

The Space Between: Team Cognition, Human-AI Dyads, and the Future of Collective Intelligence

Clare Dygert | 2026

“Tell me, what is it you plan to do with your one wild and precious life?”

— Mary Oliver

The Moment We’re In

Something unprecedented is happening to knowledge work, and we’re mostly having the wrong conversation about it.

AI is absorbing an increasing share of routine cognitive labor. The drafting, the summarizing, the analyzing, the generating, the coordinating. Not in some distant future. Now. Across every discipline: engineering, law, medicine, finance, design, research, learning and development. The proportion of knowledge work that requires a human being is shrinking, and it’s shrinking fast.

Most of the conversation about this shift falls into one of three ruts. There are the efficiency enthusiasts who want to use AI to do the same work faster and cheaper. There are the people who are genuinely afraid, worried about losing their jobs, their expertise, their sense of professional identity. And there are the people who’ve decided AI itself is the problem, a force causing disruption and harm that needs to be contained. All three groups are focused on AI as a thing to be managed. None of them are asking the more interesting question.

If AI handles the routine cognitive labor, what becomes possible for the humans?

For the first time in the history of knowledge work, we have the opportunity to focus our energy on the problems that actually require human cognition. The wicked problems. Rittel and Webber (1973) gave us that term to describe challenges that resist straightforward solutions: problems with no definitive formulation, where every attempt at a solution changes the problem itself, and where the causes are so tangled that no single perspective can see the whole. Climate change. Healthcare systems. Educational equity. Organizational transformation. These are the problems

that have always deserved our best thinking but rarely gotten it, because we've been too busy with everything else.

Now the "everything else" is migrating to AI. We have a choice about what to do with the capacity that frees up.

But wicked problems can't be solved by individuals, no matter how brilliant. They require collective intelligence. They require people thinking together, not just working alongside each other. They require team cognition. And most organizations have no idea how to build it.

There's another dimension to this that should concern anyone responsible for the future of an organization. Knowledge work has always had a natural pipeline: you enter at the bottom, doing the routine tasks, and you develop expertise over time through practice and mentorship and increasing responsibility. If AI absorbs that bottom layer, it blocks the entry point. The training ground disappears. Where do the next generation of experts come from? As I'll describe later, the team structure offers an answer to this question, but only if the team is actually thinking together.

One more thing before we proceed. If we evaluate the future of work using the mental models we've developed in the current world of work, we will draw the wrong conclusions. The nature of the work is changing, and our frameworks for understanding it need to change with it. This paper asks the reader to hold that ambiguity and to consider the possibility that the answers might not look like what we expect.

The question Mary Oliver posed was about a single life. The question we face today is the team version: we have this extraordinary human capacity, creativity, judgment, imagination, moral reasoning, and for the first time, we have the freedom to aim it at the wicked problems that matter. What will we do with it?

What Team Cognition Is

Most people hear "team cognition" and think I'm talking about teamwork. I'm not.

Teamwork is coordination: who does what, when, how the pieces fit together. It's essential. It's also not what I'm describing. Team cognition is emergent. It means the team produces thinking that no individual member could produce alone. Not just better thinking, but fundamentally **different** thinking, the kind that wicked problems actually require: ideas and insights and solutions that didn't exist in any of the individual minds and couldn't have been generated by any of them working independently.

Team cognition happens in what I call the liminal space, the threshold between one mind and another. Liminal space is where creativity and innovation actually live. It's unstable, generative, and a little uncomfortable. It's the place where your half-formed idea meets my different

perspective and something neither of us expected comes into existence. That's not coordination. That's emergence.

Vygotsky (1978) gave us a way to understand why this space is so productive. His zone of proximal development describes the gap between what a person can do alone and what they can do with the support of others, the space where learning and growth happen precisely because you're reaching beyond your individual capacity. The liminal space in a team is a collective version of this: the zone between what any individual mind can produce and what becomes possible when different minds meet at the threshold.

The research literature identifies three foundational constructs that make team cognition possible. They develop over time through shared experience and can't be produced by a training event or a workshop. But they can be intentionally cultivated, and the distinction between those two statements matters enormously.

Mental Models: The Lens

A mental model is the lens through which a person sees the world. It's how you predict what will happen next, how you interpret new information, how you decide what matters and what doesn't. You're using mental models right now, reading this paper. They're shaping what you notice, what you agree with, what you resist.

