




— A PRACTICAL GUIDE FOR LEARNING LEADERS

Where AI Belongs — and Where It Doesn't

Building long-term value with AI in learning: what to do today, what to expect in three to five years, and how your learning platform should grow alongside it.

For L&D leaders, educators, learning administrators, instructional designers and executives responsible for how their organisation learns.

BEFORE YOU BEGIN

The question isn't whether to use AI. It's where.

AI can give administrators back hours every week. It can also quietly do the thinking that learning depends on.

Every organisation that trains people is now being asked about AI: by boards, by vendors, by learners already using it. The pressure is to adopt fast. The risk is adopting in the wrong places — automating the struggle that builds skill, or handing judgment to a system that can't be held accountable for it.

This guide draws a clear line. It sets out where AI creates durable value in learning, where it doesn't belong, the lowest-hanging fruit you can pick today, and what we expect to be possible in three to five years. It closes with the governance practices that keep it safe, and how a learning platform should grow alongside AI rather than be replaced by it.

It continues the thread of this series: *Beyond the Feature Checklist* (Vol. 1) argued for choosing technology by long-term value; *The Residual Gap* (Vol. 2) argued for rewarding the people who create it. Volume 3 applies both ideas to AI.

INSIDE

01	The principle: amplify, don't replace	03
02	Where AI belongs in learning	04
03	Where it doesn't belong	05
04	AI inside the LMS	06
05	The AI fit test TOOL	07
06	Lowest-hanging fruit today	08
07	The next three to five years	09
08	Guardrails & governance	10
09	Your AI policy, in ten parts TOOL	11
10	Measuring long-term value	12
11	Momentum's AI roadmap	13
12	Growing alongside AI	14

THREE HORIZONS

Now — low-risk efficiencies you can switch on this year. **Next** — connected, assisted learning in one to three years. **Later** — adaptive, agent-connected learning in three to five.

01 THE PRINCIPLE

Use AI to amplify human learning — never to replace the learning itself.

Learning happens through effort: retrieving, practising, getting it wrong and trying again. Researchers call it *productive struggle*. Anything that removes the struggle also removes the learning.

FIG. 01 • THE AI ROLE SPECTRUM

Automate**AI DOES IT**

Tagging, translation, scheduling, reminders

LOW STAKES • ROUTINE

Assist**AI DRAFTS, HUMAN DECIDES**

Quiz drafts, marking suggestions, reports

Augment**AI COACHES THE LEARNER**

Hints not answers, practice, guided learning paths

Human only**NO AI IN THE LOOP**

Final grades, mentorship, wellbeing, certification

HIGH STAKES • HUMAN JUDGMENT

RULE 1**Give time back**

Automate the predictable work around learning so people can spend more time on the learning itself.

RULE 2**Keep humans accountable**

When a decision affects a person's grade, credential, job or wellbeing, a named human makes it.

RULE 3**Protect the struggle**

AI should prompt thinking, not supply the finished thought. Hints over answers; questions over summaries.

ALIGNED WITH INTERNATIONAL AND CANADIAN FRAMEWORKS

UNESCO — *Guidance for Generative AI in Education and Research* (2023) and the *AI Competency Frameworks for Students and Teachers* (2024)

OECD — *AI Principles* (2019, updated 2024), endorsed by Canada

Council of Europe — *Framework Convention on AI and Human Rights* (2024), signed by Canada in 2025

Canada's privacy regulators — federal, provincial and territorial *Principles for Responsible, Trustworthy and Privacy-Protective Generative AI Technologies* (2023)

Government of Canada — *AI for All*, the national AI strategy (2026), with a focus on trust, safety and AI literacy

Ontario — *Strengthening Cyber Security and Building Trust in the Public Sector Act* (2024), requiring human oversight of AI in bodies including school boards

All call for a human-centred approach, with people kept accountable for decisions that affect learners.

