitions require time, coordination, and investment.

Enabled Roles. AI will become embedded in the day-to-day activities of 23% of jobs, reshaping how tasks are performed, but not fundamentally altering how work is structured. Workers will be expected to use AI to improve efficiency, accuracy, and decision making. While routine elements may be streamlined, core responsibilities—often involving the physical presence of a human worker, interpersonal interactions, or domain-specific expertise—will remain human-led. As AI becomes a baseline capability, skill requirements will rise, making continuous upskilling essential.

Clinical assistants and lab technicians illustrate this dynamic. Their work remains hands-on or patient facing, but AI increasingly supports documentation, workflow coordination, and aspects of diagnostic analysis. For example, clinical assistants may use AI for real-time note taking and patient intake, and lab technicians can leverage AI to interpret test results or flag anomalies. Over time, effective use of these AI tools becomes part of the role, raising expectations for productivity and accuracy without fundamentally changing its nature.

Limited-Exposure Roles. In the remaining share of lower-automation jobs—34% of current jobs—both the technical feasibility of automation and the scope for AI-driven productivity gains remain limited. (Our analysis does not account for potential future breakthroughs in technologies such as humanoid robotics.) Work is often highly contextual, relationship driven, or dependent on a physical human presence in ways that are difficult to codify or scale through AI. As a result, these roles are less likely to be significantly reshaped in the near term.

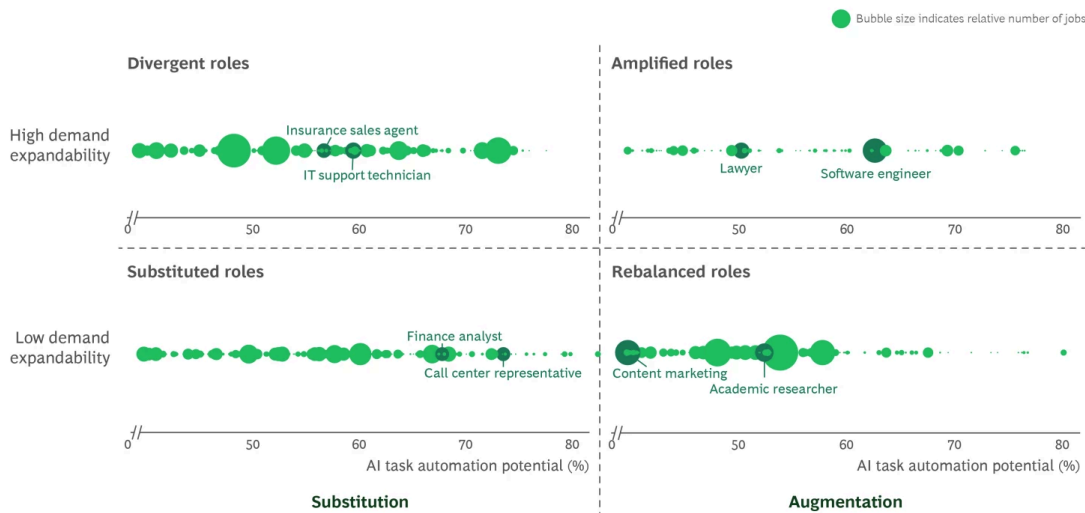
Physicians and teachers are illustrative. Both require the ability to form complex judgments, engage in interpersonal interactions, and adapt in real time to individual needs. AI may assist in limited, task-specific ways, but it will not meaningfully reshape either role. The core of the work—patient care and instruction—will remain fundamentally human.

The Scale of Impact. Across all six segments, 10% to 15% of jobs are vulnerable to elimination; these are the jobs in the substituted and divergent categories, weighted by their automation potential and adjusted for demand expansion. At the same time, 50% to 55% will be reshaped; these are the jobs in the amplified, rebalanced, and enabled categories, as well as the remaining portions of jobs in the divergent and substituted categories where work is not vulnerable to elimination. The remaining roles are largely lacking in significant automation potential in the near term. Meanwhile, an unspecified but meaningful number of jobs are likely to be created as new demand emerges and new capabilities are enabled by AI.

