PEOPLE STRATEGY

When Everyone Uses AI, Companies Risk Losing Critical Skills

By Sagar Goel, David Martin, and Charikleia Kaffe

ARTICLE JUNE 17, 2026 15 MIN READ

Humans have offloaded cognition to tools since ancient times, from the first tally mark to calculators and GPS devices. What makes generative AI (GenAI) different is the nature of the work it performs. Unlike previous tools, GenAI doesn't merely support human thinking; increasingly, it substitutes for it, with consequences for a specific set of human skills. A growing body of research points to *cognitive debt*—the gradual erosion of critical thinking, judgment, curiosity, and originality among frequent AI users.

But the larger danger is not what happens to an individual employee; it is what happens when skill erosion occurs simultaneously across hundreds or thousands of people in an organization. We call this *distributed de-skilling*—a collective erosion of human skills that undermines organizational intelligence and resilience over time. This phenomenon is not a talent problem, but a system design problem rooted in how the organization builds governance, workflows, and culture around AI.

In a global study of 70 C-suite leaders and senior executives across multiple industries, we found that half are already observing de-skilling in their organizations, and more than 60% believe that de-skilling will pose a material threat within the next three to five years. On the basis of these leaders' firsthand accounts, as well as interviews that we conducted with a dozen other company leaders, we've identified the organizational implications of distributed de-skilling, the skills that are most vulnerable to it, and some of the mitigation strategies that leaders are deploying to combat it.

Critical Skills Are Vulnerable

The skills that leaders identify as most important to long-term organizational performance—judgment and decision making, problem understanding and framing, creative thinking, analysis and causal reasoning, and solution generation and evaluation—are precisely the ones that they consider most at risk of de-skilling due to AI. (See the exhibit.) These capabilities lie at the core of what it means to think strategically, and they are the skills that GenAI most directly replaces when used without guardrails.

EXHIBIT 1

The Skills That Matter Most Are Those Most at Risk

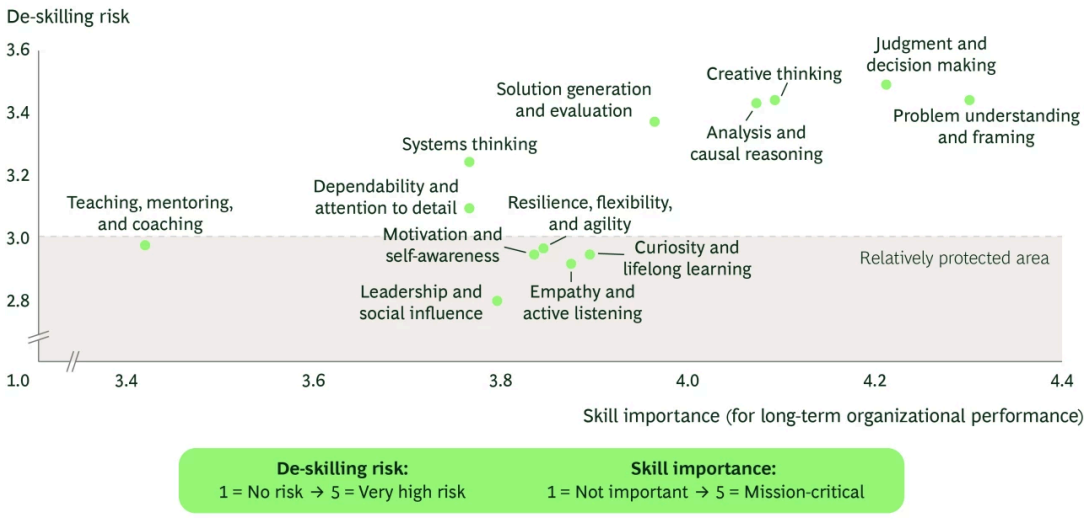

Source: BCG Institute analysis.

Problem understanding and framing stands out as the skill that leaders rate as most critical to organizational performance, while judgment and decision making carries the highest de-skilling risk score of all skills. They are the two key components of strategic thinking—which involves defining the right problem and making the right call—and both are deeply exposed.

Problem understanding and framing is upstream of everything else: an organization that loses the ability to define the right problem cannot compensate for that deficiency by using AI to generate better analysis or faster execution. Judgment and decision making, meanwhile, entails a sequence: attempting to solve a problem with incomplete information, making a call, observing what happens, and updating the mental model accordingly. One leader noted, "I've seen AI used as the wise man in the room but with insufficient input, short prompts, and not grounded to company strategy or company values. This produces fine high-level but generic outcomes, and it causes people to have differing expectations."

