

FOUNDATIONS • AIL-FP-2026-06

The Infiltration Strategy

How AI Will Take Over Education and Replace the Curriculum

Gregg Collins • Brandon Dickens • May 2026

L&D organizations are adopting AI cautiously, treating it as a productivity tool to do what they already do—create courses—a little faster and a little cheaper. They have profoundly misjudged what they have welcomed into the building. Technology revolutions, history shows, do not happen top-down through enlightened leadership; they happen bottom-up, through infiltration. New tools enter institutions in service of existing processes, then quietly outgrow them, until the original way of doing things is rendered obsolete. This paper traces the predictable ten-stage progression by which AI-driven simulations and coaches are infiltrating corporate learning—starting as supplemental practice partners, evolving into persistent personal coaches, and finally generating dynamic, on-demand learning experiences that bypass courses altogether. The implication is uncomfortable: the curriculum will not be abolished by reformers. It will be abandoned, gradually and quietly, by learners who have experienced something better.

The future is already here—it's just not evenly distributed.

— William Gibson

It's easy to fall into the trap of assuming that a new technology is very similar to its predecessors. A new technology is often perceived as the linear extension of the previous one, and this leads us to believe the new technology will fill the same roles—just a little faster or a little smaller or a little lighter.

— Michael J. Saylor

In a culture like ours, long accustomed to splitting and dividing all things as a means of control, it is sometimes a bit of a shock to be reminded that, in operational and practical fact, the medium is the message.

— Marshall McLuhan

§1. Tech Revolutions Don't Happen Top-Down

In 1968—in what has since become known as “the mother of all demos”—Doug Engelbart’s team at the Stanford Research Institute demonstrated a computer outfitted with a mouse, a graphical user interface, and a window you could click through to see multiple documents at once. In a world of blackboards, typewriters, overhead projectors, and mimeograph machines, it must have seemed like a trip to the future in a time machine.

One might think that the institutions of the day would have dropped everything to exploit this incredible new technology, but that is not what happened. Not even remotely. By 1972, key members of Engelbart’s team, who had by then moved to Xerox’s Palo Alto Research Center, developed the Alto—a fully realized personal computer. At that moment Xerox stood alone, in command of a technology that would define a whole era of human existence, and in a perfect position in the marketplace to exploit it. But in a case that will be a staple of business schools until the end of time, they decided to abandon computers to focus on their core business: copiers.

Apple came out with what was essentially a consumer version of the Xerox technology. Their machines were treated as novelties or toys. In 1979, *New York Magazine* published an article called “A Computer of One’s Own,” which included gems like: “What can you do with it once you have got it? Well, uh, you can do your taxes on it. But it’s easier to hire an accountant.” An ironic take, in retrospect, on a

technology that was about to destroy both the journalistic and publishing professions as we knew them.

Even after IBM came out with its storied PC in 1981, there were years of debate about whether computers were really a net benefit to business.

Technological revolutions don't happen top-down. They aren't driven by enlightened leaders reaching out to embrace the future.

Institutions are focused on the way they do things now. New technology can only be one of three things: a threat, a distraction, or a tool that can be exploited to improve some element of the current process.

As a result, new technology is welcomed only if it helps do something the institution is already doing. But once it's there, the people working with it start to see new uses for it, and, in time, begin to exploit its full potential. When that happens, it often renders the old way of doing things utterly obsolete.

Technological revolutions happen bottom-up—by infiltration, not institutional imperative.

This pattern is unfolding in corporate education right now, right under our noses.

§2. How to Sneak a Revolution Past the Establishment

If you talk to almost any corporate L&D department these days, you will hear about the use of AI:

“We’ve sped up content creation by 30% using AI!”

“We use AI to do content curation!”

“We have an AI student advisor that recommends courses to take.”

And so on. The persistent theme is that AI is just there to help around the edges of the current process. A useful tool that may revolutionize the *economics* of the business, but not the methodology.