Taking a deeper look at the totality of roles in the US job market and applying our methodology, we can see in Exhibits 5 and 6 how roles are distributed, with bubble size indicating the scale of impact.

EXHIBIT 5

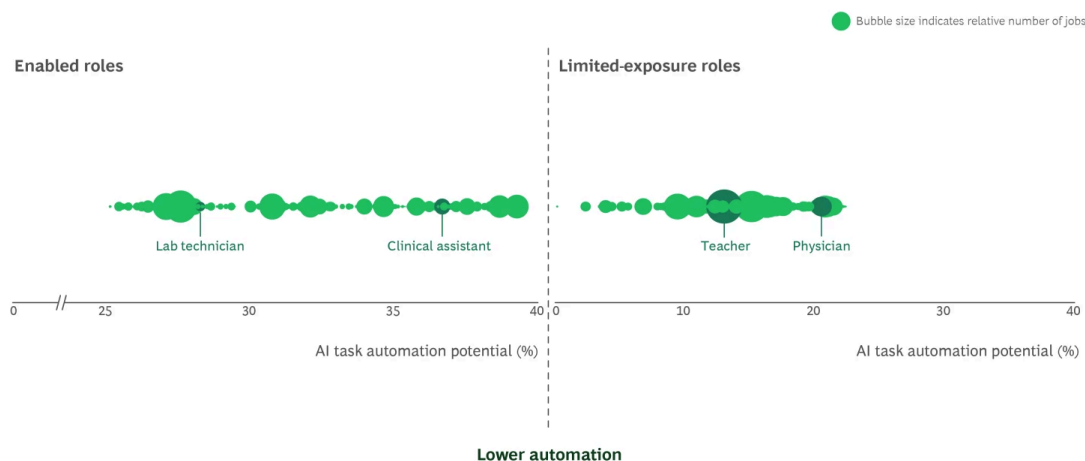
Distribution of Higher-Automation Roles Affected by Automation



Sources: Revelio Labs; O*NET; US Bureau of Labor Statistics; BCG Henderson Institute analysis.

EXHIBIT 6

Distribution of Lower-Automation Roles Influenced by Automation



Sources: Revelio Labs; O*NET; US Bureau of Labor Statistics; BCG Henderson Institute analysis.

Tackling the Great Unknowns

Our analysis enables us to categorize roles, but this isn't the whole story. Each segment—and, potentially, individual roles within each segment—will evolve and have different ripple effects that must be managed in different timeframes. As a result, stability in headline job numbers may mask critical issues for company leaders.

The Side Effects of Agentic AI Role Transformation. First, a key challenge in the AI era will not be simply the number of jobs affected, but how quickly workers can be upskilled and redeployed into redesigned roles. Absorbing a workforce shift of this magnitude will therefore require deliberate investment in reskilling, mobility, and capability building, making workforce development a leadership priority.

In the near term, some organizations may respond by slowing or freezing hiring as they reassess workforce needs and push for higher productivity with AI—creating additional pressure on existing employees to adapt and upskill. In some cases, these measures are being taken ahead of realized AI impacts and should be viewed as a temporary phase rather than a new status quo. Companies whose products and services are in high demand will continue to hire over time as productivity gains translate into growth.

Second, as AI absorbs much of the routine work that has historically justified large entry-level hiring cohorts, fewer execution-focused positions will be required. In the short term, the volume of entry-level jobs may decrease. Some companies will still choose to invest in junior workers in order to create a talent pipeline, while others may cut back. Over time, though, these jobs will be redefined and filled by candidates prepared to take on higher-order tasks, such as supervising AI outputs, managing exceptions, and contributing to more complex problem solving earlier in their careers.


In this environment, AI fluency may become an increasingly important complement to tenure in assessing readiness for augmented responsibilities, with companies prioritizing candidates who proactively adopt AI tools and rethink how work is done. In some cases, this will create opportunities for junior candidates—including recent graduates—who may be more familiar with AI than workers with more experience on the job.

