

Beyond Human-in-the-Loop

May 17, 2026



Image credit: Rik Oostenbroek - @RikOostenbroek

AI coordination tools could help organisations align, deliberate, and collaborate more effectively. But they also change the trust signals and team dynamics that human cooperation depends on.

The leadership challenge is therefore not simply to deploy AI wherever it saves time, or to insist that every AI process has a human-in-the-loop. It is to deliberately design where humans, AI systems, and hybrid teams each add the most value.

Google DeepMind's *Organising Intelligence* frames this problem well in Idea 9, "AI-driven Coordination," where it argues that research on AI collaboration is coalescing around three questions: whether AI has the social-cognitive capabilities to be a good collaborator, whether AI helps or harms human coordination, and whether human-AI hybrid teams outperform human-only teams.

The core argument of this article is that AI's organisational value lies not only in automation, but in the redesign of cooperation itself.

From automation to coordination

Most executive conversations about AI still begin with productivity: which tasks can be automated, how much time can be saved, and which workflows can be accelerated. That framing is useful, but incomplete. Organisations do not only fail because tasks are slow. They fail because teams misunderstand one another, incentives are misaligned, knowledge is trapped in silos, and disagreement hardens before it is productively surfaced.

This is why AI's deeper organisational promise may be coordination rather than automation. Sangeet Paul Choudary captures this shift in his Harvard Business Review argument that AI's large payoff may come from reducing coordination costs rather than merely replacing work (Choudary, 2026). In this view, AI becomes a translation layer between teams, a mediator across disagreement, and a mechanism for making implicit assumptions explicit.

Recent research supports this broader coordination lens. In democratic deliberation, AI systems have been shown to help people find statements of common ground that participants collectively endorse (Tessler et al., 2024). Other work suggests AI sources can increase openness to opposing views (Lu, Tormala, & Duhachek, 2025). In organisational settings, this points to a powerful possibility: AI may help groups move beyond positional argument toward structured comparison of assumptions, evidence, and trade-offs.

AI can also enrich how organisations read their own social systems. Work by Amir Goldberg and Sameer Srivastava argues that AI can deepen our understanding of organisational culture by detecting patterns in language, meaning, and interaction that are otherwise difficult to observe at scale (Goldberg & Srivastava, 2024). This matters because coordination is rarely just a formal process. It is shaped by culture: what people say, what they avoid saying, which groups translate easily across boundaries, and where shared language breaks down.

The social-cognitive question

For AI to support coordination, it must do more than retrieve information. It must participate in the social fabric of work. That requires some ability to reason about beliefs, intentions, emotions, and perspectives.

Research on theory of mind in large language models is therefore relevant. Street and colleagues find that LLMs can perform at adult-human levels on higher-order theory-of-mind tasks, which involve reasoning recursively about what other people know, believe, or infer (Street et al., 2024). Related work on empathic communication suggests that large language models may be reliable judges of empathic communication under certain conditions (Kumar et al., 2026).

But this should not be read as proof that AI “understands” organisations in a human sense. Social meaning is culturally situated. Research on facial expressions across cultures shows how complex and context-dependent the interpretation of emotional signals can be (Brooks et al., 2024). The leadership implication is that AI coordination tools should be treated as powerful but partial social instruments: useful for surfacing perspectives, summarising disagreement, and supporting empathy, but not as definitive interpreters of human meaning.

The hidden trust problem

The optimistic case for AI coordination is strong. But cooperation is not only an information problem. It is also a trust problem.

When people collaborate, they rely on visible signals of effort, judgment, and intent. We often decide whether to trust a colleague not only by the quality of their answer, but by watching how they arrived at it: Did they wrestle with the problem? Did they consider the trade-offs? Did they understand the consequences? Did they expose their assumptions?

Wojtowicz and DeDeo call this “mental proof”: the observable evidence of thinking that allows others to infer otherwise invisible traits such as skill, care, and integrity (Wojtowicz & DeDeo, 2024). AI can undermine this proof by making thinking easier, faster, and less visible. A polished AI-assisted answer may be useful, but it may also obscure whether the person presenting it has exercised judgment.

This is the paradox at the heart of AI coordination. AI can make collaboration easier at the task level while making cooperation harder at the social level. It can reduce friction in producing outputs while removing the visible struggle that helps people trust one another.

That does not mean organisations should avoid AI, quite the opposite. It means leaders need to design for trust signals deliberately. AI-supported work may need new forms of provenance: visible assumptions, audit trails, dissent logs, confidence levels, source trails, and explicit separation between human judgment and AI-generated synthesis. The point is not to slow the organisation down for its own sake. It is to preserve the social evidence that collaboration depends on.

Hybrid teams are not automatically better

The phrase “human-in-the-loop” is often treated as a universal safety principle. It sounds reassuring: AI can act, but humans will remain involved. Yet the evidence suggests that human involvement is not automatically additive.

A systematic review and meta-analysis of human–AI combinations found that hybrid teams often outperform human-only teams, but not always. In some cases, especially where the AI system already performs better than humans on a given task, adding a human can reduce performance relative to AI alone (Vaccaro, Almaatouq, & Malone, 2024). This complicates the simplistic governance answer that every important process should include a human checkpoint.

The right question is not “Should there be a human in the loop?” The better question is “What is the human there to do?”