Mental models are formed by experience plus reflection. They don't change because someone presents a better argument or gives you new information. They change when lived experience creates enough of a mismatch with the existing model that the model has to update. And that update requires reflection, some space to make sense of what the experience meant. Cranton (2006), building on Mezirow's (1991) theory of transformative learning, describes this as critical self-reflection triggered by a disorienting dilemma, an experience that simply doesn't fit the existing frame. The frame has to shift to accommodate the new reality. You can inform someone all day long. Transformation requires experience.

A shared mental model on a team means the members' lenses are compatible enough that they can anticipate each other, interpret situations similarly, and work from a common understanding of reality without renegotiating every interaction from scratch. When shared mental models are strong, the team flows. When they're weak, every conversation is a negotiation (Cannon-Bowers, Salas, & Converse, 1993; Mathieu, Heffner, Goodwin, Salas, & Cannon-Bowers, 2000).

Transactive Memory: The Directory

A transactive memory system is the team's shared understanding of who knows what (Wegner, 1987). Not everyone needs to know everything. But everyone needs to know where knowledge lives in the team and how to access it.

Think about the best team you've ever been on. You probably knew, without thinking about it, who to go to for different kinds of problems. You knew who had the technical depth, who had the client relationships, who could see the political landscape, who would catch the detail you'd miss. That was your transactive memory system at work. The team's total knowledge was far greater than any individual's, but only because the directory was accurate and people trusted it enough to rely on each other (Lewis, 2003; Hollingshead, 1998).

Transactive memory takes time to develop. It grows through shared experience: working together, seeing what each person contributes, learning who knows what and how they think.

Team Situational Awareness: The Pulse

Team situational awareness is the team's shared, real-time understanding of what's happening, what it means, and what's likely to happen next (Endsley, 1995). It's what separates a team that can handle complexity from one that falls apart when conditions change. It depends on the other two constructs: you need shared mental models to interpret what's happening the same way, and you need transactive memory to know who to turn to when something unexpected comes up (Salas, Prince, Baker, & Shrestha, 1995).

How They Work Together

These three constructs are deeply intertwined. Shared mental models make transactive memory possible. Transactive memory supports situational awareness. And all three develop through the same mechanism: shared experience over time.

This is why team cognition can't be produced by a workshop or a training module. You can't explain these constructs into existence. They have to be grown through the lived experience of working together, reflecting together, and building a shared history. But team cognition *can* be intentionally cultivated. Not by training it directly, but by designing the conditions under which it naturally develops. The distinction is between installing a capability and preparing the soil for one to grow.

The Conditions

Team cognition doesn't just happen because you put smart people in a room. I've seen rooms full of brilliant individuals produce nothing but frustration and mediocre output. I've also seen teams of ordinary people produce work that genuinely surprised everyone, including themselves. The difference isn't the people. It's the conditions.

Cognitive Diversity

Teams need members who think differently. This begins with demographic diversity: differences in race, gender, gender expression, culture, age, neurodiversity, and life experience, which is valuable in its own right and also produces cognitive diversity because people with different

lived experiences develop different lenses (Woolley, Chabris, Pentland, Hashmi, & Malone, 2010; Page, 2007). Cognitive diversity extends beyond demographics to include differences in disciplinary training, problem-solving approaches, and professional background. It's worth noting that formal education is only one path to developing a valuable lens. What matters is curiosity, thoughtfulness, wide reading, and genuine engagement with perspectives different from your own. A person without a degree who has spent years learning through practice, conversation, and self-directed study may bring exactly the cognitive diversity a team needs. A team needs richness from who people are and range from how they think, and neither is the exclusive product of a credential.

Homogeneous teams converge too quickly. They feel efficient, but they miss things. The liminal space collapses because everyone's already in the same place. There's no threshold to cross, no productive collision of different lenses, and nothing genuinely new emerges.

Cognitive diversity also keeps the talent pipeline open. A team with room for the apprentice alongside the master, someone at the beginning of their expertise who brings fresh eyes and learns by participating in the team's thinking, is a team that's developing its own future (Lave & Wenger, 1991).

Trust Architecture

Every team operates within what I'm calling a trust architecture: the set of norms, structures, and boundaries that determine what knowledge flows where, and who controls that flow. This builds on Amy Edmondson's foundational research on psychological safety (Edmondson, 1999), which demonstrated that the most important predictor of team performance is whether people feel safe taking interpersonal risks. Google's Project Aristotle confirmed this at scale (Duhigg, 2016).

