



How to **scale coaching** **with AI**, without **diluting its impact.**

Karim Hirani and Sheila Boysen-Rotelli, Ph.D, MCC



**That ceiling
just broke.**

For most of modern business history, if you wanted access to a great coach, you needed to be important enough to justify the cost. That meant senior executives, high-potentials, and the occasional lucky manager whose company had budget to spare. Everyone else developed the old-fashioned way: through experience, through mistakes, and through whatever feedback their manager had time to give.

That ceiling just broke. AI-powered coaching tools can now reach every leader in an organization, from the newest team lead to the most seasoned director, at a fraction of the cost of traditional coaching and with a speed and availability that no human coach can match. For organizations that have spent years trying to build leadership at scale, this feels like a genuine breakthrough, and a well-timed one at that. According to [Gartner research](#), 85% of learning and development leaders expect a surge in skills development needs driven by AI and digital transformation over the next three years, while simultaneously emphasizing constraints in time, resources, and leadership capacity.

But breakthroughs have a way of moving faster than our ability to think clearly about them. And right now, the field of AI coaching is moving very fast, with very little shared vocabulary, very little design rigor, and a great deal of well-intentioned noise.

We built our own AI coaching companion at BTS. The process taught us more than we expected, not just about technology, but about coaching itself. About what makes it work, what makes it fail, and what gets lost when we stop being precise about what we're doing and why. This paper is an attempt to share what we learned, honestly and practically, so that the organizations and leaders investing in AI coaching can do it in a way that actually develops people.

The original algorithm: What coaching honors

Coaching exists to help people develop the capacity to think, choose, and act for themselves. Not through advice, instruction, or even feedback, but through a relationship deliberately designed to expand what the client can see, question, and decide on their own. That means resisting the pull to advise, to reassure, to solve. It's a discipline that runs counter to most of our instincts, and it's precisely what makes coaching so powerful and so difficult to replicate.

[A 2025 study](#) from Stanford's Institute for Human-Centered Artificial Intelligence, based on a survey of 1,500 workers across 104 occupations, found that 69.4% of workers want AI to function as a collaborator rather than purely as an automation tool, and 45.2% specifically prefer an equal human-AI partnership model. People aren't just asking AI to do things for them, they're asking it to think alongside them. This is reinforced by [Anthropic's 2026 study](#) of more than 80,000 AI users, which specifically set out to understand what people want from AI in their lives and work. Notably, in professional contexts, two themes stood out: the desire for "professional excellence" (18.8%), with users seeking to offload routine tasks to focus on higher-value work, and "cognitive partnership" (17.2%), where AI supports thinking, problem-solving, and reflection.

The more valuable something is, the more consequential it becomes when that value gets diluted. And right now, coaching's meaning is being stretched in ways that quietly put that value at risk, often without anyone realizing it until the outcomes disappoint.

What we learned building our own AI coaching companion

When BTS set out to build an AI coaching companion, we thought the hard part would be the technology. It wasn't.

The hard part was the thinking. Specifically, the thinking required to make explicit what experienced coaches have always held implicitly: when to ask and when to offer, when to challenge and when to validate, when to introduce a model and when to trust the leader to find their own frame. Every design decision forced us to articulate something that had previously lived in the judgment of a skilled practitioner.

We also discovered how much the underlying model had absorbed the market's confusion. Trained on vast amounts of data in which the word coaching appears in almost every context imaginable, it had no reliable way to distinguish between a coaching conversation, teaching, practicing and advice-giving. Left to its own defaults, it would drift toward the common solutions and patterns. And the most common patterns, as we've established, aren't always the right ones.

Furthermore, the market often equated human and AI coaching, without really thinking through what each was. Human coaching is relational, presence-based, and emergent. AI does not experience meaning; it shapes the background against which humans make meaning. AI coaching is probabilistic, pattern-based, and scalable. It cannot offer emotional resonance (though can pick up on feelings), intuition (though can read between the lines) or be open to the multiplicity of what is present in each moment (though it can ask questions to open perspectives).

What we learned is that building good AI coaching requires three things:

- ✓ Clarity about what each kind of support is and does
- ✓ Grounding in the science of how people actually change
- ✓ Genuine commitment to ethical design

Here's what each of those looks like in practice.

Design principle 1:

Be clear about what you're building and the outcomes

It's Monday morning. One of your managers has a conversation they've been dreading all weekend. A direct report who keeps missing commitments. They know it needs to happen. They're not sure they're ready. They open your AI coaching tool.

What happens next depends entirely on what that tool was actually built to do.