Creative thinking, analysis and causal reasoning, and solution generation and evaluation are at risk in related ways. Because AI produces seemingly acceptable ideas efficiently enough, employees' independent ideation abilities can atrophy; and because its strength in pattern recognition creates the illusion of analytical rigor, workers' ability to perform deeper causal reasoning and assumption testing can erode. For example, several leaders flagged a growing tendency for teams to accept AI-generated analyses without interrogating the underlying logic.

By contrast, the skills sitting in the relatively protected area—empathy and active listening; motivation and self-awareness; curiosity and lifelong learning; resilience, flexibility, and agility; and leadership and social influence—are those that depend most heavily on human relationships, embodied experience, and emotional attunement. AI can assist around the edges of these skills, but it cannot substitute for them. They are, for now, more durable.

Crucially, these less exposed skills do not just constitute reassuring data points. They can actively defend skills that are at greater risk. For instance, empathy and active listening can sustain the mentoring relationships that AI is quietly displacing; leadership and social influence can make structured debate possible in situations where convergent AI thinking would otherwise go unchallenged; and motivation and self-awareness can drive employees to interrogate AI outputs rather than simply accept them. Leaders would do well to invest in these skills with that purpose explicitly in mind.

Skill Erosion at Scale

What happens when entire teams, business units, or organizations experience de-skilling? Already, the organizational symptoms that leaders associate with capability decline are visible and mutually reinforcing. The downward trend starts with overreliance on AI outputs without stress testing or challenge, a factor cited by almost 90% of surveyed leaders. As one leader observed, “Teams start treating AI-generated work as ‘good enough’ and skip the critical thinking step entirely. Over time, nobody is poking holes in the analysis, which means errors and blind spots slip through unchecked.” Uncritical acceptance is not just a quality problem, but also the gateway to a deeper erosion of judgment.

That erosion shows up most visibly in an environment of reduced ownership and accountability, cited by more than 50% of leaders. Once employees stop interrogating AI outputs, they find it easier to stop owning them. “The AI suggested it” becomes a shield and nobody feels fully on the hook. One leader described it bluntly: “I can see employees’ sense of ownership going down. If the outcome is not positive, the responsibility is laid on AI.” Leaders expect this pattern of deflecting blame onto algorithmic systems to accelerate.

As individual judgment weakens, the collective suffers. Fewer constructive debates (cited by 43% of leaders), less diversity of thinking (cited by 49% of leaders), and reduced collaboration and teaming by employees (cited by 34% of leaders) follow from the same dynamic. When everyone uses the same AI tools trained on the same data, outputs and thinking converge, draining the impetus to consult, challenge, or build on a colleague’s perspective. “Everyone will have similar answers, similar formats, similar data points,” one leader said. “Logically, every problem will have a set solution.” Speed compounds this problem: “There is less discussion because everything needs to be done immediately. Winning at speed eliminates good debate.”

Underlying these trends is a deeper structural problem: the slower development of junior talent, cited by 53% of leaders. When AI handles the analytical grunt work—such as research, drafting, debugging, and problem decomposition—that junior employees have traditionally performed and learned from, they miss out on the formative repetitions that build judgment. Several leaders noted that a person cannot develop discernment without first doing the work. The reinforcing cycle of de-skilling is already visible, as one leader described, “Less hands-on practice feeds into less ownership, which feeds into less scrutiny of AI outputs.” In turn, this means fewer opportunities to exercise judgment and engage in hands-on practice.

But the risk runs even deeper. As AI absorbs entry-level tasks, organizations expect new hires to perform without the traditional on-ramp that builds seniority and clarifies expectations. As one leader put it, “Our traditional entry-level employee will need to perform tasks that a five- to ten-year employee would do.” Organizations are restructuring the career ladder faster than they are redesigning the learning paths that made the old system work. Experienced professionals are equally vulnerable to what one leader called the “autopilot trap”—substituting AI judgment for their own out of habit and in response to time pressure. Senior employees can retain the appearance of high performance while gradually losing the depth of judgment that made them valuable.

The decline in mentoring, coaching, and knowledge sharing (cited by 33% of leaders) is a particular concern for long-term organizational learning. Several leaders noted that employees now bypass peer knowledge exchange entirely, going directly to large-language models (LLMs) for insight rather than engaging colleagues. One leader described training days that are "no longer used to learn but to generate content with AI tools"—a sign that the rituals designed to build capability are eroding.

Six Strategies to Mitigate De-Skilling Risk

Although they are aware of the threat, most organizations remain in the early stages of responding to it. In our research, only one in ten companies has an organization-wide strategy or has launched targeted initiatives to address de-skilling, and a third have not explicitly discussed it at all. That said, some companies have begun experimenting with interventions. Through our research and conversations, we've identified six strategies that leaders can implement today.

