

FOUNDATIONS • AIL-FP-2025-02

AI Solves the Most Important Challenge in Education

Educators have known the secret to exceptional learning since 1984. AI will finally unlock it for everyone.

Gregg Collins • Brandon Dickens • June 2025

Benjamin Bloom's 1984 "2 Sigma Problem" showed that one-on-one tutoring lifts the average student to the 98th percentile, yet the cost of human tutors has kept that result out of reach for four decades. This paper argues that generative AI has finally removed every constraint that made personalized tutoring unscalable, and that the organizations which treat AI as an enhancement to existing course-and-quiz pedagogy will fall behind those that reimagine learning from the ground up. Drawing on a four-stage maturity model—Accelerate, Elevate, Transcend, Unify—the paper contrasts incremental uses of AI with the transformational possibility of embedding expert-level coaching directly into the flow of work. The payoff is a workforce whose typical performance matches what we currently call exceptional.

The best way to predict the future is to invent it.

— Alan Kay

If I had to write a headline for the entire history of educational technology, it would be old wine in new bottles.

— John Seely Brown

Success is a lousy teacher. It seduces smart people into thinking they can't lose.

— Bill Gates

I want to tell you about a problem that has haunted education for forty years, and why artificial intelligence is about to solve it in a way that will make most of what we

call “training” today look as quaint as a one-room schoolhouse.

In 1984, educational psychologist Benjamin Bloom published research that should have revolutionized how we think about learning. He called it the “2 Sigma Problem,” and here’s what he found: students who received one-on-one tutoring performed two standard deviations better than those in traditional classroom settings. In plain English, this means the average tutored student outperformed 98% of students in conventional classrooms—a jump from the 50th percentile to the 98th percentile.

Think about that for a moment. Under optimal learning conditions, nearly every student could reach what we consider exceptional performance levels. The implications are staggering.

For sales professionals, a 2-sigma improvement could mean tripling conversion rates, cutting sales cycles in half, and doubling average deal sizes. For any professional role, it represents the difference between mediocrity and mastery.

But here’s the kicker: despite knowing this for four decades, we’ve done virtually nothing about it.

Why? Because one-on-one tutoring has been economically impossible at scale.

Until now.

§1. The enhancement trap

Most organizations approaching AI in learning are making a fundamental mistake. They’re thinking like the drunk searching for his keys under the streetlamp—not because that’s where he lost them, but because the light is better there.

According to industry research, approximately 70% of organizations are incorporating AI into their L&D efforts, but the vast majority are stuck in what I call

“enhancement thinking.”

They’re using AI to create traditional learning materials faster, automate administrative processes, or build basic recommendation systems. In other words, they’re using 21st-century technology to make 19th-century pedagogy slightly more efficient.

This is like using a Ferrari to deliver milk faster rather than reimagining transportation entirely.

I’ve been watching this pattern for years in instructional design. When e-learning emerged, most organizations simply converted their PowerPoint presentations into “interactive” courses with multiple-choice quizzes. When mobile learning became possible, they took those same courses and made them smaller. Now, with AI, they’re taking those same approaches and making them faster to produce.

They’re missing the point entirely.

§2. What one-on-one tutoring actually does

To understand why AI represents a revolutionary rather than evolutionary change, we need to examine why one-on-one tutoring works so well. It’s not magic—it’s a systematic approach that addresses fundamental problems with how humans learn.

Immediate feedback correction. When you make a mistake in a one-on-one setting, it gets corrected immediately, before it becomes a habit. In traditional training, mistakes often go unnoticed until it’s too late.

Adaptive pacing. A human tutor adjusts the difficulty and pace based on your responses. If you’re struggling with a concept, they slow down and provide more examples. If you’re grasping things quickly, they accelerate.

Personalized practice. The tutor creates practice scenarios specifically tailored to your weaknesses and learning style. No time is wasted on things you already know.

Emotional engagement. A good tutor reads your emotional state and adjusts accordingly. Frustrated? They provide encouragement. Bored? They increase the challenge.

Retrieval practice. Instead of just presenting information, a tutor constantly tests your recall and application, strengthening memory formation.

These elements work together to create what cognitive scientists call “desirable difficulties”—challenges that make learning harder in the moment but dramatically improve long-term retention and application.

The problem has always been scale.

How do you provide this level of personalized attention to hundreds or thousands of employees? You can't hire enough tutors, and even if you could, the cost would be prohibitive.

§3. Enter the AI tutor

Here's where most people's imagination fails them. They think of AI as a better search engine or a faster content creator. But generative AI is fundamentally different from any technology we've had before because it can engage in genuine dialogue, adapt its responses in real-time, and create unique scenarios on demand.

For the first time in history, we can provide personalized tutoring at scale.