In that moment, they might need teaching: a clear framework for the conversation, what to say, how to structure it, what to avoid. They might need practice: a chance to rehearse it out loud, hear themselves say the words, and build enough confidence to walk in ready. Or they might need coaching: someone to help them understand why this feels so hard, what assumptions they're making about conflict, what they believe about their own authority, and what kind of leader they actually want to be.

The problem isn't that any of these is wrong. It's that calling all of them coaching makes it nearly impossible to know what you're actually delivering, whether it's working and confuses participants about what coaching is. It can undermine the empowering nature of coaching, when training is offered under the umbrella of coaching. When we are clear, we can be conscious of the chosen AI, be clear on the measures and outcomes, and protect the integrity of coaching professionally. Teaching builds knowledge. Practice builds confidence. Coaching builds the self-awareness and agency that allows a leader to navigate not just this conversation, but every difficult one that follows.

That's why we distinguish between six distinct functions that AI can serve in a development context:



AI coach

supports reflection, sense-making, empowerment and goal clarity.



AI trainer

builds knowledge and capability in context.



AI practice

partner enables rehearsal through role-play and simulation.



AI assessor

provides diagnosis, competency observation, and insight.



AI guidance

offers advice, direction, or solutions.



AI decision

support explicitly supports and directs decision-making.

None of these are interchangeable. A well-designed program uses each deliberately, names them honestly, measures them against the right outcomes, and distinguishes the empowering nature of coaching from other tools. There are also AI tools that support coaching that are not based on chatbots, including scheduling, coach matching, and session summarization. As AI evolves, these will increasingly enrich the human coaching relationship, providing continuity and context between sessions.

Design principle 2:

Ground your AI in how people actually change

Consider a leader who has been through five years of leadership development programs. She can articulate what good feedback looks like. She can name every model in her organization's competency framework. But when her CFO blindsides her in a board meeting, she goes quiet, forgets the knowledge she gained and spends two weeks second-guessing her own judgment. We refer to the gaining of skills as horizontal development. It is often not enough.

What she has not yet developed is the worldview and assumptions that hold her back from speaking up – the internal stories and narratives that hold her back from speaking up with a person in authority. Horizontal skill development cannot meet this capacity to hold her ground when her worldview is challenged in real time. It does not explore the mindset and worldview that are inhibiting change.

This is particularly relevant when it comes to vertical development, one of the most important and least understood concepts in leadership development. While horizontal development focuses on acquiring skills and knowledge, vertical development refers to the evolution of how a person makes meaning. It involves growing capacity to hold multiple perspectives, navigate complexity and ambiguity, reflect on one's own assumptions, and respond rather than react. It's the difference between a leader who has learned what to do and a leader who has fundamentally changed how they see, thereby sustainably changing what they will do.

AI can support vertical development by prompting reflection, surfacing patterns over time, and helping leaders integrate their learning between human coaching sessions. It extends the reach of that deeper development without replacing the relational depth that makes it possible.

For this leader, the coaching sessions that follow would not focus on giving her new tools and techniques. They would surface her patterns: that she tends to interpret being challenged as evidence that she was wrong, not as evidence that she is being tested. That shift, from “I must have been wrong” to “I am being tested and I can handle that,” is not a skill. It is a change in how she makes meaning, in worldview or mindset. AI cannot produce that shift alone, but it can hold the thread between conversations, reflect back the pattern, and help a leader begin to see what she could not see before.

AI interventions in coaching must be built on a genuine theory of change, both behavioral and mindset-based, if they’re going to produce real impact. Our AI coaching companion is built on BTS’s research into how mindsets and worldviews shift, developed through decades of work and over 100,000 hours of coaching conversation in leadership development, strategy execution and organizational transformation. Sustainable change requires addressing not just skills and behaviors, but the underlying assumptions and meaning-making structures that shape how leaders act under pressure.

Without AI grounded in credible human mindset change research, AI tools are subject to the vastly varied landscape of approaches, mostly unproven, that respond with horizontal solutions when mindset and vertical development is actually what is needed for sustainable behavior change.



Design principle 3:

Take ethics seriously, not just seriously enough

Coaching ethics exist to protect the person being coached, ensuring their growth is not shaped by someone else's agenda, answers, or assumptions. As AI begins to take on a coaching role, this principle becomes especially important. If a tool presents itself as coaching but begins to direct, advise, or instruct, it subtly shifts the experience. The intention may still be positive, but the dynamic changes from supporting someone's thinking to influencing it. Because the label remains the same, that shift can be difficult to see in the moment, even as it impacts how leaders grow.