Third, skill thresholds will rise. Redesigned roles will demand that employees demonstrate greater expertise, oversight, and accountability, increasing the premium on domain knowledge and sound judgment. As Exhibit 7 illustrates, the more durable roles tend to require higher credentials and greater seniority. Moving into these roles requires new capabilities and, in many cases, additional training or formal credentials. Even where total employment remains steady, these higher thresholds can create barriers to entry and friction in workforce transitions.


EXHIBIT 7

More Durable Roles Tend to Require More Seniority and Higher Credentials

Seniority by role category



Credentialing requirements by role category



Sources: Revelio Labs; BCG Henderson Institute analysis.

Note: Seniority analysis was not applied to lower-automation roles because the expected disruption is more limited.

Finally, cognitive load will intensify. Many roles currently balance structured execution with higher-level thinking. As repetitive tasks are automated, the remaining work will be concentrated in problem solving, decision making, and the integration of complex inputs, increasing the cognitive intensity of work. While some workers will thrive in more judgment-driven roles, others may struggle with the shift toward continuous high-level cognitive engagement and will require upskilling. Without deliberate role design and capability building, organizations risk losing out on AI’s full productivity potential.

These four dynamics can generate short-term turnover, uncertainty, and disruption at the job level, even when aggregate employment remains stable. The labor impact of AI is therefore not only about how many jobs persist, but how accessible and sustainable those jobs become.

Open Questions on the Timeline for Change. Another key unknown is timing. Economic impact often lags model capability, as the former also depends on application maturity, workflow redesign, integration with legacy systems, and the availability of human capital capable of deploying and managing these AI systems.

Contact center tools are among the most mature applications, yet overall market penetration remains limited relative to total industry size. In our experience, full worker substitution tends to roll out more slowly than augmentation. Substitution often means that fewer humans are retained in the loop, which demands extensive process redesign and formalization of tacit knowledge. Human workers may also form a critical “escalation layer” to handle outcomes that agents can’t deliver. By contrast, augmentation can diffuse more quickly because humans remain in force and can manage context, edge cases, and ambiguity during the transition.

Scaling agentic systems requires specialized integration talent, including forward-deployed engineers, systems integrators, and project managers who tailor systems to enterprise-specific contexts. These technical experts are embedded directly with business teams, translating AI capabilities into real, working solutions. They sit between AI technology and the organization, adapting workflows, integrating legacy systems, and ensuring that outputs are reliable and

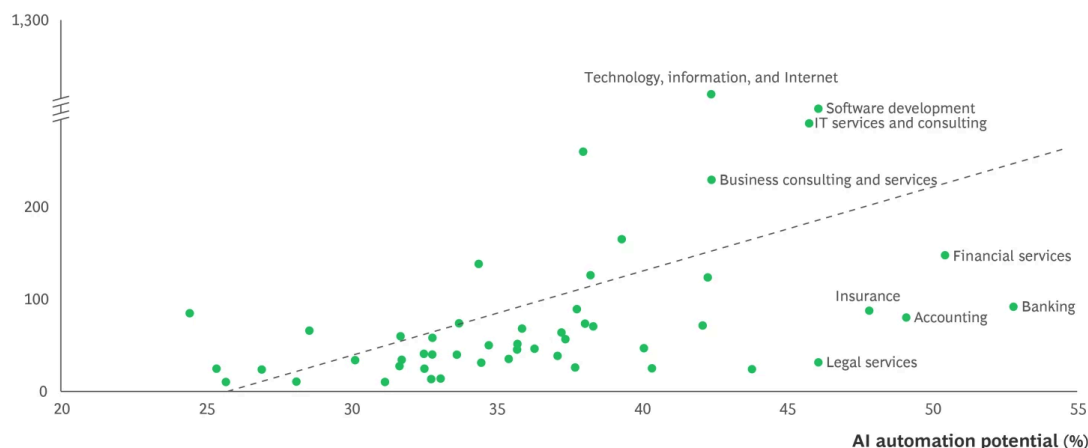
usable. Across these integration roles, supply remains limited relative to demand, making implementation capacity a key bottleneck. These roles are also an example of new jobs emerging from AI adoption, with demand expected to increase as enterprises scale deployment.

