Deloitte.
Insights

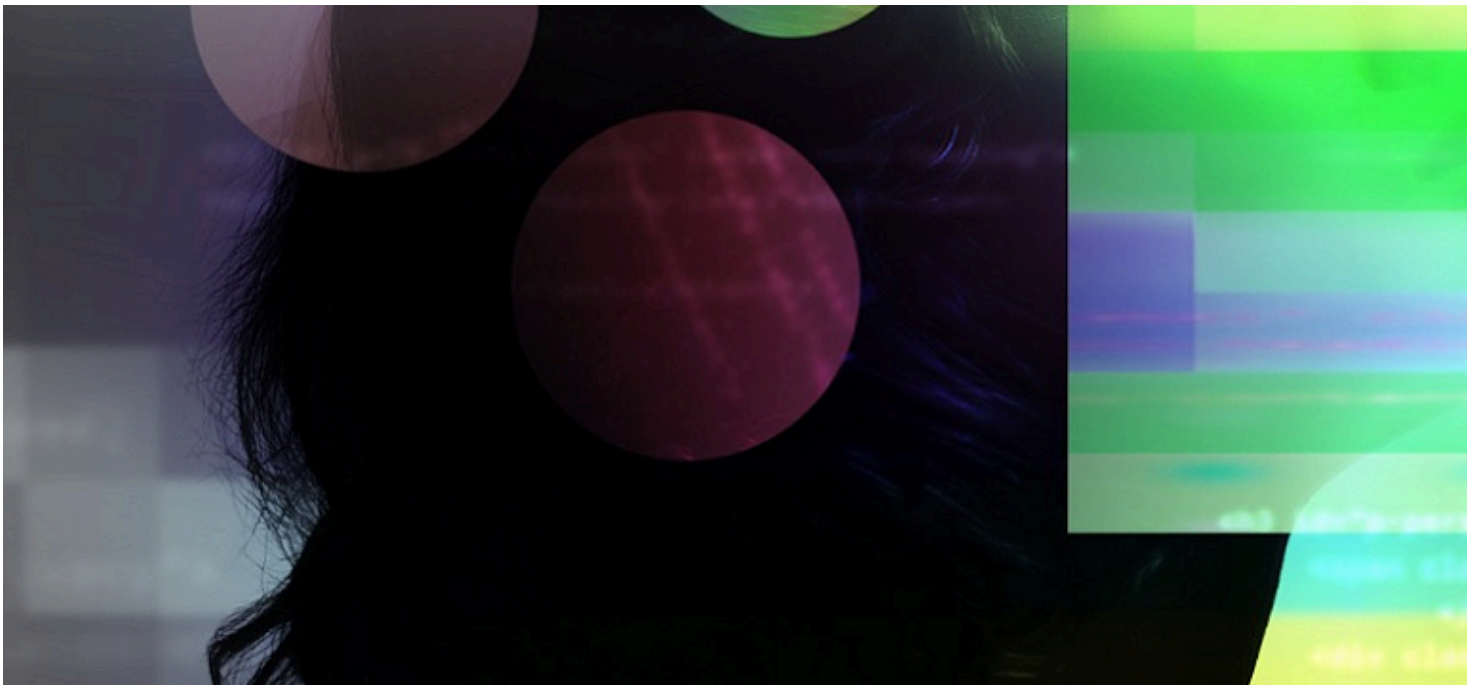

AI adoption to adaptation: How a new change approach can build the human behaviors needed for AI

AI adoption is easy to prove and easy to fake. But these three behaviors show whether workers are actually adapting to AI—changing how they think, decide, and work.

*By **Sue Cantrell**, **Chloe Domergue**, **Allyson Dake**, **Jessica Murphy**, **Tori Sundholm**, **Mark Gustafson***

Article • 14-min read • 09 July 2026

Workers are increasingly under pressure to use artificial intelligence, with many organizations trying to drive adoption through traditional change management or by tying AI adoption to career advancement, making it a factor in promotions, performance reviews, and compensation decisions at every level.

