



CEO REFLECTIONS

The New Era of Learner Empathy

How AI frees Life Sciences Learning & Development to
become what it was always meant to be

Parth Khanna

Chief Executive Officer & Co-founder, ACTO

△ July 2026

ABOUT THE AUTHOR

Parth Khanna

CEO and Co-founder, ACTO



— Parth Khanna

Parth Khanna is the CEO and Co-founder of ACTO.

Trusted by 14 of the top 50 pharmaceutical companies, high-growth biotechs, and leading medtech organizations, ACTO is redefining how therapies reach patients after regulatory approval.

ACTO unites human intelligence with empathetic AI through an AI-powered Field Excellence Platform and AI SuperAgents to accelerate product launches, strengthen field execution, and deepen customer engagement. By elevating training and coaching while agentifying up to 30% of the work embedded in customer-facing roles, ACTO frees teams to focus on what matters most: human connection, judgment, and strategic execution.

With over 12 years of experience in artificial intelligence and 10 granted U.S. patents in AI, Parth is passionate about applying AI and mobile technologies to transform how Life Sciences companies drive field performance and engage healthcare professionals.

CONTENTS

Inside this reflection

— The Prompt

Why we looked outward before we looked in

01 Missing Planks

The seven gaps in the L&D bridge

02 Rethinking the Bridge

People, process, and technology loosen at once

03 The Lumber Has Arrived

The learner, the manager, the leadership team, and you

04 Don't Forget the Foundation

An intelligent layer across the systems of record

05 Building Responsibly

Responsibility as a design constraint, not a disclaimer

06 A Blueprint for a Bridge You Can Trust

The bridge with every board restored

INTRODUCTION

The Prompt

Eight years ago ACTO approached Life Sciences Learning and Development (L&D) to help make the learning experience better. When the AI shift hit critical scale in early 2025, the instinct across was to look inward: another platform, another content refresh, another system migration. ACTO did the opposite. We looked outward.

We walked the length of the commercialization journey — sales, medical affairs, market access, patient services — and sat with the people who run it: chief commercial, chief marketing, and chief medical officers, and the line-of-business leaders who answer for direct impact on market share and P&L. Across ten cities in North America and Europe, more than a hundred of them told us how the ground was shifting under Life Sciences commercialization in an agentic era.

Then we came home and turned that outward lens on learning and development, where we started and where we've spent the better part of a decade. The perspective was earned on the journey out — and applied to our intimate experience and knowledge of L&D in Life Sciences.

REFLECTION ONE

Missing Planks

Picture three people talking by a river on a quiet afternoon. A scream cuts through the conversation — a child is in the water, struggling against the current. They dive in and pull the child out. Before they can catch their breath, another scream. Another child. Then another. They spend the whole afternoon hauling children from the current, one at a time, soaked and exhausted — until one of them stops, looks up the river, and starts walking. She follows the bank upstream until she finds the bridge. A plank is missing, and children are falling through the gap, one after another. She goes to the forest, cuts wood, and nails the bridge shut. That is L&D today. A new launch, a new indication, a new competitor, a label change: every shift creates a gap and sends another child into the water, and we dive in. We are extraordinary at the rescue. We are also exhausted by it, and the children keep coming.

ACTO's journey to the commercialization edge was a walk upstream. This reflection is an invitation to take that walk yourself, before you plan the back half of this year and 2027, and to look, for once, at the bridge itself.

Because the bridge is your L&D organization. Not just its technology — its people, its agency partners, its whole capability to carry a learner across to competence. And for years, that bridge has had missing or broken planks. Not through anyone's failure.

Through a simpler fact: you cannot nail a plank with wood you don't have. The material to fix these gaps fully did not exist yet.

"You cannot nail a plank with wood you don't have."

THE FRAMEWORK

The Seven Missing Planks

These are the planks that have been missing or broken from the Life Sciences L&D bridge. They fall into three groups: the capability we struggle to build, the impact we struggle to prove, and the capacity we never had to try. Hold them loosely for now — we will walk through them one by one.

GROUP 1 • BUILDING THE CAPABILITY

1 Lack of learner empathy

Learning is one-size-fits-all. It ignores that people learn differently, at different paces, so capability is never reliably built.

2 Coaching constraints

With field teams, ~70% of growth happens on the job with managers — but L&D only influences coaching, and managers are stretched too thin to do it well.