What most L&D organizations cannot imagine is that the technology they see as supplemental will ultimately eliminate the need for almost all of what they do now. L&D organizations are built to create courses. AI will eventually make courses obsolete. Every course is a compromise across the members of an audience who all have different needs, different starting points, different contexts, and different psychologies. AI can create a learning experience for an audience of one in real time. Once people can learn that way, what will be the point of an old-fashioned course?

While some enlightened individuals within L&D organizations see the writing on the wall, those institutions as a whole will continue to be blind to the ultimate impact of AI. So AI will enter these organizations as a tool to improve what it will ultimately destroy. The establishment will put the technology in place that will ultimately smash the curriculum without realizing what is happening.

§3. An Approach Built on Scarcity

The research on how people learn has been remarkably consistent for decades—arguably centuries—and it has *never* supported the way institutions really do training.

Humans learn through direct experience. We form schemas—mental models of how things work—and we update those schemas when reality contradicts our expectations. Learning happens when we try something, see the result, and revise our understanding of the world based on the outcome.

Many people in learning understand this at some level. But few can imagine how we could ever get around the resource constraints that have shaped the standard

approach. With classroom courses and e-learning, we've institutionalized a compromise based on the fact that teacher (i.e., expert) time doesn't scale. The only reason anyone would ever have thought a classroom where a bunch of students sit still and listen to a teacher talk (or an e-learning course where a bunch of students hit the next button and listen to a teacher talk) was a good approach is because we lacked the capacity to do better.

One-on-one tutoring using human teachers doesn't scale. In the 1980s, Benjamin Bloom addressed the issue head-on in a paper based on research that showed students taught one-on-one by expert teachers performed an astounding two “sigma” (that is, two standard deviations, or a little over 95%) better than students taught by conventional classroom methods. He called his paper the “Two Sigma Problem” because the impossibility of scaling this approach was so obvious it barely needed to be mentioned.

The compromises of the past, born of scarcity, have become so entrenched over the centuries that we've forgotten they were compromises. We see courses as the medium through which learning takes place. But there's a new medium in town.

§4. How Simulations Infiltrate: The Predictable Progression

We can predict the pattern by which GenAI will infiltrate L&D. Consider a sales training curriculum as an example.

Stage 1: The Supplemental Sim

A company has traditional sales training: three days of workshops covering objection handling, closing techniques, product knowledge—a mixture of lecture, discussion, and role-playing with other trainees. Someone proposes adding an “AI practice

partner.” Optional. Completely supplementary. It’s still clear the learning happens in the workshop.

The sim fits a traditional role—just an artificial role-play partner, as it were.

Stage 2: The Expansion

Large language models, as it turns out, make great role-players. Trained on thousands of real sales conversations, they remember context, invent plausible backstories and fact patterns on the fly, exhibit realistic human emotional responses, and hit learners with an array of nuanced challenges drawn from the collective experience of thousands of sales organizations all over the world. The AI can generate an unlimited number of individually tailored role plays, with depth and nuance:

1. The prospect burned by competitors and deeply skeptical
2. The procurement manager whose agenda conflicts with the end user’s needs
3. The technical buyer fixated on specifications that don’t matter to business outcomes
4. The friendly prospect who loves the product but has zero budget authority
5. And so on, and on.

As learners experience what the AI can do, the traditional live role play, with three or four people in a breakout session taking turns playing “the customer” with an index card’s worth of “background” to go on, begins to seem pretty impoverished.

Traditional workshop role-play:

Peer playing a prospect: I’m interested in your product, but I’m concerned about price. Can you justify the cost?

Sales Rep: I'm glad you asked. [Enumerates some memorized key features].

Peer playing a prospect: That's very helpful. What about [mentions a feature that was left out]?

Sales Rep: Absolutely. That's been extremely popular with our customers.

Peer playing a prospect: Thanks! I have a much better idea of your product's value-add now.

AI-generated scenario:

Sales Rep: So our analytics dashboard provides real-time visibility into supply chain metrics, which...

Prospect (*interrupting*): I get the concept. Your competitor promised us 'real-time visibility' last year. The project never finished, the dashboard never worked, and it ended up costing us three times the quoted price.