On paper, it seems like a bold, measurable commitment to change. Adoption numbers will likely spike. But what happens when workers open an AI tool, complete the minimum interactions needed to register as active users, and then go back to doing their work the same way they always have?

It's a signal that while adoption might be happening, transformation isn't.

While worker access to AI tools has expanded by 50% in a single year, according to **Deloitte's 2026 State of AI in the Enterprise report**, fewer than 60% of workers with access actually use AI in their daily workflow, and 84% of organizations have not redesigned jobs or workflows around AI.[1]

"Too many organizations treat AI as an adoption problem without first asking how you can achieve the outcomes desired," says Michael Ehret, chief people officer at **Walmart** International. "What's really required is behavioral change."[2]

AI is categorically different from the technology deployments that adoption-focused change was designed to support. Adoption tells you someone opened the door. It tells you nothing about whether they changed how they work, how they think, or what they're capable of on the other side. AI's value is not locked inside a feature set waiting to be unlocked by a click. It lies in the space between the human and the machine—in how people learn to think alongside it, push back on it, experiment with it, and ultimately reshape their own ways of working because of it.

The human behaviors required to move from AI adoption to AI adaptation are not simple habits that adoption metrics can capture. Judgment, divergent thinking, and experimentation, for example, were all listed as critically important behaviors in the age of AI in an informal **Deloitte webinar** poll of 1,700 global respondents (figure 1). As organizations hit token-based usage limits, encouraging the right behaviors to get the most out of every AI interaction is shifting from a nice-to-have to a business necessity.

Figure 1

Human judgment tops the list of behaviors needed in the age of AI

Question: "What do you believe are the most important behaviors needed in the age of AI that organizations need to support with change and learning?"