3 No sustainment

"The Forgetting Curve" erases most training within weeks; without reinforcement, capability evaporates before it reaches the field.

GROUP 2 • PROVING THE IMPACT

4 No demonstrable impact

The proof lives in commercial and CRM data L&D doesn't own — so even when training works, it can't be shown. The single biggest barrier.

5 The agency leak

Small teams outsource things that change the bridge — new indication, competitor, launch — to agencies; time, money, and control leak out.

6 Administrative burden

A relentless compliance load — documentation, completion logs, audit-readiness — found nowhere else in corporate learning, burdens Life Sciences.

7 Lack of governance

Bound to validated, compliant tooling (21 CFR Part 11); innovation and compliance can become a trade-off, and the overhead never lets up.

REFLECTION TWO

Rethinking the Bridge

People. Process. Technology. All three constraints are being tested at once right now. Every one of these gaps has survived for the same three reasons: what our organizations believed was possible, what our people had the capacity to do, and what our technology could actually support.

The agentic era is the first moment all three loosen together. Organizationally, we can inject new thinking instead of defending old schemas. On the human side, agents give our people genuine agency — capacity we never had. And technically, we are no longer trapped inside a single system of record; we can build an intelligent layer that sits across all of them, so data finally moves. The stakes are not abstract. Access to clinicians keeps narrowing, every field interaction carries more weight, and the industry's shareholders are watching the gap.

53%

Of HCP engagement with pharma field forces and online channels — and 62% of those still engaging interact with three or fewer companies.

VEEVA SYSTEMS, 2024 ¹

11.2_{mo}

Average time for a new sales rep to ramp to productivity; enablement and training is sales leaders' #1 priority.

SALESFORCE, STATE OF SALES ²

7.6%

Pharma's total shareholder return 2018–2024 vs. 15% for the S&P 500, raising the stakes on commercial capability.

PWC, NEXT IN PHARMA 2025 ³

THE TECHNOLOGY SHIFT, SPECIFICALLY

It's worth being concrete about the third constraint, because it's the one that moved furthest. Four ceilings defined what learning technology could do, and all four just lifted:

CEILING	THEN · WHAT BOUNDED US	NOW · WHAT'S POSSIBLE
Experience	Stuck in point-and-click systems: menus, modules, and portals the learner had to navigate.	Conversational, natural-language interaction; the interface finally speaks human.
Data	Siloed inside each system of record; nothing crossed the gap.	An intelligent layer moves data across every system it touches.
Processing	Rule-based adaptive engines with hard computational limits on what they could ingest.	LLM-based AI chews through orders of magnitude more data, faster.
Understanding	Structured deployment schemas approximated the learner by role and little else.	Models infer what a person actually needs, in the moment.

Stack the four together and you get the technical precondition for empathy: a system that can take in enough about a person to genuinely understand them — and respond in their language, at their moment of need. That is why "empathetic AI" is newly possible, not newly fashionable.

REFLECTION THREE

The Lumber Has Arrived

What follows is the walk upstream — plank by plank, told through the four people who depend on CL&D the most: the learner, the frontline manager, the leadership team, and *you*, the one we often forget.

01 THE LEARNER

Empathetic AI and the End of One-Size-Fits-All

Restores Learner Empathy (Plank 1) & Sustainment (Plank 3)

Here is the idea around which the rest of my reflections orbit. Empathetic AI is AI that aspires to augment humans, not replace them — AI built to understand what a person needs in order to help them meet it. In learning, that has a name: **learner empathy**. And the whole opportunity in front of us is a single question: how do we unite human intelligence with empathetic AI?

We have always told ourselves our learning was learner-centric. It wasn't. It was one size for everyone: check-the-box, uploaded to a system, delivered on a schedule. We tried better flavors along the way: learning in the flow of work, microlearning, learning experiences. Those weren't fads so much as honest attempts that fell short, because the real problem was never the format. The real problem is neurocognitive.


People learn differently, at different speeds, with different skills, strengths, and competency gaps. Meeting each person where they are took data we couldn't gather and processing we didn't have. It was a data problem wearing a pedagogy costume.

Now we have both. And the answer runs on two motions that only work together: **push** and **pull**.

