THE
kyndryl
institute

Trend topic: **People readiness**

By

Mohammad Hossein Jarrahi

Professor of Information Science at the University of North Carolina at Chapel Hill, Researcher and Consultant of AI and Future of Work

Mohammad Hossein Jarrahi is Professor of Information Science at the University of North Carolina at Chapel Hill and a leading scholar of human-AI collaboration, algorithmic management, and the future of work. He is widely recognized for his foundational work on the concept of human-AI symbiosis, with this body of research receiving more than 4,000 citations. His research examines how organizations can use AI to strengthen human judgment and capability in emerging hybrid structures of collaboration between humans and AI systems. In addition to his academic contributions, Jarrahi regularly engages executive audiences through invited sessions and talks with organizations such as the Food and Agriculture Organization of the United Nations and the Asian Development Bank, as well as discussions with leaders in the financial sector. He also translates research for practitioner and leadership audiences through publications in Harvard Business Review, MIT Sloan Management Review, California Management Review, and Business Horizons.

Explore more from The Kyndryl Institute
kyndryl.com/institute

Designing for human-AI symbiosis: Why collaboration is not enough

Leaders are moving from AI collaboration to human-AI symbiosis, emphasizing capability building, reciprocal learning, and coordinated orchestration to improve decision quality, strengthen trust, and build long-term organizational resilience.

It is clear that organizations now understand the importance of human-AI collaboration. The debate has moved from using AI purely to save costs or time to focusing on how it can be harnessed to create and innovate. According to the Kyndryl People Readiness Report 2026[1], 31% of leaders surveyed said they are intentionally deploying AI to create new capabilities people and AI could not create separately, compared to 8% of leaders deploying AI to cut head count. However, I would argue that the leaders pulling ahead are designing for something more strategic and demanding: human-AI symbiosis.

The cost of stopping at collaboration is already visible. Over the past couple of years, professionals across law, consulting, science, and finance have submitted AI-assisted work products containing AI slop,[2] or, worse, embarrassing instances of AI hallucination.[3] Legal briefs have cited fictional cases. Scientific reviews have summarized fabricated studies. Consulting analyses have misread data or presented phantom references. These are rarely stories of AI failure on its own. They are stories of human-AI collaboration that produces fast output, but not trusted value.

A more productive vision is human-AI symbiosis:[4] an arrangement in which humans and AI develop complementary capabilities over time. Each strengthens the other. And the unit of knowledge production becomes the joint human-AI system rather than either party alone. An example of this comes from the implementation of generative AI in customer support[5] at a Fortune 500 software firm, where the AI system learned from the practices of high-performing employees and then offered real-time response suggestions and documentation links while humans retained discretion to accept, edit, or ignore them. Not only did this improve customer service, it also aided in improving the performance of less experienced employees.

> The majority (94%) says that the organizations that will win with AI are those that invest in judgement from people, not just AI capacity.

Symbiosis as a practical vision

The framework of human-AI symbiosis[6] begins from a simple observation: humans and AI bring different strengths, and different weaknesses, to the work itself. Humans bring contextual judgment, ethical reasoning, tacit knowledge, and the ability to revise goals in response to a changing environment. AI brings scale, speed, pattern recognition across vast data, and tireless consistency. The opportunity is not to choose between them, or to use one to suppress or necessarily surpass the other. It is to build a working system in which each compensates for the other's limits and amplifies the other's strengths, a component of human-AI symbiosis often described as hybrid intelligence.[7]

Encouragingly, leaders recognize this. The majority (94%)[8] says that the organizations that will win with AI are those that invest in judgement from people, not just AI capacity. Indeed, half of leaders said people will remain the primary source of value over the next three years when it comes to decision-making in complex or ambiguous situations, as well as customer and stakeholder relations. And yet, leaders are clearly struggling to operationalize this new way of working. Only 22% of leaders report redesigning workflows to integrate AI as a collaborative partner, a gap that suggests most organizations are layering AI onto legacy structures rather than redesigning the work itself. And more than double this (45%) said redesigning roles and workflows around people and AI collaboration will be the biggest barrier to successfully deploying AI (the second biggest barrier after cybersecurity).[9]

How can leaders address this? There are three deliberate shifts they can make:

Shift 1: From automation to capability building

AI should not only save time. It should help workers reason better, learn faster, and make better decisions.

Treating AI purely as automation could result in what I call **disaugmentation**. This is the slow erosion of judgment, expertise, and discretion that occurs when AI is layered onto organizations without redesigning the work, the people, and the structures around it. Disaugmentation manifests itself, for instance, when junior consultants learn to polish AI output but not to build an argument, or when customer service employees follow suggested responses without developing the capability to handle complex cases. Disaugmentation is what AI-era leaders should be most worried about, especially as it is largely invisible on quarterly dashboards and on efficiency-driven metrics focused on cost saving. The danger is that disaugmentation can look like progress at first. Output may become faster and cheaper, while the organization quietly weakens the human expertise needed to integrate or question AI outputs, handle edge cases, train future employees, and recover when the system fails. By the time the problem becomes visible, the organization may have fewer people left with the judgment required to take corrective actions.

Treating AI purely as automation could result in what I call disaugmentation. This is the slow erosion of judgment, expertise, and discretion that occurs when AI is layered onto organizations without redesigning the work, the people, and the structures around it.

Capability-building design starts from a different question: what does this workflow do to the people who run it? It treats AI as an investment in future expertise,[10] not just a substitute for current expertise. The key question here is whether the increasing integration of AI leaves workers more capable of reasoning, judging, and solving problems over time, or merely faster at producing outputs.