Let me give you a concrete example of what this looks like in practice. Instead of taking a generic sales training course, imagine an AI system that:

- Analyzes your actual sales conversations to identify specific improvement areas
- Creates personalized role-play scenarios based on your real prospects and challenges
- Adapts the difficulty of objections based on your performance
- Provides immediate, specific feedback on your responses
- Generates unlimited variations so you never practice the same scenario twice
- Tracks your progress and adjusts the curriculum accordingly

This isn't enhancement—this is transformation. It's the difference between practicing scales on a piano and having a master musician compose personalized exercises that target your specific technical weaknesses while building toward pieces you actually want to play.

§4. The historical constraints that no longer exist

To appreciate how revolutionary this is, consider the constraints that have shaped every learning intervention you've ever experienced.

The scalability constraint. High-quality training required scarce resources—expert instructors, physical spaces, synchronized schedules. This forced us into the “batch processing” model of learning where everyone gets the same content at the same time.

The personalization barrier. Even when we knew that different people needed different approaches, creating truly personalized learning was prohibitively expensive. We settled for “learning paths” that were really just different sequences of the same generic content.

The feedback delay. In traditional training, you complete a module, take a quiz, and maybe get results later. By then, any misconceptions have already taken root. Real

learning requires immediate correction.

The practice problem. Creating realistic practice scenarios was expensive and time-consuming. We ended up with generic case studies that bore little resemblance to learners' actual challenges.

The expertise bottleneck. Access to expert knowledge was limited by the availability of actual experts. When the expert wasn't available, learning stopped.

AI removes every single one of these constraints.

§5. The maturity model: from enhancement to transformation

Based on my observations of organizations implementing AI in learning, I see four distinct stages of maturity.

Stage 1: Accelerate. Using AI to speed up traditional processes. Creating course outlines faster, generating quiz questions, automating administrative tasks. This is where most organizations are stuck.

Stage 2: Elevate. Human-AI collaboration within existing frameworks. AI helps instructors create better content or provides tutors with real-time suggestions. Still working within traditional models.

Stage 3: Transcend. Creating entirely new learning modalities. Dynamic simulations that adapt in real-time, AI coaches that provide personalized feedback, interactive scenarios that change based on learner decisions.

Stage 4: Unify. AI orchestrating complete learning journeys. The system understands each learner's goals, current capabilities, and optimal learning approach, then creates a continuously adapting experience that spans all learning modalities.

Most organizations never make it past Stage 1 because they can't imagine what Stages 3 and 4 look like. They're trying to use AI to make their current approaches better rather than asking, "If we could design learning from scratch with unlimited personalization and expert availability, what would it look like?"

§6. What this means for learning designers

If you're still thinking about "courses" and "modules" and "completion rates," you're about to become as relevant as a telegraph operator. The future of learning looks nothing like what we do today.

Instead of designing linear experiences that all learners follow, we'll be designing learning environments—spaces where AI can create infinite variations of scenarios, challenges, and interactions based on each learner's needs.

Instead of measuring seat time and quiz scores, we'll be measuring behavioral change and performance improvement. Can the sales rep actually handle objections better? Can the manager actually have difficult conversations more effectively? Can the engineer actually solve problems faster?

Instead of front-loading all the content and hoping it transfers to the job, learning will be embedded in the workflow itself. AI will provide just-in-time coaching during actual work situations, turning every challenge into a learning opportunity.

§7. The imagination gap

The biggest barrier to realizing this potential isn't technological—it's conceptual. Most learning professionals suffer from what I call an "imagination gap." They can envision AI making their current processes faster or cheaper, but they can't imagine AI enabling entirely new approaches that were previously impossible.

This gap exists because we've been constrained by limitations for so long that we've forgotten they were constraints rather than inherent features of learning. We assume that learning has to happen in courses because that's how we've always done it. We assume that personalization is too expensive because it always has been. We assume that expert guidance is scarce because it always has been.

But what if none of those assumptions are true anymore?

§8. The path forward

If you want to move beyond enhancement thinking, start by asking different questions:

Instead of “How can AI make our training faster?” ask “What would learning look like if every employee had a personal expert coach available 24/7?”

Instead of “How can we create content more efficiently?” ask “What if content was created dynamically in response to each learner's performance?”

Instead of “How can we track completion better?” ask “How can we measure whether people are actually getting better at their jobs?”

The organizations that bridge this imagination gap first will develop capabilities that their competitors won't understand, let alone match. They'll solve Bloom's 2 Sigma Problem and create workforces that perform at levels we currently consider exceptional.

The question isn't whether AI will transform learning—it's whether your organization will lead or follow in that transformation.

And if you're still designing courses with modules and quizzes, you're already following.

The future of learning isn't about making training better. It's about making training obsolete by embedding personalized, expert-level guidance directly into the work itself. The technology to do this exists today. The only question is whether we have the imagination to use it.