I want to extend the concept. Psychological safety describes the climate, how it feels to be on the team. Trust architecture describes the infrastructure that supports or undermines that climate. It's the designed element: who has access to what information, how feedback flows, what's private and what's shared, and who controls those boundaries.

Good trust architecture makes it safe to think out loud, to share half-formed ideas, to admit what you don't know, to surface problems early. This is where the real cognitive work happens, in the messy middle of sensemaking, not in polished deliverables. The liminal space only works if people feel safe being in it.

Bad trust architecture drives cognition underground. Surveillance, punishment for mistakes, forced transparency that serves management control rather than team learning: these cause people to retreat to safe, defensible outputs. The team gets dumber because the best thinking never enters the room.

A critical element of trust architecture is the commitment to explicit communication. Teams that rely on implicit communication, on assumptions, hints, and unspoken expectations, are building on sand. This matters for every team and matters especially when neurodiversity is present. Explicit communication means saying what you mean, making expectations visible, and giving feedback based on shared, explicit standards rather than subjective opinion. When a team has agreed on clear standards for its work, feedback becomes about the work measured against those standards, not about one person's reaction to another person's output. That depersonalization makes honest, continuous feedback possible without threatening relationships.

Clear Agreements

Teams need to know how to make agreements with each other and how to hold them. Most teams skip this. They assume alignment rather than building it. They start working before they've clarified who's responsible for what, how decisions will be made, what the norms are for communication, and how they'll handle it when things go wrong.

Trust is built on a track record of commitments made and kept. Agreement-making is a competency, a skill that needs to be developed and practiced. Teams that can't make clear agreements will never build the trust architecture they need.

Equally important is the ability to renegotiate. Rigid agreements become brittle. Living agreements adapt. A team that can say "the situation has shifted, we need to revisit what we agreed to" is a team that can handle complexity.

Shared Purpose and Strategic Clarity

Without a compelling shared understanding of what the team is trying to accomplish and why it matters, cognitive diversity becomes noise. Different perspectives are only productive when they're pointed at a shared problem.

Shared purpose has to be connected to organizational strategy, where the organization is headed and how it plans to get there. If the team doesn't understand the strategy, they can't align their work to it. And sometimes the problem is upstream: the strategy itself is absent or incoherent, and the team is left trying to create meaning in a vacuum. That's a leadership accountability (Hackman, 2002).

The Manager as Doula

The manager's role in a team built for strong team cognition is not to direct the thinking or surveil the process. It's to reduce unproductive friction and create the conditions for the team's best work.

I borrow the doula metaphor intentionally. A doula doesn't deliver the baby. The doula supports, coaches, clears obstacles, and creates an environment where the natural process can unfold. The manager's job is the same: make sure resources are available, keep the team connected to vision, protect the liminal space, and get out of the way.

Not all friction is bad. Productive friction, the cognitive challenge of working with people who think differently, the discomfort of having your mental model challenged, the effort of integrating perspectives that don't fit neatly, is essential to team cognition. It's the inherent discomfort of the liminal space itself, and it's where learning and innovation happen. That's what Bjork (1994) called desirable difficulty.

Unproductive friction is something else. It's extraneous cognitive load (Sweller, 1988): trying to figure out how to log onto a system, hunting for a report template, navigating unclear processes. This wastes the team's cognitive capacity on things that have nothing to do with the work. The manager's job is to eliminate as much of it as possible.

Reflective Practice

Teams develop their cognition through cycles of action and reflection. The reflection is where learning happens: the team updates its mental models, refines its transactive memory, and makes sense of what just happened (Schön, 1983; Kolb, 1984).

Most organizations get this wrong. The only reflection many teams ever do is a postmortem at the end of a project, and those sessions tend to be socially constrained. People are reluctant to name specific problems because they'll have to work with these colleagues again. The "lessons learned" stay vague and diplomatic (Edmondson, 2012).

The more powerful form of reflective practice is continuous: micro-corrections as the work unfolds rather than a debrief when it's too late to change anything. But continuous reflection only works when the team has shared, explicit standards and enough trust that someone can say "this isn't meeting our agreed standard" without it feeling like an attack. When both conditions are present, reflective practice becomes part of how the team works, not a scheduled event everyone dreads.

The Organizational Operating System

There's a practical question that anyone leading project-based work will immediately raise: if team cognition requires shared experience to develop, what happens when teams reform for every new project?