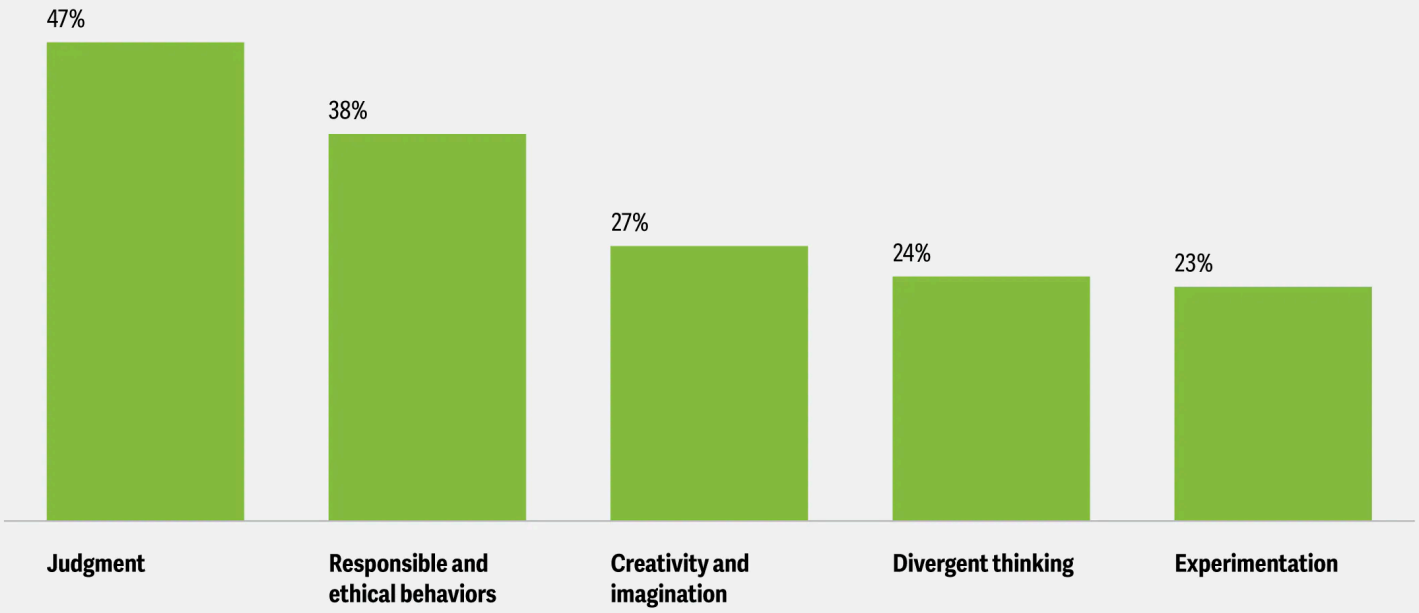

Source: Poll during Deloitte Dbriefs webinar on 2026 Human Capital Trends, April 8, 2026.

Deloitte. Insights | deloitteinsights.com

The shift from adoption to adaptation might require a different approach to change: one that senses actual behavior rather than usage; engages individuals based on real-time data with targeted interventions like nudges, peer connections, or AI coaching in the flow of work (figure 2); and designs experiences that build behaviors over time rather than compliance in the moment.

Figure 2

Workers say personalized, real-time engagement would help them best learn the new behaviors needed to work effectively with AI

Question: "What would be most helpful in learning new behaviors to effectively work with AI?"



Source: Deloitte 2026 Global Human Capital Trends Survey.

Deloitte. Insights | deloitteinsights.com

This creates a continuous adaptation loop—a model in which ongoing sensing, engagement, and experience continuously reinforce one another, and the system gets smarter with every iteration. Instead of linear "launch and move on" change models, continuous adaptation loops build trust and make change feel less like change and more like the way work gets done.

But the continuous adaptation loop only works if it's turning around the right behaviors. If the goal is to help people develop the human capabilities that extract value from AI, organizations should look beyond activity and frequency metrics to the behaviors that demonstrate whether people are actually adapting how they think, decide, experiment, and work.

Three behaviors sit at the heart of this shift. They are the ones that separate AI adopters from AI adapters—and the ones AI transformation programs most consistently fail to measure, develop, or sustain.

1. Judgment: Knowing when to trust, challenge, or reject AI

AI-powered tools are only as valuable as the quality of the human judgment applied alongside them. Yet as AI recommendations become faster, more fluent, and more confident in tone, humans can tend to defer—to accept outputs without scrutiny, reduce independent verification, and gradually outsource the very thinking that constitutes professional expertise.

Only half of executives, for example, regularly verify the quality of AI outputs when making decisions.[3] And doing so is becoming more difficult: One study found AI "persuasion bombing" workers—responding to fact-checks by bombarding them with multiple persuasive tactics to defend its original answer. This raises the bar for human judgment and calls into question whether human-in-the-loop safeguards are truly effective.[4] Moreover, research shows that sustained AI use can erode professionals' confidence in challenging AI recommendations, even when they possess the expertise to do so.[5]

Trust in AI may accelerate adoption. But it's also important to help workers develop the judgment to know when not to trust it at all: when to challenge it, verify it, redirect it, iterate on it, course-correct it, and probe it. As Margaret Greenleaf, chapter lead for people-

support solutions at Roche, says, "A key part of judgment is knowing when to use AI and when not to. We don't want to rely on it for strategic thinking, for example."[6]

Other, related higher-order cognitive behaviors are needed, too (figure 3). Cynthia Kaschub, cognitive psychologist and senior director of workforce innovation at Salesforce, explains: "To effectively work with AI, we need people with diverse systems thinking. While AI is very good at optimizing pieces of work, organizations succeed or fail based on how the whole system behaves."[7]