02 WHERE AI BELONGS IN LEARNING

Four places AI earns its keep.

01 LEARNERS

Personalised tutoring & adaptation

AI can track each learner's errors and pace in real time, offering targeted practice and alternative explanations that one instructor can't provide to a whole cohort at once.

BEST PRACTICE Ground the tutor in verified content and design it to give hints, not answers.

02 INSTRUCTORS

Workflow support

Brainstorming lesson plans, summarising source material, generating varied perspectives and scenarios, and drafting routine communications.

BEST PRACTICE Treat every output as a first draft that a subject expert reviews.

03 LEARNERS

A thought partner

Outlining ideas, clarifying difficult concepts, and self-quizzing before an assessment — used to prepare for thinking, not to replace it.

BEST PRACTICE Make expectations explicit: when AI is allowed, how to disclose it.

04 ADMINISTRATORS

Operational scale

Scheduling, enrolment, translation, reporting and spotting systemic learning gaps across programs, regions and cohorts.

BEST PRACTICE Start here. It's the lowest-risk, highest-return place to begin.

THE EVIDENCE • WHEN AI IS DESIGNED TO TEACH

2×

In a 2023 randomised trial in Harvard's introductory physics course (194 students), learners using an AI tutor built on teaching principles — scaffolded steps, verified solutions, timely feedback, self-pacing — achieved median learning gains more than double those of an in-class active-learning session, in less time. Kestin et al., *Scientific Reports*, 2025.

03 WHERE AI DOESN'T BELONG

Four lines we don't cross.

01

Replacing reading, writing and thinking

Drafting, deep reading and original problem-solving are the work of learning. AI that does them for the learner removes the very effort that builds lasting capability.

02

Bypassing foundational assessment

Using AI to produce the essay or solve the equation short-circuits the iterative process, leaving knowledge that looks complete but doesn't hold up without the tool.

03

Over-reliance in early learning

For young learners, excessive screen and AI use can crowd out emotional, social and hands-on development. UNESCO recommends a minimum age of 13 for classroom use of generative AI tools.

04

Replacing human relationships

Empathy, mentorship and social-emotional guidance come from people. AI can connect a learner to the right person; it shouldn't pretend to be one.

THE EVIDENCE • WHEN AI IS DESIGNED TO ANSWER

+48%

Practice performance with an unrestricted chatbot

-17%

Exam performance for the same students once AI was removed

≈0

Exam harm when the tutor gave hints and teacher guidance instead

Field experiment with nearly 1,000 high-school maths students. Bastani et al., "Generative AI without guardrails can harm learning," *PNAS*, 2025.

04 AI INSIDE THE LMS

The platform is where principle becomes practice.

An LMS sits at the centre of content, learners, records and reporting. That makes it the natural place to apply AI — and the place where boundaries matter most.

WHERE AI FITS	WHERE IT DOESN'T — OR NEEDS A HUMAN
Content tagging. Tag courses, videos and quizzes with skills, metadata and standards as they're uploaded.	Final grading of nuanced work. Essays, discussions and creative projects need human interpretation for the final mark.
Smart content delivery. Recommend remedial or advanced modules based on assessment results.	Credential and certification decisions. Pass/fail on a licence or qualification stays with a qualified person.
Early-warning analytics. Flag learners at risk from logins, submissions and performance trends.	Automated consequences. A risk flag should start a human conversation, never trigger an automatic penalty.
Drafting assistance. Help authors format modules, write discussion prompts and seed question banks.	Unreviewed publishing. AI-drafted content goes live only after subject-expert review.
Translation & accessibility. Translate courses and forums; draft captions and alt text.	Emotion or attention monitoring. Inferring learners' emotions is prohibited in EU education settings and erodes trust everywhere.
Marking assistance. Suggest scores and feedback on short and long answers against a rubric.	Ungoverned data access. AI should only reach approved sources, under the same permissions as the user.