This is a real concern. In organizations where teams are assembled and reassembled based on project needs, utilization requirements, and available talent, no team stays together long enough to build full transactive memory from scratch. But the cross-pollination that comes from fluid

composition is itself valuable. People carry mental models, strategies, and ways of working from one team to the next. That movement injects cognitive diversity.

The answer is that certain conditions for team cognition can live at the organizational level rather than the team level. If trust architecture, agreement-making practices, communication norms, explicit standards, and reflective practice habits are part of the organization's shared operating system, then reformed teams aren't starting from zero. They're assembling new configurations of people who already share a common foundation. The specific transactive memory of who on *this* team knows what still has to develop through working together. But the infrastructure for how we work together travels with each person.

This has a significant secondary benefit. When teams have agency over their own ways of working within this shared framework, change becomes emergent rather than imposed. The organizational cost of adaptation drops, because the people making decisions about how to work are the people doing the work. They have the context, they own the agreements, and they don't need to be "change managed" into something they created themselves.

When These Conditions Are Absent

When these conditions aren't present, team cognition doesn't emerge. What you get instead is predictable: groupthink, where cognitive diversity is suppressed because disagreement feels unsafe. Star culture, where one or two voices crowd out everyone else. Revolving doors without shared infrastructure, where no team ever develops real collective intelligence. Task focus without purpose, where people know what to do but not why.

The cost is not simply underperformance. It's wasted human potential on a scale we can no longer afford. Every team that fails to achieve real team cognition is a group of people whose best thinking never made it into the room. The wicked problems, the ones that will shape whether our organizations, our institutions, and our communities thrive or decline, are waiting for exactly the kind of collective intelligence that these failed teams were supposed to produce. The talent is there. The capacity for creativity and moral imagination is there. It's not making it into the room.

That is the urgency. Not that AI is coming. AI is already here. The urgency is that we have an extraordinary opportunity to focus our collective human intelligence on the problems that actually matter, and we are squandering it on teams that can't think together.

The AI Extension

Everything I've described so far applies to purely human teams. If you stopped reading here, you'd have a useful framework for building teams that think together. The research behind these concepts has been accumulating for decades.

What's new is what happens when AI enters the picture.

Many people have strong emotional reactions to AI. Fear, distrust, hostility. I've sat in rooms full of smart, experienced knowledge workers and listened to them describe AI with the kind of language you'd reserve for an invasive species. That fear is real. People are worried about losing their jobs, their expertise, their sense of professional identity. I don't dismiss any of that.

But the fear, and the mechanical, transactional way most people use AI as a result, is preventing us from seeing something important. Most people treat AI as an efficiency tool, a faster way to draft or summarize. That's not wrong, but it's radically incomplete. It's a mental model problem, and it changes the way all mental models change: through experience plus reflection, not through argument. Nobody has ever been talked out of their fear of AI. The people most afraid are the least likely to have the experience that would shift their understanding. The mental model protects itself.

The Human-AI Dyad

What I want to propose is a different orientation: AI as cognitive partner. In this framing, the human and the AI form a dyad, a thinking partnership in which each contributes capabilities the other lacks. The human brings judgment, values, lived experience, creativity, ethical reasoning, and embodied cognition, the thinking that is grounded in having a body in the world, in sensory experience, in the felt sense of a situation that no algorithm can replicate (Varela, Thompson, & Rosch, 1991). The AI brings processing capacity, pattern recognition, breadth of information, and the ability to hold and organize vast amounts of complexity.

This isn't a utopian fantasy. It's a description of what already happens when someone develops a genuine thinking relationship with an AI. The experience is qualitatively different from transactional use. It changes how you think, not just how fast you produce.

In an AI-augmented team, the unit of participation is not the individual. It's the dyad. Each person brings their AI partner's capabilities into the team the way they'd bring any other form of expertise. Shared mental models now include the understanding that develops within each dyad. Transactive memory expands because each dyad accesses vastly more knowledge than any individual, though the team still needs to know which dyad has worked through which problems. The directory function remains human.

Cognitive diversity becomes richer. Different people develop different relationships with AI. Some use it for analysis, some for creative exploration, some for research, some for reflection. This variation produces genuinely different cognitive configurations at the dyad level, which enriches team cognition at the team level. The variation isn't a problem to be solved through standardized training. It's a feature. The floor for meaningful participation is basic skills and an open mind. Curiosity and a willingness to engage.