Push is what the organization decides everyone must know — the baseline. It builds on a learner's demonstrated strengths and competencies, lets the strong test out and reach the field faster, and it carries sustainment: the deliberate reinforcement that answers The Forgetting Curve before knowledge decays. The decay is brutal and well documented. B2B sales reps forget roughly 70% of what they learn within a week of training and 87% within a month,⁴ a modern echo of Ebbinghaus's original forgetting curve.⁵

Pull is learner-led, at the point of need. It closes skill and competency gaps by drawing the right knowledge across every system the moment a human reaches for it. This is how people actually learn: we seek knowledge when we need it, not when a schedule assigns it. And pull has a hard prerequisite the industry keeps tripping on: findability. A rep who can't find the current approved answer, or finds two conflicting versions, hasn't been failed by pedagogy; they've been failed by content governance. Pull only works when the approved answer is a single, current, retrievable thing.

The two motions are not parallel tracks, either. They reinforce each other inside the empathy layer: what push teaches shapes what pull retrieves, and what a learner pulls reveals the next gap for push to close.




Push without pull ignores the individual and the moment of need. Pull without push has no guaranteed baseline. The craft is holding both at once — each motion feeding the other.

Map it to how competence really works. Push handles what the organization knows you need. Pull handles your known unknowns — the gaps you can feel. And the system, over time, understands your unknown unknowns — the gaps you can't yet see.

And it compounds. What a learner asks, and what a learner produces in the field, feeds what surfaces to them next, the way a good recommendation feed understands you over time. Learning that becomes more empathetic the more it is used. What makes any of this empathetic, finally, is the data: role, skills and competencies, training completions, and knowledge assessments, HRIS, CRM, the systems of record. The full learner record flows into one unified intelligence — a knowledge graph of the learner — and personalized, well-timed learning flows back

out. Under the hood the two motions differ: push is algorithmic and deterministic, assisted by LLMs; pull is agentic, with governance, compliance, observability, and testing built in.



"We seek knowledge when we need it — not when a schedule assigns it."

— Head of Commercial Learning & Development, top-20 pharma

Coaching Is Where Learning Becomes Development

Restores Coaching (Plank 2)

If you accept 70/20/10 — that roughly seventy percent of a rep's growth happens on the job, after onboarding ends⁶ — then the frontline manager is the most important learning system in your company. And you don't own it.

That is the uncomfortable truth about coaching. It is development, and development is our mandate, but the manager drives it, not us. So our real job is to support the manager, and we spend enormous energy chasing their buy-in, one relationship at a time. Meanwhile, the manager is drowning.

12.1

Average direct reports per manager in 2025, up from 10.9 in 2024 — a near-50% climb since 2013's 8.2.

GALLUP, STATE OF THE GLOBAL WORKPLACE 2025 ⁷

30–60%

Share of frontline managers' time consumed by administrative work and meetings.

MCKINSEY ⁸

90% vs 62%

Leaders who believe they coach at least monthly vs. reps who report receiving it — the coaching perception gap.

SALES ENABLEMENT RESEARCH ⁹

Layoffs have stretched territories. Admin work has climbed as resources have thinned. And a large share of managers are player-coaches, carrying their own quota and holding their own key HCPs. The time to coach well simply isn't there. Worse, coaching lives in systems we don't control: CRM, field ops, the call report. Which is exactly why the CRM vendors are walking into coaching from the system-of-record side.¹⁰ If we try to meet them there, on the channel, we lose.

Training does not stop at certification. A passing score is a snapshot; competence in front of an HCP is a moving picture. Coaching is the mechanism that carries a person from certified to capable: the ride-alongs, the field coaching reports, the

debrief after the call that went sideways. It is where the organization's thinking, delivered through push, finally meets a human brain having its own experiences in the field — and where pull gets its cues, because a good coach surfaces the gaps a learner can't yet name.

Only about a quarter of sales managers spend five or more hours a month coaching,⁹ and in our own industry roundtable conversations, a recurring theme was that a large share of field professionals say they receive feedback — but not real coaching: development that is episodic, subjective, and dependent on whichever manager they happened to draw.²¹ L&D absorbed the blame for inconsistent development it never had the authority to fix. That is the plank: not that coaching is absent, but that it is unowned, unmeasured, and starved of time.

So L&D's play has to be different, and it is defensible. Reduce the manager's load, and raise the quality of what coaching does happen. The restored plank runs on three moves:

INSIGHT

See the gap, don't guess it.