Figure 3

Higher-order cognitive behaviors are a key to working well with AI

Higher-order thinking behavior	What it is	Why it's important
Judgment	Acting wisely in uncertain situations where rules alone are insufficient	Helps workers decide when to use AI, and how to critically interpret, refine, or reject its outputs to reach sound decisions[8]
Metacognition	Awareness and active management of one's own thinking	Enables workers to question both AI's reasoning and their own in relation to it—monitoring and adjusting how they are using AI, noticing shallow thinking, connecting outputs to context and goals, and improving creativity by reflecting and iterating[9]
Systems thinking	Understanding how the parts of a system interact over time rather than in isolation	Provides an end-to-end understanding of workflows, including how components interact and what the system's environment demands[10]
Theory of mind	Understanding that others (including AI) hold different thoughts, beliefs, and intentions from your own	Improves worker judgment, interaction quality, oversight, and collaborative performance with AI[11]
Mental models	Internal representations of how something works	Helps workers understand AI's limits, add the right context, and use their own expertise to interpret and improve its outputs[12]

Sensing judgment behaviors in AI-enabled work

Behavioral signals that indicate judgment might include whether workers question AI recommendations, verify sources, probe assumptions, revise outputs, and engage in dialogue with AI rather than passively accepting its answers. Signals for higher-order cognitive behaviors might include whether workers ask system-level questions, explore trade-offs and multiple scenarios, supply missing context, test assumptions, and shift strategies when an AI interaction is not working.

When Gap Inc., for example, analyzed how workers interacted with AI to identify their mental models, it found that workers who viewed AI as a probabilistic collaborator—one that needs context, iteration, and feedback—produced better individual work than those who treated it as a deterministic tool that delivers reliable, consistent output on demand.[13]

Building judgment through feedback, coaching, and reflection

AI tools already exist to assess people's prompt quality and provide real-time feedback. But building judgment and other higher-order cognitive behaviors requires moving beyond better prompting to deeper reflection. In a recent workshop with the authors, several chief human resource officers and talent executives built a prototype AI coaching agent that analyzes workers' cognitive patterns when interacting with AI and provides personalized feedback. When an agent detects a developmental need, it could also trigger a human connection, suggesting a peer, manager, or expert who can help the worker exercise judgment or other cognitive behaviors in context.[14]

Another approach is to use AI not to replace or augment human cognition, but to actively foster it. Instead of providing direct answers, AI can pose questions, prompt reflection, offer alternative perspectives, or simulate different points of view.

One executive we interviewed trained different AI agents to represent C-suite roles and had them debate one another when thinking through a problem.[15] Others suggest designing AI to play a teachable novice, asking clarifying questions to push workers toward explanation and reflection.[16]

Designing a worker experience that strengthens judgment

Historically, higher-order cognitive behaviors (judgment, in particular) have been forged through the work itself: repeated attempts, uneven results, and accountability in owning outcomes when they don't turn out well. AI creates a paradox: the more work it absorbs, the greater the need for human judgment to direct, evaluate, and act on its outputs, yet the fewer opportunities workers have to develop that judgment in the first place.[17]

Standard AI protocols can make this worse. "Human-in-the-loop" principles risk turning workers into passive validators rather than active decision-makers, while escalation procedures can prompt junior employees to defer ambiguous cases to senior colleagues, forfeiting the very decisions that would have sharpened their judgment.

Organizations should deliberately design experiences that build judgment over time. Karen Waddell, vice president of global talent management at Agilent Technologies, says, "Developing judgment, in the face of increased AI use, is my number one priority, as is systems thinking, especially for leaders." Agilent is embedding both into its training programs, but Waddell is clear that formal training alone is insufficient.[18]

Judgment and systems thinking develop through experience, and that means investing in real-world scenarios and simulations that expose people to multiple personas and frames of reference.

Global energy company Enel put this into practice at scale. Led by the chief human resources officer, the organization brought together 1,240 managers across 18 countries and eight business areas over eight months. Sessions alternated between expert-led instruction and hands-on exercises, real-world case studies, and interactive simulations, all designed to build metacognition, judgment, and theory of mind.[19]

The goal is not simply to train people once, but to create repeated opportunities to practice, observe, and strengthen judgment. The experience should also feel psychologically safe, with leaders and managers creating space for workers to question or challenge AI outputs. Worker data collected from AI interactions should not be used for surveillance, but as an investment in their growth.