Sometimes, the human should decide. Sometimes, the human should audit. Sometimes, the human should set goals, values, and constraints. Sometimes, the human should monitor for exceptions. Sometimes the AI should generate options, critique assumptions, or translate between expert groups. And in some low-risk, well-bounded contexts, the AI may be best positioned to execute without continual human interruption.

This is where the literature on AI teaming becomes important. Schmutz and colleagues argue that AI teaming requires rethinking collaboration itself in the digital era, including roles, interdependence, trust, and team design (Schmutz et al., 2024). AI is not merely a tool dropped into an existing team. It changes the team.

The Task Tensor: a richer design language

My work with Anil Doshi on the “Human–AI Task Tensor” offers a useful way to move beyond crude human-in-the-loop thinking. The framework organises human–AI work across eight dimensions: task definition, AI integration, interaction modality, audit requirement, output definition, decision-making authority, AI structure, and human persona.

This matters because different work requires different human–AI configurations. A high-volume, well-defined, easily auditable task may justify substantial automation. A strategic, ambiguous, politically sensitive task may require AI support but human decision authority. A creative or deliberative process may benefit from AI-generated alternatives, but still require humans to preserve accountability, interpret context, and negotiate meaning.

We outline 20 levels of decision authority in human–AI dyads. It distinguishes, for example, between a human giving approval “in the loop,” a human holding veto authority “on the loop,” a human observing “near the loop,” and progressively more autonomous configurations.

That distinction is crucial. **Human-on-the-loop** is not the same as **Human-in-the-loop**. In a human-in-the-loop process, the human is part of the active decision cycle. In a human-on-the-loop process, the AI may act with greater autonomy while the human supervises, intervenes, or vetoes when necessary. This gives leaders a richer design vocabulary for matching oversight to task risk, auditability, performance, and accountability.

What leaders should design for

The practical lesson is that AI coordination should be designed, not improvised.

First, leaders should map tasks rather than jobs. A single role may contain tasks that require different human–AI configurations. Some tasks are suitable for automation; others need augmentation; others require deliberation, contestation, and human accountability.

Second, leaders should define the AI’s coordination role. Is the AI acting as a translator between functions, a mediator between disagreeing groups, a critic of assumptions, a recorder of decisions, a recommender of options, or an autonomous executor? Each role creates different risks and trust requirements.

Third, leaders should preserve mental proof. When AI helps generate an output, teams should still be able to see the reasoning process: assumptions, alternatives considered, sources used, judgments made, and points of uncertainty. In high-trust teams, AI may accelerate collaboration. In low-trust teams, invisible AI assistance may make people more suspicious unless the process is made legible.

Fourth, leaders should treat hybrid teams as teams. That means designing roles, handoffs, incentives, escalation paths, and accountability norms. The question is not whether AI is present. The question is whether the human–AI system is coordinated well.

Finally, leaders should resist the false binary between automation and human control. The Task Tensor points toward a more nuanced architecture: humans may define goals, AI may generate options, humans may deliberate, AI may simulate consequences, humans may decide, and AI may monitor execution. In other cases, AI may decide within boundaries while humans supervise on the loop.

Conclusion

The next generation of AI-enabled organisations will not be the ones that automate the most decisions. They will be the ones who best design the division of cognitive labour between humans, AI systems, and hybrid teams.

AI coordination tools can help organisations deliberate, align, translate, and collaborate. But they also reshape the social evidence through which people decide whether others are competent, sincere, and trustworthy. Used naively, AI may make work faster while weakening cooperation. Used thoughtfully, it can become a new coordination architecture: one that improves collective judgment while preserving the human signals that make collaboration possible.

The goal is not faster work by default. The goal is better-designed cooperation.

References

- Brooks, J. A., et al. (2024). Deep learning reveals what facial expressions mean to people in different cultures. *iScience*, 27(3), 109175.
- Choudary, S. P. (2026). AI's big payoff is coordination, not automation. *Harvard Business Review*.
- Doshi, A. R., & Moore, A. P. (2026). Toward a Human–AI Task Tensor: A Taxonomy for Organizing Work in the Age of Generative AI. In *Handbook of Artificial Intelligence and Strategy*. Open version: SSRN.
- Goldberg, A., & Srivastava, S. B. (2024). How Artificial Intelligence Can Enrich Our Understanding of Organizational Culture. *Management and Business Review*, 4(2), 32–37.
- Google DeepMind. (2026). *Organizing Intelligence: 16 big ideas from frontier AI research that will redefine how we build, lead, and scale*.
- Kumar, A., et al. (2026). When large language models are reliable for judging empathic communication. *Nature Machine Intelligence*, 8(2), 173–185.
- Lu, L., Tormala, Z. L., & Duhachek, A. (2025). How AI sources can increase openness to opposing views. *Scientific Reports*, 15, 17170.
- Schmutz, J. B., et al. (2024). AI-teaming: Redefining collaboration in the digital era. *Current Opinion in Psychology*, 58, 101837.
- Street, W., et al. (2024). LLMs achieve adult human performance on higher-order theory of mind tasks. *arXiv*.
- Tessler, M. H., et al. (2024). AI can help humans find common ground in democratic deliberation. *Science*, 386(6719), eadq2852.
- Vaccaro, M., Almaatouq, A., & Malone, T. W. (2024). When combinations of humans and AI are useful: A systematic review and meta-analysis. *Nature Human Behaviour*, 8(12), 2293–2303.
- Wojtowicz, Z., & DeDeo, S. (2024). Undermining mental proof: How AI can make cooperation harder by making thinking easier. *arXiv*.