Skills and competency data — assessments, role-play scores, field signals — flows to the manager, via an agent, before the ride-along, so coaching starts from evidence, not anecdote.

PRACTICE

Reps without the calendar cost.

AI role-play and certification give every learner objective, private, always-available practice — and give the manager an observable signal to coach against.


ASSIST

The agent carries the rote.

Coaching prep, the field coaching report, tracking and follow-ups move to the agent, so the manager spends their scarce hours on the one thing only they have: tacit, in-field judgment.

Notice what the three moves protect. The manager is not being automated out of coaching — they are being returned to it. The agent takes the administration that made coaching unaffordable; the human keeps the relationship, the judgment, and the call. And every coaching interaction writes back into the learner's model, so push

and pull re-personalize, which means coaching is what makes learner empathy continuous after onboarding ends. The old silos — learning here, coaching there — existed because our tools forced them; the data never crossed the gap. Now it can.



"Your most important learning system is a manager who doesn't have time to teach."

— SVP, Commercial Excellence, global biopharma

From Execution Shop to Strategic Partner

Restores Impact (4), the Leak (5), Admin Burden (6) & Governance (7)

Your commercial, medical, and market access leaders rely on you to advance their objectives. Yet L&D keeps getting cast as the execution-and-compliance shop: necessary, rarely strategic. And every two or three years, you need a new story to justify the function to them. Here is the reframe.

Your strategic problem is not "how do we engage HCPs." That is the chief commercial officer's problem — or another CXO's — and the systems vendors already own the channel it runs on. Chase it and you walk into a fight you can't win. Your strategic problem is showing, in their terms, how your function moves that agenda: proving that field capability is the reason engagement stands apart. Building that capability at scale is strategic work — but leaders only see it that way when you can draw a straight line from capability to outcome.


"We help reps execute better" is not a boardroom story. "Our field's capability is the reason our KOL engagement stands apart — and here is the data trail" **is**.

Four planks stand between you and that line.

4 **Demonstrable impact — the measurement we could never do**

We never proved training's effect on performance. Not for lack of trying, but because the proof lived in data we didn't own. Impact shows up in the CRM, in commercial operations, in systems L&D has no control over. The industry numbers say the quiet part out loud: only about 29% of enablement teams can tie programs directly to revenue, and 35% of chief sales officers admit they don't know what measurable improvement they seek from training.⁹

Name what proof would actually look like. Call it the **Impact Line**: one measurable chain. We built this capability in our reps → their field behavior changed → the HCPs they engaged moved further along their journey — trial, adoption, advocacy → and it showed up in the business. Today, L&D can measure only the first link — completions and quiz scores — and nothing after. Everything else is a claim; the Impact Line is evidence. Of the seven planks, this one is the single biggest barrier to the strategic seat.



What changed is plumbing. Data warehouses — Snowflake, Databricks, and their peers — let the learning data and the commercial data finally sit together,¹¹ so we can go past Kirkpatrick's first two levels into the third and fourth,¹² with metrics drawn straight from CRM fields. Voice-of-customer becomes the outer signal on the same loop.

And we measure it honestly — against the human baseline, which was never zero. Reps forget, tire, and carry bias; we already accept a knowledge-assessment score as "good enough" proof a human is ready. AI deserves the same test: a real before-and-after against the old way of working, not a demand for perfection. Honesty also means conceding what a warehouse gives you: correlation you can act on and defend, not laboratory-grade causation.¹² Say that plainly and the CFO leans in instead of pushing back.

5 The agency leak — and the queue that feeds it


There is a leak in the function: time, money, and resources draining out to external agencies. Two things cause it. L&D lacks the broader strategic perspective, and it lacks internal bandwidth. The pattern is familiar: any change — new indication, new competitor, new launch — sends a ripple, and the team reacts. To absorb the ripple, they call an agency. Reactive work becomes outsourced work, quarter after quarter. Industry-wide, roughly 62% of training instruction and facilitation is now outsourced.¹³

Agentic AI changes the economics. The reactive, produce-it-again work can move in-house to agents, and the leak closes. As for the strategic-perspective deficit, it closes the same way ours did: by walking upstream to the commercialization edge. The journey isn't a detour from the day job; it's the cure for it.