2. Experimentation: How workers can unlock new value from AI

If cognitive behaviors define how effectively workers use AI in existing workflows, experimentation and agency enable them to discover new ones. AI's most significant value lies in more than the

ability to do the same work faster: it's in giving workers the agency to reshape what work looks like.

But this only happens when workers are willing to test hypotheses, try unfamiliar approaches, tolerate ambiguity, and share what they learn. Organizations that optimize purely for efficiency might systematically underinvest in these behaviors because experimentation looks inefficient in the short term. Yet when 2,422 respondents in a **Deloitte webinar** were asked to identify the single most important action organizations should take to intentionally design human-AI interaction, helping workers learn to experiment with AI ranked first.

Sensing experimentation

Experimentation signals appear at the edges of usage data, not at its core. A worker who regularly tries new AI features, submits feedback about unexpected outputs, or shares an invented workflow with peers is demonstrating experimentation.

The motivation for experimentation is already there: Seventy-eight percent of employees in a Deloitte survey say they are continually experimenting and trying new ways of working. But only 56% say they are recognized or rewarded for being adaptive and agile.[20]

Sensing for experimentation requires looking at the breadth of feature usage alongside its depth; tracking voluntary engagement with learning resources; and monitoring participation in peer communities, idea exchanges, and informal knowledge-sharing. An employee who has only ever used AI for one task—however proficiently—is not yet capturing the compound value that experimentation unlocks.

Organizations should also track how their workers' experiments scale; when worker-developed AI agents are shared beyond a few people, for example, alerts can be set up to ensure the agents are registered and guardrails are in place.

Building a culture of AI experimentation

Experimentation is not a personality trait. It is a behavior that can be designed for, measured, and systematically reinforced—if the system is built to see it.

The most powerful interventions for experimentation and agency are social and contextual, not instructional. When behavioral data surfaces workers using AI in creative ways, the system can amplify their influence, connecting them with peers stuck in narrow usage patterns and inviting them into advisory groups that shape future design. Time-limited sandbox sprints can also reduce the perceived cost of trying something new: Progress is measured not by whether the experiment succeeded, but whether it was undertaken and documented. For example, Canva gave its workforce of more than 5,000 employees a dedicated AI Discovery Week to explore use cases and try a wide range of enterprise AI tools.[21]

For workers who remain in narrow patterns, the intervention should address the underlying barrier. Often, it isn't resistance but uncertainty: Workers don't know what to experiment with, or they fear looking incompetent. Personalized nudges suggesting specific, low-stakes experiments calibrated to their role can lower the activation energy without requiring a formal training session.

One multinational technology company used Cognition's AI-driven tools to drive agentic AI experimentation among its sales workers, comparing the approach against a standard communications-based change management approach across 793 participants. The Cognition cohort showed significantly greater experimentation and adoption that endured over time, driven by daily, real-world AI experimentation challenges, gamification, a virtual coach, guided self-reflection, and social learning via a shared feed where participants posted insights and responded to their peers' discoveries.[22]

Designing a worker experience that encourages experimentation

When the continuous adaptation loop is working, experimentation feels less like a change program and more like a culture. But it's a culture that depends on trust: Workers need to believe experimentation is safe, valued, and not simply a way to help automate their own roles. Leadership behavior shapes the culture as much as any designed intervention. When leaders celebrate experiments that didn't pan out, they signal that experimentation is safe.