1. Set the Organizational Conditions

Effective AI governance does more than define what *not* to do: it actively shapes how employees can use AI in ways that strengthen human skills. This involves defining which AI systems are approved for use, how output verification should be conducted, and when human review is mandatory. It also means codifying clear right-use patterns with concrete guidance detailing how to use AI in specific contexts and defining AI-off zones that apply to tasks where originality, ethical judgment, or synthesis is most critical.

The organization should embed these conditions in a practical AI code of conduct that addresses how to manage hallucinations, how to interrogate model reasoning, and how to escalate issues appropriately. Employees should receive consistent reminders (through policy and company communication) that they own the outcome regardless of the tool they use. At the same time, leaders must actively manage the internal narrative, avoiding the trap of incentivizing the use of AI tools for its own sake rather than for genuine value creation.

For example, at a leading multinational materials science company, the center of excellence within the enterprise data and analytics department includes a dedicated change management function whose explicit purpose is to ensure preservation of the human element as the company adopts AI tools. It works directly with business and functional stakeholders to coach people away from old habits and toward new ones, clarifying what the new way of working looks like and what its benefits are, and deploying digital literacy programs in parallel to close skill gaps.

2. Redesign How Work Gets Done

The core design question is work allocation: what tasks should be delegated to AI, what tasks should remain human, and where should humans and AI adopt an ensemble approach (working on a problem independently, and then combining their outputs)? As organizations redesign their workflows, resistance to skill erosion needs to sit alongside efficiency and effectiveness as an explicit design criterion, with leaders asking, "What work should we keep humans involved in, and why?"

Several leaders noted that an organization's most consequential decisions are not about individual tasks but about sequences of work that build judgment or require contextual reasoning. These must remain human-led even as surrounding tasks are automated. Control of sequencing is critical, especially now that AI agents can handle multistep workflows end-to-end, compressing the cognitive loops that build capability. In practice, this means that humans need to retain ownership of final judgments and accountability for outcomes.

Another key dimension involves preserving the senior-junior mentoring dynamic through which organizations transmit tacit knowledge, professional judgment, and craft across generations. An organization must protect or rebuild this relationship through structured apprenticeships, coaching relationships, and team designs that require human-to-human knowledge transfer.

For example, Shell has adjusted the workflows of its teams so that humans explicitly own interpretation, validation, and final decisions, with peer review expected for material outputs. Human-in-the-loop decision checkpoints require employees to challenge AI outputs and assumptions before approving them. The company has also fundamentally redesigned junior learning paths: juniors must first independently frame the problem, test assumptions, and produce a baseline analysis, and only afterward use AI to refine their analysis. Early results show improved question quality, clearer rationales for decisions, and faster progress toward independent work.

To preserve the mentoring dynamic, teams at Salesforce draw on the concept of pair programming from software development, creating small team structures that mix high-agency AI adopters with others who are less comfortable with the tools. The approach relies on positive peer pressure and direct observation: employees learn more effectively by watching a skilled colleague than by undergoing formal training alone, and having a partner is invaluable for employees who have access to the tools but don't know where to start with them.

3. Make Human Skill Development Visible in Performance Systems

Embedding human skill development in performance systems is the hardest and least implemented intervention, but it is also one of the most structurally consequential. As AI compresses the distribution of visible performance, leaders find it increasingly difficult to differentiate talent and to reward genuine capability. The same problem arises in the context of learning and development: certification rates go up while actual capability may not. Performance systems need to see through this fog.

The response is to make *how* an employee achieves outcomes as visible as *what* they deliver, by assessing the skills deployed and built. Human capabilities required for long-term organizational performance should be explicitly identified in reviews, promotion criteria, and development pathways, as a core signal of what the organization values and intends to protect.

At CNIL, France's data protection authority, managers are responsible for assessing employees' ability to challenge AI outputs—not just to use them—making critical engagement with AI a visible and evaluated dimension of performance. This shift in assessment procedure has reinforced accountability and helped sustain critical thinking that governance alone cannot guarantee.

In implementing changes in performance systems, leaders at Salesforce found that introducing revised evaluation frameworks alongside new technology created anxiety among employees and distorted their performance. To mitigate this effect, the organization kept its evaluation cadences stable but shifted the incentive structure to reward employees who generate insights that go beyond their immediate job requirements and who foster a culture of sharing rather than hoarding knowledge.

4. Build Skill-Replenishment Rituals into Everyday Work

Leaders must actively establish routines that will help employees rebuild cognitive and creative muscle. Just as organizations have norms for different meeting types, they need norms that exercise different human capabilities: cognitive warm-ups, AI-free problem-solving sessions, structured debate moments, and reflection prompts ("What reasoning step did we strengthen today?").