Sales Rep: That sounds... uh... not good. Can you tell me more about what went wrong?

Prospect: Integration with our legacy ERP. They said it would be 'seamless.' It wasn't. So before you show me all the cool features your dashboard has, I need to know how it integrates with SAP R/3 version 4.6C.

Sales Rep: That's a... uh, very specific question. I'd need to check with one of our technical folks...

Prospect: Maybe you should do that right now. My VP wants a decision tomorrow.

Sales Rep: I'm not sure if I...

Prospect: ...and, by the way, you didn't hear it from me, but she's golfing buddies with your competitor's rep, so... jollyng me along is not going to help you. I need a guarantee that you will perform the integration, to spec, within schedule and budget. And I need a *stiff* penalty clause to back that up. Otherwise, we'll be dancing with the devil we know.

Sales Rep (*to AI coach*): Uh... help?

Stage 3: The Simulations Take Center Stage

The first time the new course is run, the AI practice partner is optional. Some learners try it, and they come back to the group sessions energized: "This thing remembered what I said three exchanges ago and called me on a contradiction." "It played a prospect who was nice at first but got hostile when I couldn't answer a technical question."

On the second run of the workshop, everyone wants access.

The facilitators notice something. People are now comparing notes: "The AI gave me a prospect whose company just got acquired and all decisions are frozen—how would you handle that?" The generic scenarios in the live role-play suddenly feel thin.

So on the third run, the training team adds more scenarios. Not just difficult prospects, but specific *types* of difficult prospects: the risk-averse buyer who needs three vendor comparisons, the innovation champion with no budget authority, the technical gatekeeper protecting their own empire. Each one pulls from real deal patterns.

Learners devour them. The feedback is unanimous: "This is the most realistic training we've ever had."

By the fourth run, the AI scenarios aren't just "practice" anymore—they're becoming the primary learning experience. The live role-plays are abandoned. The lecture content that used to fill the mornings is compressed. The bulk of workshop time is now learners working through increasingly sophisticated simulations, with facilitators circulating to debrief and coach.

The training team gets more ambitious. They add scenarios that chain together: "You lost the Miller deal. Six months later, Miller calls you back—the solution they bought isn't working. How do you rebuild the relationship?" They create scenarios where the learner's earlier choices have consequences: botch the technical objection in scenario one, and in scenario four the prospect says, "I talked to your tech team after our last call, and honestly, I'm not confident you understand our architecture."

Each workshop run, the simulations get deeper. More realistic. More consequential. More tailored to each learner's specific gaps.

But the institution still thinks this is a "workshop enhanced by AI." They haven't grasped that the AI has become the primary learning mechanism. The "workshop" is just the scheduling container.

Stage 4: The "Workshop" Morphs into an On-Demand Learning Experience

Eventually someone asks: "Why are we still pulling people from the field for three days? With the AI we can do the whole thing asynchronously."

The workshop becomes virtual. The now-much-shorter lecture sessions are done as web meetings, or even pre-recorded. A feature is added that allows learners to debrief with the coach after a simulation run, and learners report that this is extremely valuable.

Rep: Man, I was doing great until the CTO threw that crazy technical question at me.

AI Coach: Can you think of anything you could have done differently?

Rep: Sure. I could have known the answer.

AI Coach: Is it reasonable to expect you to answer a question like that?

Rep: Not in my book. I'm a sales guy, not a cloud engineer.

AI Coach: Exactly.

Rep: So what am I supposed to... oh! I get it. He *knew* I didn't know the answer—he just wanted to see how I handled the question. And I tried to BS my way around it.

AI Coach: Right. It's no shame to admit a question is beyond your expertise...

Rep: I should have just said, "I don't know, but we've got people who do, and I'll get the answer for you."

AI Coach: That's a good insight. Ready to try again?

Rep: You bet!

The three-day workshop evolves into an always-available, personalized, on-demand, simulation-based practice session.

A radical change has occurred, not just in the content of a course, but in the entire teaching framework. The optional AI-based add-on has subverted the original approach into something unrecognizable—and much better. There's no going back.