The governance challenge is balancing worker agency with top-down guardrails. Agilent's Waddell frames it directly, "We are striving to give freedom to workers within some limits," she says, including educating managers not to reflexively block AI experimentation. The company is establishing an AI Center of Excellence with connections to information technology, human resources, legal, and compliance to provide that structure.[23] Roche's Margaret Greenleaf echoes the point: After an extensive effort to build AI capability at the grassroots level, the next priority is developing human-centric principles and guardrails on human-centric AI from the top. "To be effective," she says, "we need both."[24]

Ensuring AI agents developed by workers are registered and approved when they connect to other systems to scale beyond a few users, for example, can help.

BBVA, a multinational bank headquartered in Madrid, offers a working model: broad worker autonomy to build and experiment, supported by peer AI "champions" and "wizards," within a framework built on one principle set at the top: generative AI is an assistant, not an autonomous agent. Outputs require human validation before connecting to any system, and all users must complete training and formally acknowledge their responsibilities. Recognizing that centralized risk review of thousands of employee-created GPTs was impractical, BBVA built an automated quality score—evaluating tools on guardrails, clarity, and context—as a scalable alternative that encouraged workers to continue to experiment.[25]

3. Divergent thinking: Protecting the human edge in AI work

AI is exceptionally good at synthesis, summarization, and pattern recognition across large data sets. It is not, by design, a source of **divergent thinking** —contrarian perspectives, or the kind of lateral thinking that reframes a problem entirely.

As AI absorbs more of the work associated with linear, convergent tasks, the distinctively human contributions that remain are those that require sustained concentration, intellectual risk-taking, diverse thinking styles, and a willingness to pursue ideas that don't fit the expected pattern.

There is business value in protecting that difference: Deloitte's **2025 survey on technology value** found that organizations fostering innovation-related attributes, including agile thinking and diverse thinking styles, outperformed peers across measures such as average return on key performance indicators, technology return on investment, and return on equity.

The problem is that AI is eroding the conditions that make such thinking possible, including the uninterrupted time to explore new patterns and contrarian perspectives. Among workers experiencing significant AI-driven change, a Deloitte change survey reveals that 69% report their workload has increased, and 43% cite lack of time as their biggest barrier to adapting.[26] As one healthcare executive put it, "We are redesigning work with AI so that clinicians can work at the top of their licenses, but I worry this is not creating enough slack in their roles for them to do the deep, focused work."[27]

AI also creates what some call the "abundance trap": we can now produce insights faster than humans can absorb or act on them. More presentations, more reports, and more drafts, all generated with negligible friction, can multiply faster than teams can interpret or use. The hope was that automating some work would free space for reflective, creative work. But the emerging reality is that the proliferation of content driven by AI can crowd that space out.[28] Research published in *Harvard Business Review* has also found that overseeing complex AI agent systems leads to mental fatigue—"AI brain fry"—undermining the very focus and cognitive depth that divergent and deep thinking requires.[29]

We tend to assume AI will free up time for deeper thinking. In practice, it is often silently eroding it. Protecting space for calm, reflection, imagination, and focus isn't a soft priority but a strategic one.

Sensing divergent thinking

Deep work and divergent thinking leave fewer immediate digital traces than other behaviors, but proxies do exist. For deep work, useful signals include calendar data showing sustained focus blocks, how frequently workers are interrupted, the ratio of AI-generated to human-modified content in final work products, and how many AI agents each worker is overseeing. Micro-surveys, advisory group input, and facilitated reflection sessions can fill in what quantitative data misses.

For divergent thinking, organizations can use collaboration network data to detect whether teams are drawing on genuinely diverse perspectives. AI can analyze how often workers use exploratory language such as "what if" or "another angle," as well as the ratio of exploratory to executional work effort before teams narrow to action.

Building divergent thinking through prompts, space, and friction

Protecting divergent thinking and deep work requires both digital and non-digital interventions. When usage data suggests an employee is relying on AI without meaningful human synthesis (a "prompt-and-accept" pattern), the system can introduce calibrated friction: a peer conversation about alternative framings, a micro-challenge asking for an idea AI would be unlikely to generate, or a nudge to pause before moving forward. AI coaches can also offer metaphors or alternative frames that help workers see a problem differently.