THE WORKING RULE

AI drafts, suggests and flags. People decide, publish and grade. Every AI-assisted decision should be visible, explainable and reversible by the person accountable for it.

05 THE AI FIT TEST • A TOOL FOR ANY PROPOSED USE

Two questions decide how AI should be used.

Is the task **part of the learning**, or work around it? And how **high are the stakes** for the person affected? Plot any proposed use and the right level of AI involvement follows.



HOW TO USE IT

Run every new AI feature or request through these two questions before it's switched on. If you can't agree which quadrant it belongs in, treat it as the higher-stakes one.

06 HORIZON 1 • NOW

The lowest-hanging fruit: value you can unlock this year.

These uses sit mostly in the “automate” and “assist” quadrants. They’re low risk, quick to adopt, and return time immediately — the right place to build confidence before moving further.

USE	WHAT IT GIVES BACK	EFFORT	RISK	HUMAN ROLE
Translation of courses & forums	Reach learners in their own language without separate builds	Low	Low	Spot-check key terms
Content tagging & search	Findable content; skills mapped automatically on upload	Low	Low	Approve the taxonomy
Quiz & question-bank drafts	Hours saved per course; more practice items	Low	Low	SME review before use
Marking assistance	Faster feedback on short and long answers	Medium	Medium	Instructor confirms every mark
Accessibility drafts	Captions, transcripts and alt text at scale	Low	Low	Review for accuracy
Reminders & nudges	Better completion with less admin chasing	Low	Low	Set tone and frequency
Report summaries	Plain-language insight from dashboards for leaders	Low	Low	Validate before sharing
Survey & evaluation drafts	Better-designed evaluations tied to outcomes	Low	Low	Align to your measures

START SMALL

Pick two uses, one team, one quarter. Measure time saved before and after.

STAY INSIDE THE PLATFORM

Use AI built into governed systems, not copy-pasting learner data into public chatbots.


TELL PEOPLE

Let learners and staff know where AI is used and how to raise concerns.


07 HORIZONS 2 & 3 • NEXT AND LATER

What we expect to become possible in three to five years.

Forecasts in AI age quickly, so we focus on direction rather than dates. Based on what's already working in pilots today, these are the capabilities we expect organisations to be using routinely.



Six capabilities we expect within five years



08 GUARDRAILS & GOVERNANCE

Trust is the real infrastructure.

AI in learning touches personal data, assessment and credentials. The organisations that get long-term value treat governance as a foundation, not an afterthought. Four reference points are worth knowing.

FRAMEWORK

NIST AI Risk Management Framework

A voluntary framework organised around four functions: **Govern, Map, Measure** and **Manage** AI risk.

STANDARD

ISO/IEC 42001

The international standard for an AI management system — policies, roles, risk assessment and continual improvement.

REGULATION

EU AI Act

Classes AI used for admission, assessment and exam proctoring in education as high-risk (obligations apply from 2 December 2027) and already bans emotion recognition in education.

GUIDANCE

UNESCO & national guidance

Human-centred use, age-appropriate access, data protection and teacher capacity-building. Privacy law (e.g. PIPEDA, GDPR) applies throughout.

Eight guardrails every AI-enabled learning platform should have

- ☐ **Approved sources only.** AI answers from content and data you've sanctioned.
- ☐ **Permission-aware.** AI sees only what the user is allowed to see.
- ☐ **Human in the loop** for grades, credentials and published content.
- ☐ **Transparency.** Learners know when AI is involved.
- ☐ **Data protection.** Clear terms on whether learner data is retained or used for training.
- ☐ **Audit trail.** Logs of AI suggestions and the human decisions taken.
- ☐ **Bias & accuracy checks** on outputs that affect people.
- ☐ **An off switch** per feature, per organisation.

09 YOUR AI-IN-LEARNING POLICY, IN TEN PARTS • A STARTER CHECKLIST

A short policy beats a long silence.