As a result, there will likely be a multiyear lag between automation potential and realized labor market impact. Exhibit 8 clearly illustrates this gap. Industries with high automation potential do not always demonstrate high levels of scaled AI adoption. In others, such as technology and software, adoption is already well above average, although today's average remains far from the full potential for most industries. Still other industries, including financial services and legal services, have substantial automation potential but scaled deployment has yet to catch up. Diffusion will vary by sector, company size, and the availability of integration expertise, with larger enterprises typically moving faster than smaller organizations owing to resourcing, system, and data access issues.

EXHIBIT 8

The Gap Between AI Potential and Scaled Adoption Varies Across Industries

AI adoption versus average (%)



Sources: Revelio Labs; O*NET; HG Insights; BCG Henderson Institute analysis.

Note: AI adoption is indexed to the weighted average across all US employees; it is measured using the share of employees with AI skills, the share of AI-related positions, the share of AI job postings, and the number of AI technology vendors used by the company.

What CEOs Need to Do Now to Prepare

There are four critical starting points for leaders.

Embed workforce strategy into competitive strategy. AI will reshape competitive dynamics and enable new business models. Companies that respond by accelerating innovation, redesigning

offerings, or reconfiguring their delivery of value at the enterprise, business line, product, and functional levels will fundamentally change the nature of work inside their organizations. Workforce strategy cannot sit downstream of automation. It must be embedded in strategic planning. Remaking the workforce is a competitive imperative.

Leaders should avoid reactive actions on cost driven by headlines or peer behavior. Workforce decisions must reflect the specific mix of automatable and augmentable work within the specific enterprise. Copying another company's reductions without understanding one's own exposure risks undermines both productivity and long-term competitiveness.

Refocus automation on redesign, not just cost reduction. Agentic AI isn't a blunt instrument. Different roles require different approaches. Sometimes cost is a dominant factor. At other times, other priorities dominate, such as speed of work and quality. Cost actions—such as headcount freezes or reductions—are visible and straightforward, with explicit OpEx impact. But when AI drives productivity rather than cuts, ROI becomes harder to define and defend.

In these situations, leaders must redesign workflows and rethink how performance is measured. Capturing value requires new, domain-specific KPIs that link productivity gains to tangible outcomes, such as higher revenue per FTE, more product shipped, or stronger customer impact. Task turnover within a role will be a measure of how fast roles can evolve toward higher-value work. Realizing AI's upside often means growing into the improvement—translating productivity gains into financial results through budgeting and workforce planning.

Put upskilling, reskilling, and redeployment at the center of workforce strategy. Just because a job remains does not mean employees are prepared for it. Upskilling, reskilling, and structured redeployment pathways must become central to workforce strategy. Moreover, workers will likely require more frequent (versus one-time) upskilling as the technology evolves. Leaders should prepare their workforce in phases aligned with how roles are evolving and create clear pathways to adjacent, augmented, and emerging positions. In each of the five AI Labor Disruption Segments in which AI is expected to materially impact work, leaders should keep the following considerations in mind for each category:

- **Amplified Roles.** Leaders should focus on retaining and developing this talent, reinventing team-based workflows to maximize human–AI collaboration and redefining what excellence looks like. As productivity expectations rise, leaders should remain mindful of cognitive overload and actively manage workloads. Reinforcing the narrative that AI enhances long-term value will be critical to sustaining engagement and performance.
- **Rebalanced Roles.** The opportunity lies in role redesign. Leaders should identify repeatable components that can be automated and reinvest the time savings in higher-value activities. Upskilling in AI fluency and structured experimentation is essential to ensure that higher-order skills translate into measurable improvements in output quality.