Content velocity — readiness at the speed of the market

When a label changes or a competitor launches, the message and the training both wait in the same queue: medical-legal-regulatory review. In mid-to-large pharma those cycles routinely stretch 50–60 days per asset,¹⁴ and while the queue backs up, the field is selling on last quarter's story. The absurd endpoint of the current model: 77% of approved pharma content is almost never used by field teams, even as production volumes keep climbing.¹⁵

This is not an argument for weakening review. It's an argument for redesigning the workflow around it. Agentic systems that pre-check claims against the approved registry, assemble first drafts from already-approved components, and route only genuine judgment calls to human reviewers are already compressing regulated timelines by 50–65% in adjacent workflows.¹⁶ The plank isn't "make more content." It's "make readiness arrive when the market moves" — and stop paying agencies to re-produce what an agent can assemble from approved parts.



Before: 50–60-day cycles per asset → After: 50–65% compression in regulated agentic workflows

Figure 5 · Closing the leak — agentic assembly, human judgment reserved for exceptions

6 Administrative burden — the root of the leak

Underneath the leak sits its root cause: a relentless compliance load found nowhere else in corporate learning. Documentation. Completion logs. Certification cycles that never stop. Audit-readiness, always.¹⁷ None of it is optional, and none of it is strategic; it simply consumes the hours strategic work would need. Which is exactly why every ripple gets outsourced: a team that spends its capacity keeping records has none left to absorb change.

The fix is not more headcount. It's agents that automate the certification chase, the quizzes, the record-keeping, and the reporting — logged, validated, and audit-ready by default. When the administrative weight shifts off human shoulders without weakening the record, the leak's root cause goes with it.

7 Governance — the constraint that's uniquely ours

Planks 6 and 7 together are a weight that no other industry's L&D function carries. Beyond the daily load sits the structural constraint: we are bound to validated, compliant tooling — 21 CFR Part 11 and its relatives: audit trails, electronic signatures, validated releases, read-and-understood tracking.¹⁷ And the envelope is expanding: GAMP AI guidance, EU Annex 22, and the FDA's own agentic-AI posture now shape how AI itself must be validated in GxP contexts.¹⁸

For years, governance was a gate; it throttled what we were allowed to try. The shift is to treat governance as an operating model instead: something that lives inside the system, so innovation and compliance stop being a trade-off. That is where responsible AI stops being a slogan and becomes architecture.

"'We help reps execute better' is not a boardroom story. 'Here is the data trail from capability to outcome' is."

— SVP of Sales, top-10 pharma

The One Lens We Often Leave Out

We've talked about your learners, your managers, your leaders. The fourth lens is the one we often forget: **you**.

The same shift that reshapes the work reshapes the people who design it. The reactive operator — producing content and tracking completions — is not the role that matters next. The role that matters is the one who brings responsible, agentic AI into the enterprise and operationalizes it for the field. The function that teaches the whole organization to learn should be the first to learn AI itself. And the gap is wide open: 89% of firms say their workforce needs better AI skills, yet only 6% have meaningfully upskilled.¹⁹ Nobody in the enterprise is better positioned to close that gap than L&D — and almost nobody has claimed it.

So run the exercise on your own job. Separate what an agent can take — the assembly, the tracking, the first drafts, the rote retrieval — from what only you can hold: judgment, strategy, the human relationships, and the call when something doesn't look right. What's left after you strip out the rote isn't a smaller job. It's a bigger one, if you claim it.

ONE JOB DESCRIPTION, SEPARATED

WHAT THE AGENT TAKES

Content assembly & first drafts

Completion tracking & logs

Rote retrieval & reporting

Re-production after every change

WHAT YOU KEEP

Judgment & strategy

The human relationships

The Impact Line story

The ethical override — the call

Figure 6 · Strip out the rote and what remains isn't a smaller job — it's a bigger one, if you claim it


That is career advancement in plain terms. The L&D leader who can draw the Impact Line from capability to outcome, and who can govern AI responsibly, is the one who earns the strategic seat.

REFLECTION FOUR

Don't Forget the Foundation

Architecture is everything. One idea holds this section together: we are no longer bounded by systems of record.

For years, every capability lived trapped inside its own system, and the tools quietly shaped the teams around them — learning in the LMS, engagement in the CRM, content behind MLR, coaching in field ops, none of them speaking. The agentic era puts an intelligent layer across those systems. Data moves. And learning, coaching, and measurement finally share one ground.