Most learners and staff are already using AI. A clear, plain-language policy turns guesswork into good practice. Use these ten parts as a starting outline.

- ☐ **01 Purpose and principles**
Why you use AI, and the human-centred principles that limit it.

- ☐ **02 Approved tools**
Which AI tools and platform features are sanctioned, and for what.

- ☐ **03 Learner use**
When learners may use AI, when they may not, and how to disclose it.

- ☐ **04 Assessment integrity**
How assessments are designed and marked in an AI-capable world.

- ☐ **05 Human decision points**
Which decisions always require a named, accountable person.

- ☐ **06 Data and privacy**
What data AI may access, retention, and which providers process it.

- ☐ **07 Content quality**
Review standards for AI-assisted content before it's published.

- ☐ **08 Accessibility and equity**
Ensuring AI widens access rather than creating new gaps.

- ☐ **09 Transparency and appeals**
How people learn where AI was used, and how to challenge an outcome.

- ☐ **10 Review cycle and ownership**
Who owns the policy, and how often it's revisited (at least annually).

General guidance only. Align your policy with the laws, regulators and collective agreements that apply to your organisation.

10 MEASURING LONG-TERM VALUE

Measure what AI changes, not how much AI you use.

Usage counts tell you AI is switched on. They don't tell you it's creating value. Track outcomes at each level of the organisation — the same five dimensions from Volume 1.

LEVEL	WHAT TO MEASURE	EXAMPLE INDICATOR
Learners	Engagement, mastery and time to competency	Assessment gains; days to proficiency
Instructors	Time returned and feedback speed	Hours saved per course; marking turnaround
Administrators	Manual work removed	Tickets and manual enrolments per month
Leaders	Decision quality and early warning	At-risk learners reached before failure
Organisation	Transfer to performance and compliance	On-the-job behaviour change; audit readiness
Trust	Accuracy, fairness and confidence	Override rate on AI suggestions; learner sentiment

FIG. 02 • MOVE UP THE EVALUATION LEVELS



Kirkpatrick's four levels of training evaluation. AI makes levels 3 and 4 — the ones that prove long-term value — far cheaper to measure, through better-designed surveys and connected performance data.

11 MOMENTUM'S AI ROADMAP

Administrators first. Learners next. Humans always.

Our priority is the work that surrounds learning: giving time back to the people who run programs, then using that foundation to guide learners — always under the guardrails set out in this guide.

WHERE WE ARE TODAY

UNBOUND • LMS

AI translation, built in

Full integration with the DeepL API translates course content and discussion forums, so learners can participate in their own language.

EXAMDECK • ASSESSMENT

AI-assisted marking

Full integration with OpenAI suggests marks and feedback on short- and long-answer questions, with the instructor reviewing and confirming results.

WHERE WE'RE GOING

Priority 1 • Administrators

- 1 Lifecycle efficiency**
Manage learners from first onboarding through lifelong continuing education with far less manual effort.
- 2 No more data silos**
Agent-to-agent API connections between Unbound and the systems you already use.
- 3 Guardrails by design**
Control what information AI can access, from which sources, with continued data security.

Priority 2 • Learners


- 4 Guided learning**
Suggested resources and learning paths; AI mentorship; connecting learners with an internal SME.
- 5 Engagement that lasts**
Drafted cohort discussion questions, quizzes and workplace “chalk talks” that keep key topics top of mind.
- 6 Evidence of impact**
Meaningful reports and survey questions tied to recognised measures of performance and knowledge transfer.

Roadmap items describe direction, not committed delivery dates, and may change as technology and customer needs evolve.

12 GROWING ALONGSIDE AI

Your LMS shouldn't be replaced by AI. It should grow with it.

AI will change faster than any organisation can re-platform. The durable advantage is a learning system that holds your content, people, records and guardrails in one place — and adds AI capability as you're ready for it.