- **Divergent Roles.** The challenge here is structural. Leaders should redesign career pathways with deliberation, creating accelerated development tracks and clear routes into higher-skilled responsibilities. Preserving entry points, investing in young talent, and managing progression thoughtfully is essential to avoiding long-term skill gaps.
- **Substituted Roles.** Leaders should reimagine processes end to end around AI agents. Workforce transition planning must be developed in parallel with AI deployment. Redeployment pathways, transition support, and change management should be designed alongside automation initiatives.
- **Enabled Roles.** Leaders should focus on embedding AI into day-to-day workflows and building broad-based AI fluency across the workforce. The priority is consistent adoption: removing barriers to usage, standardizing tools and workflows, and setting clear expectations so that gains are realized widely.

Shape the AI narrative to unlock performance. Sequencing and signaling matter. Focusing first on highly substitutable roles may deliver short-term efficiency gains, but it can create a demoralized atmosphere that undermines broader transformation. When employees associate automation with displacement, engagement declines and the motivation to upskill erodes. They may resist augmentation efforts, even when designed to elevate their roles. Leaders must communicate clearly that if workers upskill, AI in most roles will be about value creation, not displacement. The narrative set at the top will shape whether the workforce embraces transformation or resists it.

AI creates a massive opportunity for business leaders but also significant uncertainty in terms of how to unlock it. Moreover, this is all happening in a charged environment. In some cases, restructurings that would have occurred regardless, as part of the normal business cycle, are likely to be attributed to AI, which will create fear at a societal level.

At the same time, AI's impact will vary significantly across companies. Some will lean into AI to drive innovation and growth, while others will focus on efficiency and automation. This could lead to very different talent strategies, with some reducing headcount while others hire aggressively.

For CEOs, the imperative today is to focus on achieving the right balance of automation, upskilling, and deliberate talent planning—to deliver enterprise ROI at scale and help their employees develop the skills they need to thrive in the AI era.

The authors would like to thank Djon Klein, Philip Ventura, Nissim Ray, and Hanson Wong for their valuable contributions to the research and analysis that informed this article.

The BCG Henderson Institute is Boston Consulting Group’s strategy think tank, dedicated to exploring and developing valuable new insights from business, technology, and science by embracing the powerful technology of ideas. The Institute engages leaders in provocative discussion and experimentation to expand the boundaries of business theory and practice and to translate innovative ideas from within and beyond business. For more ideas and inspiration from the Institute, please visit our [website](#) and follow us on [LinkedIn](#) and [X \(formerly Twitter\)](#).

Authors



Greg Emerson

Managing Director & Senior Partner
San Francisco - Bay Area



Matthew Kropp

Managing Director & Senior Partner
San Francisco - Bay Area



Julie Bedard

Managing Director & Partner;
BCG Henderson Institute Fellow
Boston



Lisa Kraye

Principal
Philadelphia



Viacheslav Romanov

Partner & Associate Director,
PIPE Tech Capital
Boston



Megan Hsu

Project Leader
San Diego



Luis Sanchez Boedo

Consultant
Raleigh-Durham



Diya Mohnot

Associate
Philadelphia



ABOUT BOSTON CONSULTING GROUP

Boston Consulting Group partners with leaders in business and society to tackle their most important challenges and capture their greatest opportunities. BCG was the pioneer in business strategy when it was founded in 1963. Today, we work closely with clients to embrace a transformational approach aimed at benefiting all stakeholders—empowering organizations to grow, build sustainable competitive advantage, and drive positive societal impact.

Our diverse, global teams bring deep industry and functional expertise and a range of perspectives that question the status quo and spark change. BCG delivers solutions through leading-edge management consulting, technology and design, and corporate and digital ventures. We work in a uniquely collaborative model across the firm and throughout all levels of the client organization, fueled by the goal of helping our clients thrive and enabling them to make the world a better place.

© Boston Consulting Group 2026. All rights reserved.

For information or permission to reprint, please contact BCG at permissions@bcg.com. To find the latest BCG content and register to receive e-alerts on this topic or others, please visit bcg.com. Follow Boston Consulting Group on [Facebook](#) and [X \(formerly Twitter\)](#).

¹ This number assumes 165 million US jobs filling 1,500 roles.

* We sourced real output (value added), industry price indices, and real GDP from the US Bureau of Economic Analysis.