In practice, this means a different geometry of work. Some tasks should remain human-led by design, to preserve the development of judgement people will need five years from now. Some should run in parallel, with humans and AI working the same problem so comparison itself becomes a learning surface. Others should be sequenced deliberately. The ScreenTrustCAD breast-screening trial[11] in Sweden is a useful parallel workflow: two radiologists and an AI system independently read the same mammograms, without seeing one another's judgments; cases flagged by any reader then went to consensus discussion. The workflow made human-AI disagreement visible, turning comparison itself into a surface for review, calibration and learning.

The deeper point is that AI augmentation and human development are not separate budget lines. In fact, they are the same.

Shift 2: From oversight to reciprocal learning

Capability building in shift 1 facilitates moving from oversight to reciprocal learning. Humans should not only approve AI outputs. They should provide feedback, context, correction, and judgment that improve both AI performance and their own.

Models such as human-in-the-loop remain useful for accountability and risk control, but they are too narrow. They implicitly cast people as reviewers, approvers, or exception handlers. Symbiosis demands a more bidirectional flow: AI improves what humans can do, and humans actively shape how AI is used, corrected, and contextualized within the organization.

This reframing changes how trust gets built. We have seen two opposing failure modes in AI applications: cognitive complacency, where people defer to AI when they shouldn't, as when professional recruiters over-rely on AI résumé-screening[12] recommendations rather than independently evaluating candidates; and algorithmic aversion, where they may reject AI output that could help, as when forecasters abandon algorithmic predictions[13] after seeing them make mistakes.

Both are forms of broken reciprocity and miscalibrated trust. In the first, humans stop learning from the system; in the second, the system stops learning from humans. Calibrated trust, knowing when to defer, when to override, and when to investigate — is built through reciprocal learning. Key to establishing any kind of trust when using AI, however, is clear and effective governance, and this remains a serious challenge. Less than a quarter of leaders said their governance and compliance framework is currently ready to successfully leverage AI.[14] Leaders need to give their employees a chance to even get out of the gate when it comes to confidently and effectively using AI.

Shift 3: From projects to symbiotic enterprise orchestration

Reciprocal learning in shift two facilitates moving from projects to symbiotic enterprise orchestration. AI should not only be deployed project by project. It should be orchestrated across the enterprise, spanning agents, humans, and external stakeholders.

Most AI investments today are scoped to specific products, processes, or business units. That scope is convenient for ROI calculations, but inadequate for symbiosis. Real value, and real risk, lives at the seams: In how AI agents interact with clients, partners, and suppliers; in how decisions made by one AI system propagate into another; and in how workflow changes in one division silently degrade work in an adjacent one.

Symbiotic enterprise orchestration treats AI deployment as a portfolio of interacting agents and humans,[15] rather than a series of independent automation tasks. It also treats stakeholder asymmetries as a first-order problem. When AI is introduced into a multi-party process, parties with less visibility or power tend to absorb the costs of poor design while others capture the gains. Holistic redesign

asks who exactly benefits, who is at a disadvantage, and whether those asymmetries are acceptable to the organization.

The diagnostic question for executives is therefore not "Does this AI deployment save money?" It is "Does this design produce a more capable organization a year from now, across all the people and partners it touches?"

This also changes what leaders measure. Usage, speed, and cost savings matter, but they are incomplete. Symbiotic orchestration requires metrics for decision quality, skill development, escalation quality, accountability, and whether AI use is making the organization more resilient or more dependent over time.

The diagnostic question for executives is therefore not "Does this AI deployment save money?" It is "Does this design produce a more capable organization a year from now, across all the people and partners it touches?"

Designing for symbiosis

Most organizations were designed for a world in which humans did the work and digital technology served as infrastructure. The rise of AI breaks that assumption in both directions. AI does not sit politely behind the workflow, and it cannot be treated exactly like a human either, a trap researchers often describe as anthropomorphism.[16] Symbiosis is the design paradigm that takes the asymmetry between humans and AI seriously, takes advantage of it through careful organizational design, and avoids collapsing into techno-optimism or human-replacement narratives.

The leaders ready for the next phase are not those with the largest AI budgets. They are the ones asking harder questions: whose judgment am I protecting? Whose skills am I developing? Whose accountability am I preserving? Before scaling an AI system, leaders should ask: What human capability will this strengthen? What expertise might it weaken? Where will correction and feedback enter the workflow? And how will we know whether the organization is becoming more capable, not merely faster?

References

1 People Readiness Report, Kyndryl, 2026

2 https://hbr.org/2025/09/ai-generated-workslop-is-destroying-productivity

3 https://www.ibm.com/think/topics/ai-hallucinations

4 https://humansplus.ai/podcast/mohammad-hossein-jarrahi-human-ai-symbiosis-intertwined-automation-augmentation-race-machine-tacit-knowledge-ac-ep62/

5 https://academic.oup.com/qje/article/140/2/889/7990658?guestAccessKey=

6 https://www.researchgate.net/publication/320812037_Artificial_Intelligence_and_the_Future_of_Work_Human-AI_Symbiosis_in_Organizational_Decision_Making

7 https://journals.sagepub.com/doi/10.1177/20539517221142824

8 People Readiness Report, Kyndryl, 2026

9 Ibid

10 https://mitsloan.mit.edu/ideas-made-to-matter/how-can-we-preserve-human-ability-age-machines

11 https://pubs.rsna.org/doi/10.1148/radiol.242566

12 https://anacanhoto.com/wp-content/uploads/2024/08/554ee-fallingasleepatthewheel-fabriziodellacqua.pdf

13 https://psycnet.apa.org/buy/2014-48748-001

14 People Readiness Report, Kyndryl, 2026

15 https://cmr.berkeley.edu/2025/07/rethinking-ai-agents-a-principal-agent-perspective/

16 https://link.springer.com/article/10.1007/s43681-024-00419-4