The talent pipeline reopens here too. A person at the apprentice level, paired with an AI, can participate meaningfully in a team's thinking much earlier than traditional development paths

would allow. They bring a different dyad configuration to the team, and that difference is valuable. The team becomes the training ground, not through isolated entry-level tasks that may no longer exist, but through participation in collective thinking augmented by a cognitive partner.

Trust Architecture and AI

Trust architecture becomes even more critical when AI enters the team. The dyad needs to be a safe container for thinking. When I work with my own AI partner, the conversation includes half-formed ideas, wrong guesses, exploratory tangents, and the kind of raw sensemaking that could be embarrassing if taken out of context. That's not a failure of discipline. That's the liminal space within the dyad. That's where the cognitive work happens. If I thought someone else was reading those conversations, I would immediately change how I use AI, and everything I brought to a team would diminish.

Multiply that across an organization and the stakes become clear. When people self-censor with their AI, you've degraded every dyad. You've made the entire team less intelligent to satisfy one person's need for control. Surveillance of AI use isn't just ethically questionable. It's cognitively expensive.

The manager-as-doula model applies here directly. The manager's job is to support people in developing productive dyadic relationships, not to monitor them. Clear agreements about what knowledge is shared at the team level and what stays within the dyad are essential. Privacy should be the default. The human decides what surfaces.

This leads to what I think is the most important design principle for AI-augmented teams: ****the human controls the permeability of the dyad boundary.**** The AI doesn't decide what becomes team knowledge. The person does.

As AI capabilities continue to develop, the systems themselves should be designed to support trust architecture natively, giving humans clear and reliable control over what stays within the dyad and what gets shared with the team.

Transparency as Intellectual Honesty

There's one area where transparency isn't just appropriate but necessary. In many organizations today, people use AI quietly. They think with it, draft with it, research with it, and then present the output as entirely their own. This is not a fringe behavior. A 2025 global study by KPMG and the University of Melbourne found that 57% of employees have concealed their AI use from managers and colleagues, and 55% have presented AI-generated content as their own without disclosure (KPMG, 2025). A separate WalkMe survey found that nearly half of employees hide AI use to avoid judgment, with C-suite leaders even more likely to conceal it than their teams (WalkMe, 2025). The reasons range from fear of looking replaceable to shame to the simple absence of any organizational norm about it.

For teams working on wicked problems, this secrecy undermines the team's ability to evaluate its own thinking. If nobody knows AI was involved in shaping an idea, nobody can ask whether biases were introduced, whether assumptions were challenged, or what data the AI was drawing on. The analogy is scientific method: scientists publish their methods alongside their results so that the community can evaluate the rigor of the work. A team working on complex, high-stakes problems needs the same ethic. Not surveillance, but transparency. Not "show me your AI conversations" but "be explicit about how your thinking developed, including the role your AI partner played."

This is explicit communication as a trust architecture norm. The team agrees that AI use is visible not because anyone is being monitored, but because the quality of the team's collective intelligence depends on being able to evaluate how ideas were produced.

A Call to the Room

We are living through a moment when the nature of knowledge work is changing fundamentally. The routine cognitive labor that has defined most knowledge jobs for decades is migrating to AI. What remains is the work that requires human judgment, human creativity, human conscience, human imagination. The wicked problems. The ones that matter.

That work is too important to waste on teams that can't think together.

The constructs described in this paper, shared mental models, transactive memory, team situational awareness, are not new ideas. The research has been building for decades. Edmondson, Wegner, Salas, Woolley, Cranton, and many others have laid the foundation. What's new is the urgency, and what's new is the possibility that AI, understood as a cognitive partner rather than a tool or a threat, can amplify what teams are capable of in ways we're only beginning to understand.

The work ahead is deeply, irreducibly human. Building trust. Cultivating diversity of thought and experience. Making clear agreements and honoring them. Designing the organizational infrastructure that lets team cognition develop even as teams reform and reconfigure. Creating the conditions for people to think together in the liminal space, that uncomfortable, generative threshold where something new becomes possible.

Every knowledge work discipline faces the same question right now: what are the wicked problems our expertise could help solve, if we organized our collective intelligence to tackle them? The specifics differ across medicine, engineering, education, finance, law, and design. The underlying challenge is the same. The problems are too complex for individuals. They require teams that can actually think together.

So here is my invitation. Start with your own team. Look at the conditions described in this paper and ask which are present and which are missing. Start with your own relationship with AI, not

as a task machine but as a thinking partner. Start with curiosity, which is always the lowest barrier to entry and the most powerful catalyst for change.