§5. The Coaching Jailbreak

More and more simulations are added to the curriculum, and, just as with the sales workshop, they evolve to become more and more central to the learning experience,

while the vestigial trappings of the original teach-by-telling course shrink or disappear entirely.

But that's only the beginning.

Stage 5: The Linked Coach

At first, each simulation has its own AI coach, whose scope is limited to the learner's actions within that simulation. But as simulations get linked together, it only makes sense to link the coaches as well, so that from the learner's point of view they are talking to the "same" coach from one simulation to the next in a sequence.

Learners like this feature—it's nice to be remembered, and the coaching is more powerful when the coach has a broader context to work from.

Stage 6: The Persistent Coach

Eventually someone gets the idea of connecting the coaches even when the simulations are *not* linked. This creates an even better user experience—when learners start a new simulation, the coach already knows them, and vice versa. It starts to feel like they have their own personal AI coach, who shows up whenever they need it.

Having a dedicated coach for each learner in the system starts to offer all kinds of possibilities for honing the behavior of that coach to support the specific learner it is working with. Over time the coach can build a more and more accurate model of the learner—what their long-term goals are, what interests them, what kinds of coaching they respond best to, where they struggle, what they already know, and so on. So not only does the coach become more and more familiar to the learner, it gets better and better at coaching them.

Stage 7: Freeing the Coaches

The next step is inevitable: learners start to ask, “Can I talk to my coach when I’m not in a simulation?” And the answer is... of course. Why not? Free the coaches!

At first it seems like a minor convenience. But it changes everything. There’s something deeply frustrating about having a coach who knows you, who’s helped you through difficult challenges, who’s seen you grow—and then being cut off from that coach the moment you leave the training environment.

Now the coach is available for learning in the flow of work. Need to prep for a meeting? The coach can help. Get a challenging new assignment from your boss? The coach is there for you.

Usage statistics for the coaches start to spike as employees find more and more ways their coach can help them on the job.

Stage 8: The Coach Versus the Curriculum

The unleashing of the coaches now begins to subvert the curriculum in a really interesting way. While it wasn’t anyone’s intent, the coach itself immediately becomes a competitor to every conventional course in the curriculum.

Consider Mia, an HR manager who signed up for a Statistical Methods course because her director asked her to analyze an employee engagement survey. She’s halfway into the course and completely lost:

Mia: I’m taking this Statistical Methods course for a project I’m working on, and I’m completely stuck on something they call heteroske... heteroskedis...

Coach: Heteroscedasticity?

Mia: That! Wow, I can't even say it. You wouldn't believe this crazy formula I'm staring at.

Coach: Does it look something like $\text{Var}(\hat{\beta}) = (X'X)^{-1} X' \Omega X (X'X)^{-1}$?

Mia: That's it!! Help!

Coach: Don't panic—the idea is simpler than it sounds. Heteroscedasticity just means your data is more spread out in some places than others. Like if you're predicting income from age, young people might cluster tightly around similar salaries while older people are all over the map. That unevenness can throw off your predictions, so you need to adjust for it. That's what this formula does.

Mia: Okay... that actually makes sense, but...

Coach: Hang on... before we go deeper on that, can I ask what the project is?

Mia: My director wants me to analyze our employee engagement survey results for a leadership presentation. Stuff about how different locations compare, differences between remote and in-office employees—that kind of thing. And he wants to know if the results are “statistically significant.”

Coach: How was the survey done?

Mia: SurveyMonkey. We got ratings from 1–5 on around 30 questions from about 800 employees.

Coach: Good news! You don't need to worry about heteroscedasticity. You don't need regression at all.

Mia: I don't?

Coach: For questions like “how does Austin compare to Dallas,” you're just filtering and comparing averages. I can show you how to do that with SurveyMonkey's built-in

tools.

Mia: And statistical significance...?

Coach: Let me ask you something first. Did you survey everyone in the company, or just a sample?

Mia: Everyone in our division. We had a 93% response rate!