Ethics break down when there is no visibility

Imagine a global financial services firm piloting an AI coaching solution for its director-level population. When the CHRO asks what the tool is actually doing when it responds to someone, the L&D team cannot answer. The vendor's documentation describes the product as "ICF-aligned" and "grounded in positive psychology," but there is no way to verify what that means in practice. There are no disclosed frameworks, no explained logic, and no way to audit why the tool said what it said to a leader in distress.

Six months in, the team notices a pattern in the engagement data. Leaders are returning frequently, sessions are long, and satisfaction scores are high, but 360 feedback results have not moved. The tool is clearly providing value in some way. It just may not be the kind of development the organization intended. This gap between perceived support and actual growth is what can emerge when ethics is treated as a compliance checkbox rather than a design constraint.

A well-supported leader who isn't continuing to grow ultimately limits the return on an organization's investment in their development. To unlock that growth, it is essential to understand what is holding them back, which is not possible if the process remains a black box.



Quality is an ethical issue

Quality is not separate from ethics. When an AI tool consistently validates a leader's perspective or offers encouragement where challenge might be more useful, it may feel supportive but can subtly reinforce existing thinking rather than expand it.

In human coaching, growth comes from the balance between support and challenge. Skilled coaches calibrate this carefully, creating enough safety while introducing the right level of stretch. It is this balance, and the judgment behind it, that enables meaningful development. A consistently "validating" AI persona can struggle to hold that tension.

At scale, the risk is a growing gap between feeling supported and actually developing, one that can be difficult to detect until confidence in the investment begins to erode.



Transparency is what makes ethics actionable

A lack of transparency compounds both problems. Most commercial AI coaching solutions do not disclose the frameworks, research base, or ethical guidelines shaping their outputs, which makes it nearly impossible for L&D professionals to conduct meaningful due diligence.

If you cannot determine whether an AI's responses reflect established coaching methodology or simply patterns absorbed from internet-scale training data, you cannot govern it, align it to your values, or hold it to any professional standard.



What this requires in practice

Our approach addresses all of these challenges. It includes alignment with professional coaching ethics and competencies, transparency about the role the AI is playing at every moment, curated research-backed content using retrieval-augmented generation, and robust safeguarding protocols for situations where interactions surface psychological distress or complexity that requires human judgment.

We also partner closely with the organizations we work with to learn and adapt continuously, because we do not treat AI coaching as a finished product. We treat it as an evolving discipline.

Will AI coaches replace human coaches?

Here's a more useful frame than the replacement debate: AI coaching is genuinely effective for structured, model-based development.

[Research by Terblanche et al. \(2022\)](#) shows it can match human coaches on goal attainment in those contexts, and its ability to deliver that kind of support at scale is a real advantage. But the same research identifies deeper levels of coaching that AI cannot replicate, the kind that requires a coach who can read a room, sense what's unspoken, and be genuinely present to the complexity of another human being. That's not a limitation AI will engineer its way out of anytime soon.

The answer isn't to choose between AI and human coaching, but to design them intentionally based on what each does best. Our experience deploying both has shown that each brings distinct strengths, and the real opportunity lies in how they are applied in different contexts.

At scale, AI can provide consistent, always-on support, handling repetition, availability, and continuity in day-to-day development. It can offer a range of roles transparently based on the context. Human coaches, by contrast, bring relational depth, navigate complexity, and offer the presence required for deeper transformation. The value comes not from assuming they are always combined, though they can be a powerful conversation in the right context, but from understanding and using them when and how each is most impactful.

Questions we're still sitting with

We don't have all the answers, and nobody does yet. The field is moving fast, and the most useful thing we can offer alongside our experience is a shared sense of curiosity about what comes next. These are the questions we keep coming back to, and we think they're worth your attention too:



How do we design AI and human coaching to genuinely complement each other, not just coexist in the same program?



Which AI interventions are actually building leaders, and which just feel like they are?



How do we build feedback loops that improve AI coaching based on real developmental outcomes, not just satisfaction scores?



What does AI's role in vertical development look like as the technology matures?



What role should professional coaching bodies play in setting standards, and how do we make sure L&D professionals have a seat at that table?

Think back to that manager on Monday morning. Done well, their AI experience wouldn't feel like a series of interventions. It would feel like genuine support, moving naturally from reflection > learning > practice > action. The distinction between those modes matters to the people designing the system. For the leader using it, it just feels like growth.

AI didn't just democratize access to coaching. It forced the field to get serious about what coaching actually is and how it can create development at scales we've never seen. That's not a threat to the profession. It's an invitation to be more precise, more honest, and more intentional about one of the most powerful tools for human development we have. The organizations that accept that invitation will build stronger leaders, healthier cultures, and more resilient organizations.