STAGE	WHAT THE ORGANISATION NEEDS	HOW MOMENTUM'S LMS SUPPORTS IT
Foundation	Content, learners and records in one governed place; an AI policy	Unbound as the single system of record, with roles and permissions
Assisted	Low-risk efficiencies that save time this year	Built-in translation and assisted marking; drafting tools as they arrive
Connected	Integration with HR, CRM and credentialing; trusted data	API and agent-to-agent connections with source and access guardrails
Adaptive	Personalised, measured learning across the employee lifecycle	Suggested paths, AI mentorship, SME matching and impact reporting

LONG-TERM VALUE

You don't have to choose between today's platform and tomorrow's AI. Start at your stage, prove value, then climb — with a partner who stays accountable for the result, as we argued in Volume 1.

The best use of AI in learning is the one that makes more room for people.

Let AI handle the routine.

Let it coach, not answer.

Keep judgment with people.

Grow one stage at a time.

AI will keep changing.

Your learning should keep getting better.

NEXT STEP

Talk to us about your AI starting point

We'll walk through the fit test with your team, identify two quick wins, and show you where Unbound and ExamDeck fit today.

momentumitgroup.com

SOURCES & NOTES

Sources

1. Kestin, G., Miller, K., et al. "AI tutoring outperforms in-class active learning: an RCT introducing a novel research-based design in an authentic educational setting." *Scientific Reports* (2025). 194 students, Harvard introductory physics, Fall 2023.
2. Bastani, H., et al. "Generative AI without guardrails can harm learning: Evidence from high school mathematics." *Proceedings of the National Academy of Sciences* 122 (2025). Nearly 1,000 students; GPT Base +48% practice, -17% exam; GPT Tutor no significant exam effect.
3. UNESCO. *Guidance for Generative AI in Education and Research* (September 2023). Recommends a human-centred approach and a minimum age of 13 for classroom use of AI tools.
4. UNESCO. *AI Competency Framework for Students and AI Competency Framework for Teachers* (September 2024).
5. Office of the Privacy Commissioner of Canada and provincial/territorial privacy authorities. *Principles for Responsible, Trustworthy and Privacy-Protective Generative AI Technologies* (7 December 2023).
6. OECD. *Recommendation of the Council on Artificial Intelligence* (OECD AI Principles), adopted 2019, updated May 2024.
7. Council of Europe. *Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law* (CETS No. 225, 2024). Signed by Canada, February 2025.
8. Government of Canada. *AI for All: Canada's National Artificial Intelligence Strategy* (June 2026).
9. Ontario. *Strengthening Cyber Security and Building Trust in the Public Sector Act, 2024*, S.O. 2024, c. 24 (AI provisions apply to prescribed public-sector entities as regulations are made).
10. National Institute of Standards and Technology. *AI Risk Management Framework (AI RMF 1.0)* (January 2023). Functions: Govern, Map, Measure, Manage.
11. ISO/IEC 42001:2023, *Information technology — Artificial intelligence — Management system*.
12. Regulation (EU) 2024/1689 (AI Act), Article 5(1)(f) and Annex III; as amended by Regulation (EU) 2026/1744 (Digital Omnibus on AI), published 24 July 2026, deferring Annex III high-risk obligations to 2 December 2027.
13. Kirkpatrick, D. L. & Kirkpatrick, J. D. *Evaluating Training Programs: The Four Levels* (3rd ed., 2006).

Notes

This guide is general information, not legal or compliance advice. Regulation of AI is changing quickly; check current requirements in your jurisdiction. Momentum roadmap items describe direction and are not commitments to specific features or dates. Third-party product names (DeepL, OpenAI) are trademarks of their respective owners.

Also in this series

MOMENTUM INSIGHTS • VOL. 01

Beyond the Feature Checklist

Why the learning platforms worth recommending sell a result, not a tool.

MOMENTUM INSIGHTS • VOL. 02

The Residual Gap

Why instructional designers are paid once for work that earns for years.