Three properties make the layer trustworthy rather than merely clever. **Bounded:** every agent operates inside a defined role and scope. **Permissioned:** agents inherit the data permissions of the systems they touch — nothing more. **Governed:** every action is logged, auditable, and escalated to a human wherever the consequence is real. These are the same principles regulators themselves are converging on.¹⁸

REFLECTION FIVE

Building Responsibly

Everything I've shared so far assumes you are dealing with AI you can actually trust. And in the Life Sciences industry, "responsible" isn't an ethics label borrowed from tech; it's defined by who you serve. The bar is higher here because the stakes are patients, regulators, and human judgment. A marketing generator that hallucinates is an embarrassment; an agent that puts the wrong thing into a field workflow is a different order of risk. So responsible AI can't be a paragraph at the end. It is a design constraint on every plank we've walked:

- **Empathy** depends on learner data handled responsibly — collected with consent, minimized, and never quietly repurposed to judge the people it was meant to help.²⁰
- **Coaching-assist** depends on the human keeping the consequential call.
- **Measurement** depends on auditability.
- **Velocity** depends on the approved-claims boundary holding.
- **Governance** has to live inside the system, not gate it from outside.

The regulators, notably, are moving in the same direction. The FDA deployed agentic AI to its own workforce in December 2025 with an explicit governance blueprint — human oversight, transparency, and data that stays inside its permission boundary — and the EMA's lifecycle guidance emphasizes the same principles.¹⁸ The direction of travel is set. The only question is whether your function leads it or absorbs it.

"Responsibly' is a design constraint,
not a disclaimer."

REFLECTION SIX

A Blueprint for a Bridge You Can Trust

What does a modern, empathetic, agentic L&D organization look like — plank by plank, across people, process, and technology?

THE BRIDGE WITH EVERY BOARD RESTORED

PLANK	PEOPLE	PROCESS	TECHNOLOGY
1 • Learner empathy	Learners met at their pace, strengths, and gaps	Push baseline + pull at point of need, unified	Empathy layer over role, skills, HRIS, and CRM data
2 • Coaching	Managers coach with insight, not instinct	Agents carry prep and rote; humans keep judgment	Skills & competency data flowing to the manager via an agent
3 • Sustainment	Knowledge that survives the month	Reinforcement native to push, timed to decay	Adaptive spaced-repetition at scale
4 • Demonstrable impact	L&D fluent in commercial metrics	The Impact Line: capability → behavior → HCP journey → business	Learning + commercial data joined in the warehouse
5 • The agency leak	Team time shifts to strategic work	Reactive re-production and the MLR queue move to agentic assembly	Agentic content assembly from approved parts

6 • Administrative burden	Humans freed from the certification chase	Records, quizzes, and reports automated — audit-ready by default	Agentic record-keeping on validated rails
7 • Governance	Compliance owned as a capability, not a chore	Governance as operating model, not gate	Part 11-native audit trails inside the system

IN CLOSING

Until Next Time

I hope you have enjoyed this pause to reflect on where we are as an industry, in this agentic world. Before you plan the back half of this year, and before you plan 2027, take this walk again. Stop diving for one child at a time. Step back, and look at the bridge. The planks have been missing for years — not because we failed, but because the wood didn't exist. It does now.