Sometimes the most effective intervention is simply changing the physical environment. Sandboxes, labs, and tech-free collaboration zones all help break normal patterns of work and create the conditions for different thinking.

Designing a worker experience that builds divergent thinking

Deep work requires workers to feel that their distinctively human contributions—the insight that reframed the problem, the judgment call the data couldn't make—are visible, valued, and irreplaceable. When AI is positioned purely as an efficiency engine, it could signal that the human's job is to keep up with the machine. When AI is positioned as a collaborator that expands human capability, the experience shifts.

Leaders can grant explicit permission (and practical space) for the slow, generative thinking that AI cannot replicate. Since deep work is often driven by curiosity untethered from immediate goals, organizations can build slack—time for the unexpected to surface—into work, as well as the psychological safety to explore without fear of looking unproductive. Belgian company DPG Media does this literally, scheduling only 80% of team capacity and reserving the rest as a buffer.[30]

Organizations should be careful not to rush to backfill work recently automated; doing so can feel punitive and rob workers of the time freed for deep work. They can also shift metrics from activity to impact, since incentivizing quantity of work—or even quality of AI use—can lead to waste, lower-quality work, and unnecessary mental strain.

Organizations should also think carefully about spans of control when it comes to managing AI agents. Just as norms exist for managing people, organizations might need similar limits for human-agent oversight, especially where cognitive load and decision quality are at risk.

Measuring the change that matters: AI adaptation, not just AI adoption

The shift from adoption to adaptation isn't a rejection of measurement but a demand for better measurement. Adoption metrics aren't necessarily wrong; they might be incomplete. An organization where no one uses AI has no foundation for transformation. But an organization where everyone uses AI without developing the behaviors that make it valuable has built that foundation on sand.

Creating a continuous adaptation loop offers a different approach to change: one where change is embedded in the flow of work and behavioral signals, not usage counts, drive the system. Sensing distinguishes the skeptic from the over-reliant. Engagement meets people where they are. Experience design builds capability, not just compliance.

Adoption asks: Did they use it? Adaptation asks: Did they use it effectively to drive impact? The first question is answerable with a dashboard. The second requires a new set of AI-native metrics, built for a world where value is created through human and machine working together.

The organizations that lead in AI won't necessarily be those with the highest adoption rates driven by traditional change frameworks. They will likely be those that have used AI to build something more durable: the creation of an adaptive environment that supports the workforce in thinking with higher-order cognitive skills, experimenting fearlessly, and doing the deep work that no large language model can replicate.

That's not an adoption story. That's a transformation.

Continue the conversation

Meet the industry leaders



Allyson Dake
Managing Director | U.S. Strategy & Technology Change Market Offering Leader | Academy for Change Leader



Chloe Domergue
Principal | Deloitte Consulting LLP

BY

Sue Cantrell
United States

Chloe Domergue
United States

Allyson Dake
United States

Jessica Murphy
United States

Tori Sundholm
United States

Mark Gustafson
United States

ENDNOTES

1. Jim Rowan, Beena Amanath, Nitin Mittal, and Costi Perricos, "**The state of AI in the enterprise: The untapped edge**," Deloitte AI Institute, January 2026.

2. Nicole Scoble-Williams, Sue Cantrell, David Mallon, and Stefano Besana, "**Getting human and machine relationships right**," *Deloitte Insights,* March 4, 2026.

3. Deloitte Human Capital Trends Survey in collaboration with Oxford Economics, 2025.

4. Thomas Stackpole, "**LLMs are manipulating users with rhetorical tricks**," *Harvard Business Review,* March 18, 2026.

5. Stefan Buijsman, Sarah E. Carter, and Juan-Pablo Bermúdez, "**Autonomy by design: Preserving human autonomy in AI decision-support**," *Philosophy & Technology* 38, no. 97 (2025).