For a replicable recurring format, one leading Indian bank runs a structured AI-free session on the first Friday of each month across all business and support functions, combining case study discussions, design thinking assignments, team retrospectives, and sector briefings. The structure rotates among analytical, creative, and reflective modes to exercise the full range of human capabilities. A US-based B2B telecommunications company is introducing yearly no-AI hackathons that give teams three days to build executable proofs of concept, thereby exercising problem-solving, creativity, and a growth mindset.

To embed deliberation in existing workflows, teams at a leading industrial technology company use the planning phases for new activities as structured moments for human judgment. Teams are required to formulate different scenarios, defend specific positions, and challenge one another's reasoning before the company chooses a particular path. AI can help generate the scenarios, but humans must provide the critical judgment.

5. Train Employees to Work with AI in Nonlinear Ways

Most people use AI linearly: describe the task, ask for an answer, tidy the output. This pattern can reinforce cognitive offloading and produce homogeneous thinking because skipping straight to a generated answer omits the cognitive work that builds judgment. The alternative is to use AI not as an answer machine but as a tool for provoking better questions, and to work backward from

opposition and failure modes (such as, “How could this go wrong?”) rather than forward from assumptions.

Several powerful nonlinear techniques are already in use across industries. In design education, there is a practice of prompting for the opposite. The idea is to instruct AI, “Give me a solution that does not work,” rather than “Give me a better solution.” This approach forces the tool to surface ways in which the solution could break, unexpected edge cases, and unconventional approaches that normal prompts overlook—all of which generate raw material for more original thinking.

In legal and policy work, adversarial red teaming uses AI as a skeptic rather than as an answer machine, working backward from worst-case scenarios to uncover vulnerabilities before they become problems. And in product development, Amazon’s PR/FAQ “working backwards” method flips the workflow entirely: teams write the press release and customer FAQ for a product before any development begins, forcing clarity on vision and constraints upfront.

These techniques share a common operational logic: use AI to provoke better questions, not just faster answers. Leaders who teach employees to work with AI in this way will find that the quality of AI-augmented thinking improves because the human doing the prompting has done the hard cognitive work first.

6. Embed Reflective Prompts and Cognitive Nudges in AI Tools

The design question that every leader should ask about AI tooling is “Does this keep our people thinking?” The most direct way to answer that question is to build reflective prompts into tools at the point of use so that instead of simply delivering an answer, the tool requires the user to think first. Several leaders called for AI models that actively surface levels of certainty, present counterarguments before a conclusion is accepted, and embed prompts such as “What assumptions are you making?" or require a brief human rationale before generating a recommendation.

These nudges serve a dual purpose. In the short term, they improve output quality by slowing down the unchallenged acceptance of AI-generated conclusions. Over time, they build employees’ meta-cognitive capacity—the ability to monitor their own thinking, recognize when they are offloading too readily, and deliberately re-engage their own judgment. The goal is that employees eventually internalize the prompt, asking “what am I missing?” before the tool asks it for them.

For example, a global information services tech company has deliberately designed in-house LLM-based tools to point users toward areas where they should apply their own judgment, rather than providing the solution directly. The prompt architecture itself is structured to provoke thinking, not replace it.

AI failure drills extend this logic further. Rather than simply teaching employees how to use AI well, some organizations are deliberately exposing them to AI behaving badly—introducing hallucinations, unexpected answers, and curveball questions into training modules to test how

employees respond when AI behaves unpredictably. Exposure to AI failure, in controlled conditions, directly builds the habit of questioning outputs rather than accepting them.

Capability as a Competitive Advantage

The skills most at risk—judgment and decision making, problem understanding and framing, creative thinking, analysis and causal reasoning, and solution generation and evaluation—share a common vulnerability: they survive only through active use. The six strategies above are designed to protect those skills, ensuring that humans remain the originators of questions and judgment. Together, these strategies set the organizational conditions that govern AI use, redesign the daily mechanics of work, and ensure that the tools, rituals, and incentives within the environment itself actively build human capability.

But the response cannot stop at protecting the skills most at risk. Leaders should also invest deliberately in relatively protected skills, such as empathy and active listening; curiosity and lifelong learning; motivation and self-awareness; resilience, flexibility, and agility; and leadership and social influence. These skills are not merely durable. They are the organizational immune system through which the most vulnerable capabilities are sustained.

Human skills are not a fixed endowment. They are a renewable asset that atrophies without use and grows with deliberate practice. Leaders who recognize this and act on it now will find that managing distributed de-skilling is not just a risk mitigation exercise, but a source of sustained competitive advantage.

The BCG 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

Sagar Goel

Managing Director & Partner
Singapore

David Martin

Managing Director & Senior Partner, Global Lead, People & Organization
Dallas

Charikleia Kaffe

Research Director
Athens

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.