Coach: Then statistical significance doesn't apply. Significance testing is for when you're trying to generalize from a small sample to an entire population. You surveyed your whole division, which *is* the entire population. It's a kind of perverse way to put it, but when you survey the whole population, *every* result is significant. The numbers are what they are. In any case, common sense is usually enough to interpret a simple survey like this—high-powered statistics are probably not going to help much. Shall we start with the built-in tools and see how far we can get?

Mia: Yes! Thank goodness I don't have to finish this stupid course!

Employees start by getting the coach to help them get through existing courses, but they soon start to realize that, more often than not, they can get custom-tailored tutoring on the exact skills they need to learn from their coach without having to even sign up for a course.

The revolution is underway. Just as AI simulations squeezed out other course content, AI coaches begin to squeeze out entire courses.

§6. The Inversion

The earlier AI simulations were carefully designed artifacts—the sales role-play, the technical troubleshooting scenario. Someone had to build them. But a coach that truly knows its domain doesn't always need pre-built simulations. It can create a

practice environment on the spot, tailored to exactly what the learner needs right now.

Stage 9: The Coach Creates Simulations on the Fly

As learners get used to talking to their coach first whenever they have a learning need, the isolated simulations start to feel almost as disconnected as the traditional courses. The AI tutor can create a light version of a simulation of almost any domain on the fly.

Jordan: Need some help, coach.

Coach: What's up?

Jordan: In our team meeting this morning, my manager asked if anyone knew how to set up a CI/CD pipeline for the new project. And I raised my hand.

Coach: And... do you?

Jordan: I mean... I... I've looked at a YAML file...

Coach: So... no. When do you need it?

Jordan: Demo is Friday.

Coach: Okay... what does the pipeline need to do?

Jordan: Run some tests automatically when someone submits code, then deploy to our staging server if all the tests are passed.

Coach: Okay, let's break that down. The first task is to tell GitHub *when* to run the pipeline, right?

Jordan: I guess, yeah.

Coach: I've set up a practice environment—just type in what you think you would enter into GitHub and you'll see the response you would get. The pipeline lives in a configuration file in your project. Where do you think it goes?

Jordan: Some kind of config folder?

Coach: Try creating it.

Jordan: .github/pipeline.yml

Sim Response: ✖ Workflows must be in .github/workflows/

Jordan: .github/workflows/pipeline.yml?

Sim Response: ✖ Created .github/workflows/pipeline.yml

Jordan: Now what?

Coach: When should it run?

Jordan: When someone submits a pull request.

Coach: Write that.

Jordan: when: pull_request

Sim Response: ✖ Invalid key: 'when'

Jordan: on: pull_request

Sim Response: ✖ Valid trigger

Coach: But you probably want this to run only for pull requests to your main branch, not every branch, right? How would you specify that?

Jordan: on: pull_request: branch: main

Sim Response: ❌ Invalid key: 'branch'

Coach: Did you mean 'branches'? It expects a list of branches.

Jordan: But I only want one branch.

Coach: So, how would you handle that?

Jordan: Oh, duh... just a list with one element. on: pull_request: branches: [main]

Sim Response: ❌ Trigger configured: pull requests to main

Jordan: Okay! So now whenever someone opens a pull request to main, this pipeline will run.

Coach: That'll work. Problem one solved.

Jordan: You know what, this isn't as complicated as I thought it was.

Coach: There's no reason you can't figure it out. Ready for the next step?

Jordan: Let's go.

What just happened? The coach wove a simulation into a natural Socratic conversation, keeping the entire interaction anchored to the practical problem Jordan actually faced. No syllabus, no prerequisites, no modules to click through—just direct, guided practice on the task at hand. Employees start to realize this kind of interaction with their coach is often the quickest way to build the necessary skill to carry out a challenging task or project.

Stage 10: Coaches Running the Show

Heavy-duty simulations still have their uses. They're more immersive, more concentrated, more realistic. They make sense when learners need to master a

complex skill set thoroughly. But when and how should they be utilized?