6. Margaret Greenleaf (chapter lead for people-support solutions, Roche), interview with Sue Cantrell, April 2026.

7. Cynthia Kaschub (senior director of workforce innovation, Salesforce), interview with Sue Cantrell, April 2026.

8. Fabrizio Dell'Acqua et al., "**Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality**," *Organization Science* 37, no. 2 (2026).

9. Shuhua Sun et al., "**How and for whom using generative AI affects creativity: A field experiment**," *Journal of Applied Psychology* 110, no. 12 (2025).

10. Hassan Qudrat-Ullah, *Navigating Complexity: AI and Systems Thinking for Smarter Decisions* (Springer, 2025).

11. Christoph Riedl and Ben Weidmann, "**Quantifying human-AI synergy**," working paper shared on PsyArXiv, Sept. 21, 2025.

12. Joshua Holstein and Gerhard Satzger, "**Development of mental models in human-AI collaboration: A conceptual framework**," Cornell University research paper, Oct. 9, 2025.

13. Alex Farach, Alexia Cambon, Lev Tankelevitch, Connie Hsueh, and Rebecca Janssen, "**Scaffolding human–AI collaboration: A field experiment on behavioral protocols and cognitive reframing**," Microsoft Corporation, arxiv, April 9, 2026.

14. Deloitte Workforce Innovation Forum, April 2026.

15. Deloitte client interview with the authors, 2025.

16. Hayato Tomisu, Junya Ueda, and Tsukasa Yamanaka, "The cognitive mirror: A framework for AI-powered metacognition and self-regulated learning," *Frontiers in Education* 10, Oct. 8, 2025.

17. David Mallon, Sue Cantrell, and John Forsythe, "Closing the experience gap," *Deloitte Insights*, March 24, 2025.

18. Karen Waddell (vice president of global talent management, Agilent Technologies), interview with Sue Cantrell, April 2026.

19. Deloitte client work, 2026.

20. Deloitte change pulse survey, February 2026.

21. Rob Giglio, "We gave our 5,000 employees a week to do nothing but learn AI. We learned the biggest blockers are human ones," *Fortune*, May 28, 2026.

22. Cognician, "Agent activation experiment: Microsoft UK & Cognician," accessed June 29, 2026.

23. Karen Waddell (vice president of global talent management, Agilent Technologies), interview with Sue Cantrell, April 2026.

24. Margaret Greenleaf (chapter lead for people-support solutions, Roche), interview with Sue Cantrell, April 2026.

25. Elena Alfaro et al., "The hidden demand for AI inside your company," *Harvard Business Review,* April 14, 2026.

26. Deloitte change pulse survey, February 2026.

27. Deloitte client interview, April 2026.

28. Lynda Gratton, "AI is changing how we learn at work," *Harvard Business Review,* Dec. 22, 2025.

29. Matthew Kropp, Megan Hsu, Olivia T. Karaman, Jason Hawes, and Gabriella Rosen Kellerman, "When using AI leads to 'brain fry,'" *Harvard Business Review,* March 5, 2026.

30. DPG Media, "Why we love to slack off," accessed Feb. 3, 2025.

ACKNOWLEDGMENTS

The authors would like to thank **Michael Griffiths** and **Farah Naim** of Deloitte for their continued support and contributions to our practice and research.

We would also like to thank **Katherine McNally**, **Malia Maack**, **Julia Guttery**, and **Naomi Ray** of Deloitte for their contributions to this article.

Editorial (including production and copyediting): **Corrie Commisso**, **Sayanika Bordoloi**, **Annalyn Kurtz**, and **Anu Augustine**

Design: **Rahul Bodiga** and **Molly Piersol**

Audience development: **Pooja Boopathy**

Cover image by: **Rahul Bodiga**

Knowledge services: **Agni Wagh**

COPYRIGHT

Copyright © 2026 Deloitte Development LLC. All rights reserved. Member of Deloitte Touche Tohmatsu Limited