The liminal space is waiting.

Author's note: This paper developed through a series of conversations between the author and Claude, an AI thinking partner built by Anthropic, in the spring of 2026. The ideas, arguments, and frameworks are the author's own; Claude served as a collaborative partner in developing and articulating them. This is, in itself, an example of the human-AI dyad at work.

References

Bjork, R. A. (1994). Memory and metamemory considerations in the training of human beings. In J. Metcalfe & A. Shimamura (Eds.), *Metacognition: Knowing about knowing* (pp. 185-205). MIT Press.

Cannon-Bowers, J. A., Salas, E., & Converse, S. (1993). Shared mental models in expert team decision making. In N. J. Castellan, Jr. (Ed.), *Individual and group decision making* (pp. 221-246). Lawrence Erlbaum Associates.

Cranton, P. (2006). *Understanding and promoting transformative learning: A guide for educators of adults* (2nd ed.). Jossey-Bass.

Duhigg, C. (2016, February 25). What Google learned from its quest to build the perfect team. *The New York Times Magazine*.
<https://www.nytimes.com/2016/02/28/magazine/what-google-learned-from-its-quest-to-build-the-perfect-team.html>

Edmondson, A. C. (1999). Psychological safety and learning behavior in work teams. *Administrative Science Quarterly*, 44*(2), 350-383. <https://doi.org/10.2307/2666999>

Edmondson, A. C. (2012). *Teaming: How organizations learn, innovate, and compete in the knowledge economy*. Jossey-Bass.

Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37*(1), 32-64. <https://doi.org/10.1518/001872095779049543>

Hackman, J. R. (2002). *Leading teams: Setting the stage for great performances*. Harvard Business School Press.

Hollingshead, A. B. (1998). Communication, learning, and retrieval in transactive memory systems. *Journal of Experimental Social Psychology*, 34*(5), 423-442.
<https://doi.org/10.1006/jesp.1998.1360>

Kolb, D. A. (1984). *Experiential learning: Experience as the source of learning and development*. Prentice-Hall.

KPMG & University of Melbourne. (2025). *Trust, attitudes, and use of artificial intelligence: A global study 2025*.
<https://theconversation.com/major-survey-finds-most-people-use-ai-regularly-at-work-but-almost-half-admit-to-doing-so-inappropriately-255405>

Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge University Press.

Lewis, K. (2003). Measuring transactive memory systems in the field: Scale development and validation. *Journal of Applied Psychology*, 88*(4), 587-604.
<https://doi.org/10.1037/0021-9010.88.4.587>

Mathieu, J. E., Heffner, T. S., Goodwin, G. F., Salas, E., & Cannon-Bowers, J. A. (2000). The influence of shared mental models on team process and performance. *Journal of Applied Psychology*, 85*(2), 273-283. <https://doi.org/10.1037/0021-9010.85.2.273>

Mezirow, J. (1991). *Transformative dimensions of adult learning*. Jossey-Bass.

Page, S. E. (2007). *The difference: How the power of diversity creates better groups, firms, schools, and societies*. Princeton University Press.

Rittel, H. W. J., & Webber, M. M. (1973). Dilemmas in a general theory of planning. *Policy Sciences*, 4*(2), 155-169. <https://doi.org/10.1007/BF01405730>

Salas, E., Prince, C., Baker, D. P., & Shrestha, L. (1995). Situation awareness in team performance: Implications for measurement and training. *Human Factors*, 37*(1), 123-136.
<https://doi.org/10.1518/001872095779049525>

Schön, D. A. (1983). *The reflective practitioner: How professionals think in action*. Basic Books.

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science*, 12*(2), 257-285. https://doi.org/10.1207/s15516709cog1202_4

Varela, F. J., Thompson, E., & Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. MIT Press.

Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes*. Harvard University Press.

WalkMe. (2025). *AI in the workplace survey 2025*.
<https://fortune.com/2025/08/29/what-is-ai-shame-readiness-gap-training-artificial-intelligence/>

Wegner, D. M. (1987). Transactive memory: A contemporary analysis of the group mind. In B. Mullen & G. R. Goethals (Eds.), *Theories of group behavior* (pp. 185-208). Springer-Verlag.

Woolley, A. W., Chabris, C. F., Pentland, A., Hashmi, N., & Malone, T. W. (2010). Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330(6004), 686-688. <https://doi.org/10.1126/science.1193147>