In the old model, the course structure specified what simulations you did when and in what order. But the rise of the coaches has eliminated most of the pre-set curriculum in favor of in-the-moment decision-making by a coach who knows what activities are going to be most useful for the learner right now. Ironically, standalone simulations may end up being the last vestige of a pre-set curriculum.

And then someone will realize: there's no reason the coach shouldn't make this decision too.

The coach can send a learner to a simulator with a detailed specification of what they're supposed to do there, in much the way a pilot trainer today might send a trainee into a physical simulator to practice some specific scenarios. Simulator training becomes, from the learner's point of view, just another mode of coaching.

§7. The New World of Learning

And with that, the last shreds of course and curriculum are gone. The AI dynamically generates learning experiences for each individual learner. Traditional course boundaries are meaningless. The “negotiation course,” the “statistics course,” the “communication course”—all dissolved. Elements of what each course used to contain will be introduced to individual learners when and if needed, under the orchestration of a dedicated AI coach.

What remains is continuous, personalized learning. Subject boundaries exist only as administrative artifacts.

The old curriculum wasn't just inefficient. It was actively counterproductive—creating artificial barriers between areas of knowledge, forcing everyone through the

same predetermined sequence regardless of individual needs, prior knowledge, and learning patterns.

To a new generation of learners, all of this will seem like a curious anachronism. Taking a course? Sitting through material you already know to get to the part you need? Waiting until Module 4 to ask the question you had in Module 1? How did people learn that way?

In the new world, learning is something that happens through interacting with your personal AI coach. All of the things that used to happen in courses now happen through the coach—but 100% customized to the learner's exact needs at the moment. There is no one-size-fits-all anything. There are no curricula, no courses, no artificial distinction between training time and being on the job.

Just a coach who knows you, ready to help whenever you need it.

§8. The Revolution Nobody Planned

It starts with a simple role-play simulation added to a sales workshop. The simulation gets a coach. The coach gets linked across simulations. The coach gets freed from simulations entirely. The coach starts creating its own learning experiences on the fly. And somewhere along the way, the course—that ancient technology for delivering learning at scale—quietly becomes irrelevant.

This is how technology revolutions actually happen. Not through bold proclamations, but through infiltration. The institutions that are about to be transformed welcome the new technology because it helps them do what they're already doing.

By the time anyone realizes what's happening, it's too late to prevent the destruction of the status quo. The better alternative has already taken root. Learners have experienced what it's like to have a coach who knows them, who meets them where

they are, who helps them solve real problems in real time. Going back to sitting through a course feels like going back to a card catalog after you've used Google.

Will it unfold exactly this way in every organization? Of course not. Some will move faster, some slower. Some will skip stages, others will linger. But the underlying dynamic is inexorable: AI coaching is simply better than courses for most learning needs, and once people experience the difference, they won't go back.

The curriculum won't be abolished. It will be abandoned—gradually, quietly, one learner at a time. The establishment will have built the very thing that makes it obsolete.

References

Bloom, B. S. (1984). The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher*, 13(6), 4–16.

Christensen, C. M. (1997). *The innovator's dilemma: When new technologies cause great firms to fail*. Harvard Business School Press.

Collins, A., Brown, J. S., & Newman, S. E. (1989). Cognitive apprenticeship: Teaching the crafts of reading, writing, and mathematics. In L. B. Resnick (Ed.), *Knowing, learning, and instruction: Essays in honor of Robert Glaser* (pp. 453–494). Lawrence Erlbaum Associates.

Engelbart, D. C. (1968). The demo [Conference presentation]. Fall Joint Computer Conference, San Francisco, CA.

Gibson, W. (1999, November 30). The science in science fiction [Interview]. *Talk of the Nation*, NPR.

McLuhan, M. (1964). *Understanding media: The extensions of man*. McGraw-Hill.

Piaget, J. (1952). The origins of intelligence in children. International Universities Press.

Saylor, M. J. (2012). The mobile wave: How mobile intelligence will change everything. Vanguard Press.

VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. Educational Psychologist, 46(4), 197–211.