Parth Khanna

CEO & Co-founder, ACTO · July 2026

References & Notes

- 1 Veeva Systems, HCP access research (2024): HCP engagement with pharma field forces and online channels at 53%; 62% of engaging HCPs interact with three or fewer companies; 27% of HCPs find pharma communication relevant and personalized.
- 2 Salesforce, State of Sales (survey of 5,500+ sales professionals): improving enablement and training as the top priority for sales leaders; average new-rep ramp of 11.2 months.
- 3 PwC, Next in Pharma 2025: equal-weight index of 50 pharma companies returned 7.6% total shareholder return 2018–Nov 2024 vs. 15% for the S&P 500.
- 4 Gartner research on B2B sales knowledge retention: ~70% of training content forgotten within one week; 87% within one month.
- 5 Ebbinghaus, H. (1885), Über das Gedächtnis — the original forgetting curve; modern replications confirm steep decay absent spaced reinforcement and retrieval practice.
- 6 McCall, Lombardo & Eichinger, Center for Creative Leadership — the 70/20/10 model. The percentages are directional, not prescriptive; we use the model as a lens, not a law.
- 7 Gallup, State of the Global Workplace 2025 (citing BLS data): average direct reports per manager rose from 10.9 (2024) to 12.1 (2025); 13% of managers now oversee 25+ reports.
- 8 McKinsey & Company, frontline-manager research: 30–60% of frontline managers' time consumed by administrative tasks and meetings.
- 9 Sales enablement and CSO research (multiple sources): ~90% of leaders believe they coach at least monthly vs. ~62% of reps reporting it; ~29% of enablement teams can tie programs to revenue; 35% of CSOs cannot name the measurable improvement they seek from training; ~25% of sales managers spend 5+ hours/month coaching.
- 10 Veeva launched AI Agents for Vault CRM and PromoMats (Dec 2025) with coaching, KAM and back-office agents on the 2026 roadmap; Gartner's inaugural Magic Quadrant for Revenue Enablement Platforms (Nov 2025) and subsequent market consolidation (Seismic–Highspot; Showpad–Bigtincan) mark the CRM-into-coaching trend.
- 11 Veeva added connectors to Snowflake, Databricks, Redshift and Power BI, enabling BI/AI on Vault data alongside commercial data.
- 12 Kirkpatrick, D. — four levels of training evaluation. Level 3–4 attribution remains genuinely hard (see Thalheimer's critiques); warehouse-joined data yields defensible correlation, not laboratory causation. Pair with Phillips ROI / Return on Expectations where rigor is demanded.
- 13 Training Magazine, 2025 Training Industry Report: ~62% of training instruction/facilitation outsourced; ~7% of training budgets to outsourcing.
- 14 Indegene (via Vodori) benchmarking: MLR review cycles of 50–60 days per content piece at mid-to-large pharma under current workflows; European average ~20 days at 1.3 revision cycles.
- 15 Veeva Pulse: 77% of approved pharma content is almost never used by field teams; content production grew 7% globally and 29% in the U.S. (2023).
- 16 McKinsey & Company: leading pharma companies compressed regulatory submission timelines by 50–65% through AI-enabled automation and workflow redesign.
- 17 21 CFR Part 11 (electronic records and signatures); EU Annex 11; read-and-understood and certification tracking obligations for regulated training records.
- 18 FDA deployment of agentic AI to all employees (Dec 1, 2025) following the Elsa LLM tool, with explicit emphasis on human oversight, security, transparency, and permission-bounded data; EMA/HMA final guidance on AI across the medicinal product lifecycle (2025); ISPE GAMP AI guidance and EU Annex 22 for validating AI in GxP contexts.

- 19** BCG: 89% of firms say the workforce needs better AI skills; only 6% have meaningfully upskilled. McKinsey: 46% of C-suite cite skill gaps as the top barrier to AI value.
- 20** Ingesting HRIS and CRM data triggers GDPR/CCPA obligations: DPIAs, consent, data minimization, and caution around automated decision-making about employees.
- 21** ACTO global executive roundtable series (2025–26): conversations with 100+ commercial, medical, and market access leaders across ten cities in North America and Europe on responsible, agentic AI in life sciences commercialization; field-development themes include feedback-without-coaching and inconsistent, manager-dependent development. Blinded quotations in this paper are drawn from these conversations, edited for concision, and attributed by role only.

Figures labeled as vendor-reported in the underlying research (e.g., role-play onboarding-time reductions) are deliberately excluded from this paper. All market statistics are date-stamped to their source period; the regulatory and vendor landscape is moving quickly and claims should be re-verified before major republication.

ABOUT ACTO

Intelligent Field Excellence for Life Sciences

ACTO is an AI-forward technology company purpose-built for Life Sciences. Trusted by 14 of the top 50 pharmaceutical companies, high-growth biotechs, and leading medtech organizations, ACTO unites human intelligence with empathetic AI through an AI-powered Field Excellence Platform and AI SuperAgents — accelerating product launches, strengthening field execution, and deepening customer engagement.

By elevating training and coaching while agentifying up to 30% of the work embedded in customer-facing roles, ACTO frees teams to focus on what matters most: human connection, judgment, and strategic execution.



AI SUPERAGENTS

Learn more about SuperAgents

Scan the code to see how ACTO's AI SuperAgents put humans at the center of Life Sciences field excellence.



Intelligent Field Excellence for Life Sciences